The Classification of Coherence Relations and their Linguistic Markers: An Exploration of Two Languages

Alistair Knott

Human Communication Research Centre, University of Edinburgh, 2 Buccleuch Place, Edinburgh EH8 9LW, Scotland

Ted Sanders

Utrecht Institute of Linguistics OTS, University of Utrecht, Trans 10, NL–3512 JK Utrecht, The Netherlands

It has become popular among discourse linguists to explain a text's coherence by identifying 'coherence relations' which apply at various levels between its component spans. However, there is currently no overall agreement about how to define a standard set of coherence relations, and even about what the coherence relations in a text are intended to represent. In this paper, both questions are addressed: we outline a conception of relations as modelling psychological constructs used by readers and writers, and suggest how a limited set of categories of coherence relations can be identified.

We relate two independent methods for investigating relations, one drawing mainly on psycholinguistic experiments on Dutch speaking subjects, the other starting from a study of the 'cue phrases' used to signal relations in English text. Both approaches lead to classifications of relations and cue phrases. We examine to what extent these classifications converge—and to what extent they accord with the psychologically motivated classification—in a comparative study of a set of cue phrases in English and Dutch. Interesting similarities are noted on both counts.

1 Introduction: Coherence Relations

The starting point for any theory of discourse is the oft-noted observation that what we call 'a text' is more than just a collection of random sentences. Consider the following two passages:

(1)   Tim must love that Belgian beer. The crate in the hall is already half empty.

(2)   Tim must love that Belgian beer. He's six foot tall.

While passage 1 is easy for a reader to understand, passage 2 poses problems. It is hard to find a connection between its two sentences, and hard to see why they have been placed next to one another as if they *are* somehow connected. The notion of 'coherence' is often invoked to distinguish between texts such as these—although our intuitions about coherence are not always categorical, readers are likely to agree that text 1 is coherent while text 2 is not.

In developing an account of what makes a text[1] coherent, an influential suggestion has been that a fixed set of methods exists for linking one portion of text to another. For instance, one span of text can elaborate on the span which precedes it, or give some justification for it, or express a conclusion which follows from it. The idea is that a choice amongst a finite set of alternatives must be made when juxtaposing two portions of text, and hence that a finite set of *coherence relations* will be sufficient to enable an analysis of every coherent text. The coherence of text 1 might, for instance, be attributed to the presence of an EVIDENCE relation between its two sentences, as in Figure 1.

EVIDENCE

Tim must love that Belgian beer.   The crate in the hall is already half empty.

Fig. 1. An Analysis of Text 1

A number of recent theories have made use of relations in explaining coherence (e.g. Longacre 1983, Hobbs 1985, Grosz and Sidner 1986, Mann and Thompson 1988, Martin 1992). However, while the general idea of coherence relations is attractive, no consensus has emerged about a single set of relations; as is well documented by Hovy (1990), the alternative sets that have been put forward are very different from one another. The number of relations proposed varies from 2 (Grosz and Sidner) to over 100 (Martin); they are defined variously in terms of semantics (Longacre, Hobbs), intentions (Grosz and Sidner), or a

---

[1] In this paper, we confine our attention to monologues, and primarily (but not exclusively) written ones.

combination of both (Mann and Thompson); and a number of different top-level organising principles are employed in their classification. As a result of the many different proposals, the question of 'which set of relations to choose' has recently been the subject of an ongoing debate. The present paper can be seen as a contribution to this debate.

We begin from a hypothesis about the nature of coherence relations—namely that they should be thought of in psychological terms, as a set of conceptual relations used by readers and writers when processing text. This conception of relations is sketched in Section 2. The rest of the paper sets out two proposals for investigating relations thus conceived: the cognitive approach to coherence relations pursued by Sanders, Spooren and Noordman (1992, 1993), and the linguistic study of cue phrases, which can be defined as the set of lexical signals which make coherence relations explicit in surface text, including connectives, conjunctions, and subordinators, by Knott and Dale (1994). The common ground between the proposals is the suggestion that evidence for psychological text-structuring mechanisms can be obtained via a study of the linguistic devices people use to signal relations in surface text. In text 1, for instance, the writer could have used the connectives *since* or *because* to express the EVIDENCE relation, but never the connectives *however* or *moreover*. In both proposals, such restrictions on the use of connectives are used as a source of evidence about the underlying coherence relations.

Even though they are both concerned with relations and cue phrases, the two proposals discussed below have different starting points, and rely on different arguments. In the first, outlined in Section 3.1, it is argued that the set of coherence relations should be organized, and it is proposed to do this in terms of a set of cognitive primitives, such as polarity (positive or negative) which are common to all relations. Arguments for the saliency of these primitives are provided from several sources; in particular from experimental studies. The second approach, outlined in Section 3.2, begins by arguing in principle that there are likely to be strong correspondences between the set of cognitive relations used by readers and writers and the set of linguistic devices for signalling them, and uses this argument to motivate a purely linguistic study of the substitution relationships between connective phrases.

In Section 4, some interesting similarities are noted between the results of the two approaches. For one thing, both studies suggest a decomposition of coherence relations into a number of independent parameters. What is more, there are some indications that the two studies motivate similar sets of parameters; a fact which is particularly interesting considering that the experiments in Section 3.1 are on Dutch subjects while the substitution study in Section 3.2 is of English connectives. The main aim of this paper is to examine more closely whether the two approaches can indeed be said to converge. To this end, small substitution studies have been carried out on a similar set of phrases in both

English and Dutch, and the parameters which result from the two studies have been compared in detail. Sections 4.2.2 and 5 describe the two studies and their preliminaries; Section 6 contains a discussion of the similarities and differences which are found, and of some of the potential problems with the substitution methodology which are highlighted.

2   A Psychological Conception of Relations

Some debate surrounds the issue of the cognitive relevance of coherence relations. It is sometimes claimed (e.g. Grosz and Sidner, 1986) that readers of a text do not need to determine the coherence relations in a text in order to understand it, and that relations are useful principally as 'analytic tools' for describing text structure. There are also proposals (e.g. Hobbs, 1979, 1983, 1990; Mann & Thompson, 1986; Sanders *et al.*, 1992, 1993) in which it *is* hypothesized that coherence relations should be considered as cognitive entities.

The hypothesis we would like to explore is that coherence relations can be thought of as modelling cognitive mechanisms operative in readers and writers when they process text. According to this view, when a particular relation is posited between two spans of text, a claim is being made about the mechanism used by the writer to join these two spans together, and about the mechanism used by its readers to interpret them.

As regards discourse interpretation, the psychological claim can be elaborated as follows. When readers process a text, they construct a representation of the information it contains. A crucial property of this cognitive representation is that it integrates the individual propositions expressed in the text into a larger whole, just as our representation of an individual situation in the world draws together a number of separately perceived events and states (see, for instance, Noordman and Vonk, 1993 and several contributions to Gernsbacher and Givón, 1995). The psychological claim about coherence relations is that they should be seen as modelling the different ways in which this integration between propositions can occur. According to this idea, determining the coherence relations in a text is part of the process of understanding it.

Is there an empirical way of testing this claim? There are at least two hypotheses that could be investigated (Sanders, 1992). First, if coherence relations affect the reading process, the processing of the information should depend on the type of coherence relation (e.g. causal versus additive or contrastive). Second, one may assume that the linguistic marking of relations will influence processing as well. And indeed, there is some evidence for both hypotheses in the literature.

4

Several experimental studies suggest a processing difference between causal and non-causal relations. For instance, causally related events in short narratives are recalled better (Black and Bern, 1981; Trabasso and Van den Broek, 1985; Trabasso and Sperry, 1985). Also, causally related sentences are read faster (Haberlandt and Bingham, 1978).

Also, there are many studies that support the idea that readers make use of connectives in processing extended text. For instance, Haberlandt (1982) shows that reading time for a sentence is speeded if it is preceded by a linguistic marker which makes the relation with the preceding text explicit. Segal *et al.* (1991) show that the presence of connectives in a narrative text facilitates the categorisation of its relations.

Sanders (1992, chapter 4), and Sanders and Noordman (submitted for publication) have conducted reading experiments in which both the type of coherence relation between two segments (Problem-Solution versus List) and the linguistic marking of these relations (implicit versus explicit) was investigated. Results show that both factors influence text processing; i.e. segments connected by Problem-Solution relations are processed faster and recalled better, and signalling leads to the faster processing of the directly following text segment.

Finally, there is evidence of priming effects between relations. Traxler *et al.* (1996) have found, for instance, that reading time for a two-sentence text containing an inferential relation is improved by prior presentation of an inferential relation; likewise, that reading time for a text containing a causal relation is improved by prior presentation of a causal relation.

The claim for the psychological basis of coherence relations as it applies to text production is less well substantiated. It is certainly possible to conceive of relations as having a dynamic role in text production; they are frequently employed as planning operators in natural language generation systems (see e.g. Hovy 1993, Moore and Paris 1993). But whether they play a similar role in the human text production mechanism has yet to be investigated. There are general indications that planning operators are involved in multisentential text production (see for instance Flower and Hayes, 1980). There is also a certain amount of evidence that connectives are useful in the production of text: for instance, in a study of children's writing, Bereiter and Scardamalia (1987: p62) found that prompting children with phrases such as *even though*, *also* and *for example* helped them to generate relevant content in their compositions. But much more work is needed to flesh out these claims.

While all of the above studies provide some support for the notion of psychologically real relations, they certainly do not count as conclusive evidence for this notion. For instance, there is a measure of circularity in using a study of the effect of sentence/clause connectives on reading as evidence for cog-

nitive relations. To do so requires the assumption that connectives are good experimental indicators of cognitive relations, and this assumption virtually takes for granted that such cognitive structures actually exist. In fact, there are many possible explanations as to why connectives facilitate the processing of text. They might do so only indirectly, by making some 'non-relational' structure such as a schema easier to access; or they may only do so in the sense that *any* additional information about the text is likely to facilitate its interpretation. In summary, the issue of the psychological reality of coherence relations is far from settled yet.

In this paper we want to explore the psychological basis of coherence relations from a different angle. We will present two independent methods for investigating relations, each of which adopts a cognitive conception of relations as a working assumption, and consider to what extent the methods converge on the same classification of relations. Though both methods involve attention to connectives, they make use of them in quite different ways; moreover, the two methods have been applied to connectives in different languages. Any convergence between the results obtained would therefore be interesting evidence for a cognitive conception of relations.

3 Investigating Coherence Relations: Two Alternative Approaches

In this section, we outline two independent methods for pursuing the claim that relations model cognitive constructs, and note some interesting similarities between their findings.

3.1 A Relational Account: Classification in terms of cognitive Primitives

Sanders, Spooren and Noordman (1992, 1993) have argued that many existing accounts of coherence relations do not allow for a cognitive interpretation because they do not allow for plausible hypotheses about how language users construct a cognitive representation. For instance, how should a reader arrive at the interpretation of a particular relation, such as EVIDENCE? If each relation is thought of as a separate cognitive primitive, we must assume that in order to interpret a stretch of discourse, readers use their instant knowledge of all these relations (30? 100? 1000?), to determine its structure. It is far more attractive, they argued, to assume that readers understand a piece of discourse because a notion like EVIDENCE is composite. It consists of more elementary notions, such as causality, and readers make use of their knowledge of this limited set of basic notions to derive the appropriate coherence relation.

Four of such basic notions were hypothesized, common to all coherence relations, each of which can take two alternative values. These are as follows:

- BASIC OPERATION: every relation is deemed to have either a CAUSAL or an ADDITIVE component. CAUSAL relations are those where a 'relevant' causal connection exists between the spans; all other relations are ADDITIVE.
- SOURCE OF COHERENCE: every relation is coherent on SEMANTIC or PRAGMATIC grounds. It is semantic if the spans are related in terms of their propositional content and pragmatic if they are related because of their illocutionary force. [2]
- POLARITY: a relation is POSITIVE if its basic operation links the content of the two spans as they stand, and NEGATIVE if it links the content of one of the spans to the negation of the content of the other span. Negative polarity relations typically involve either a violation of expectation, where the expectation derives from a causal basic relation; or a contrast, where the basic relation is additive.
- ORDER OF SEGMENTS: this distinction only applies to CAUSAL relations; they are deemed to have BASIC order if the antecedent is on the left, and NON-BASIC order if it is on the right.

The four cognitive primitives are combined to generate classes of coherence relations. The combination of the four primitives results in a classification scheme in which 12 classes of relations are characterized. In this way, the set of relations can be organized in terms of their own 'meaning characteristics': if causality is found in a certain relation it is classified as belonging to a different group than when a relation is not causal. For instance, the relation CONSEQUENCE-CAUSE is defined as BASIC OPERATION = CAUSAL, SOURCE OF COHERENCE = SEMANTIC, ORDER = NON-BASIC, POLARITY = POSITIVE. An example of this relation is given in text 3:

(3)   Last week it rained a lot in Scotland, because there was low pressure over Ireland.

Text 4 is an example of a 'list' relation, defined as SEMANTIC, POSITIVE and ADDITIVE:

(4)   Last week it rained in Scotland. The weather was bad in the Netherlands as well.

A NEGATIVE POLARITY relation is presented in text 5. It is defined as NEGATIVE, ADDITIVE and SEMANTIC.

---

[2] The term EPISTEMIC (Sweetser, 1990) can also be used to describe PRAGMATIC relations; the difference is that Sweetser also differentiates between SPEECH ACT and EPISTEMIC relations, two relation types both incorporated in the category of PRAGMATIC relations (see Sanders, 1997).

(5)   Last week, the weather was bad in Scotland, whereas in the Netherlands the sun was shining.

Finally, a CLAIM-ARGUMENT relation like that in text 1 would be categorized as CAUSAL, POSITIVE and PRAGMATIC, with BASIC order.

The four primitives are important cognitive categories, prominent in research on language and language behavior. For the distinction between causals and non-causals, for instance, we have already mentioned several studies on discourse processing above. The primitive of polarity is also a well-known factor in psycholinguistic literature: for instance, negative polarity utterances are processed more slowly than their positive counterparts (Wason and Johnson-Laird, 1972; H. Clark, 1974).

Support for the saliency of these four primitives was provided by a number of experiments. In one experiment (Sanders *et al.*, 1993) trained discourse analysts were explicitly asked to make direct comparisons between different relations. In general, the similarity judgments conformed to the categorizing principles. In another experiment (Sanders *et al.*, 1992) discourse analysts were given definitions and examples of all 12 relations, and were asked to decide which relations were appropriate in a number of sample texts. In a third experiment, 'naive' Dutch-speaking subjects were used, who did not know about the relation definitions. They were shown sample texts without explicit connectives, and their task was simply to decide which connective word was most suitable. The researchers had identified prototypical markers for each relation—for instance, *because* is a prototypical marker of CAUSE-CONSEQUENCE—and these markers provided an experimental window on the relations being used by the subjects.

The experiments were designed to test how much agreement there was between subjects about the relations in a text. In all cases, a fair amount of agreement was found. But the most revealing results come from an analysis of the cases where subjects disagreed. In each experiment, it was found that where there was disagreement over which relation to use, it tended to be over the value of a single parameter only, rather than more than one. This result provides good support for the independence of the decisions about the different parameters.

3.2   A Language-Based Account: from Cue Phrases to Categories of Relations

A second approach to investigating relations focuses not so much on the relations themselves, nor on the disagreements between subjects about the relations present in a text, but rather on a study of the linguistic devices that are used to signal relations explicitly. Sanders *et al.* (1992) noted a number of correspondences between connectives and the features they signal; for instance,

but typically signals NEGATIVE POLARITY relations, while *and* typically signals ADDITIVE relations. There is no one-to-one mapping between relations and connectives (see also Mann and Thompson, 1986, 1988). However, these restrictions imply an organization of the relations connectives can express; they do not just co-exist as a set of relations on one and the same level. Sweetser (1990) has also argued that connectives differ with respect to the 'domains' in which they can be used to express meaning; for instance, *since* is hard to interpret in the 'content' domain (similar to Sanders *et al*'s SEMANTIC domain).

This linguistic argument is developed at a finer level of granularity by Knott and Dale (1994). Their central idea is that a classification of the set of connectives found in a given language can be used in its own right as evidence about the relations used by speakers of that language.

The argument rests on a conception of relations as constructs which are actually *communicated*, via a text, from writers to readers. It adopts as a working assumption the proposal outlined in Section 2, that readers make use of the constructs modelled by relations when they interpret a text. A consequence of this scenario, it is argued, is that there is likely to exist a method for *making explicit* each of these constructs in surface text. Clearly, many of the relations in a given text are easily inferrable by the reader, and so need not be made explicit; but it is unlikely that any relation exists which is easily inferrable in *any* text where it is used. (The 'inferrability' of a relation is not a characteristic of the relation itself, but of the context in which it appears, and the amount of background knowledge posessed by the reader.) Hence it would be beneficial if a language contained ways for a writer to identify each relation individually. Given that language is a flexible communicative tool, well adapted to the needs of its users, we can hypothesise that a study of the methods of signalling relations in a language can give us direct evidence about the strategies which speakers of that language actually use.

An obvious place to begin such a study is by looking at the set of cue phrases. Knott and Dale 1994 formulate a definition of this category of phrases, which encompasses sentence and clause conjunctions, subordinators, and a wide variety of sentential adverbials. [3] They then classify this set using a functional metric, into groups of phrases which can be used interchangeably in a given discourse. A 'test for substitutability' is proposed to compare two cue phrases X and Y. The test calls for the judgements of a writer, who is given a text in which the cue phrase X appears. The writer is to imagine that X must

---

[3] The definition is based on a test. The phrase under consideration is located in a piece of naturally-occurring text, and isolated in the context of its host clause. If this isolated clause is uninterpretable without its surrounding context, but *becomes* interpretable if the phrase is removed, it is taken to be a cue phrase. See Knott and Dale (1994) for a more detailed formulation of the test.

be replaced by another cue phrase (perhaps for reasons of elegant variation, X having already been used in the preceding text). Can Y (the 'candidate phrase') be used in its place? Such decisions must often be made during the course of normal writing; the test is thus oriented towards the judgements of ordinary writers, rather than calling exclusively on the theoretical intuitions of a discourse analyst.

An example is given in text 6.

$$(6) \qquad \text{It was a hot day,} \left\{ \begin{array}{l} so \\ \checkmark\, therefore \\ \#\ however \end{array} \right\} \begin{array}{l} \text{they ate outside on} \\ \text{the patio.} \end{array}$$

In this notation, the original cue phrase is *so*; *therefore* is represented as substitutable for *so* in this context; and *however* is represented as not substitutable for *so*. [4]

If we generalise over all contexts, there are four possible substitutability relationships between two cue phrases X and Y.

– X is synonymous with Y if in any context where one can be used, the other can also be used.
– X and Y are exclusive if they can never be substituted for one another in any context.
– X is a hypernym of Y—and Y is a hyponym of X—if whenever Y can be used, so can X; but there are some contexts where X can be used and Y cannot.
– X and Y are contingently substitutable if there are some contexts where they can be substituted, other contexts where X can be used and not Y, and still other contexts where Y can be used and not X.

3.2.1 A Taxonomy of Cue Phrases

Starting from the hypothesis that substitutability relationships can be ordered systematically, the relationships between cue phrases can be represented in diagrams, as shown in Figure 2. Using this notation, a hierarchical taxonomy of cue phrases can be constructed, making use of inheritance to represent the substitutability relationships between any set of cue phrases. For example, consider the taxonomy in Figure 3. This indicates (among other things) that B and C are contingently substitutable; that F is a hyponym of A; and that F and E are exclusive. Note that the taxonomy represents in a systematic

---

[4] Note that the hash sign does not necessarily signal ungrammaticality or incoherence in this notation—it merely indicates that a writer would not be prepared to substitute the 'hashed' cue phrase for the original one.
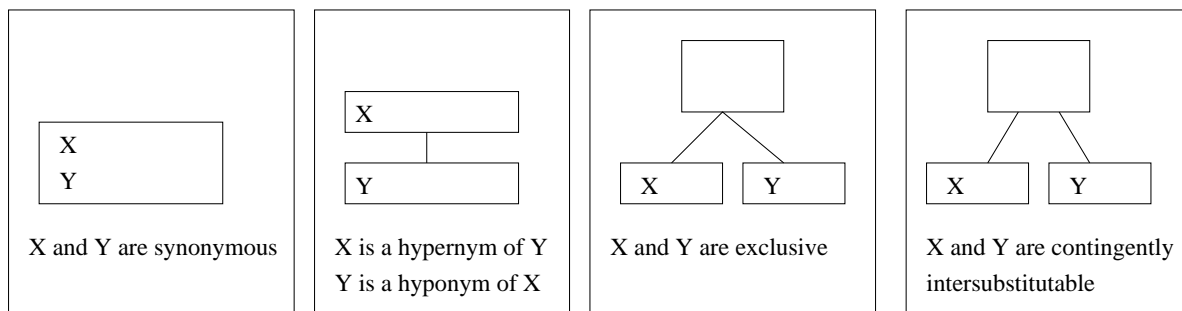
Fig. 2. Diagramatic Representation of Substitutability Relationships

Fig. 3. A Taxonomy of Substitutability Relationships

way the relationship between every cue phrase and every other cue phrase; it is thus an extremely rich source of information.

A large taxonomy of around 150 English cue phrases has been created (see Knott, 1996). Table 1 gives an idea of the range of phrases which figure in the taxonomy; and a simple extract from the taxonomy itself is given in Figure 4. Some motivating examples are given in Texts 7 to 9:

admittedly	if	in case	as soon as
in that	all in all	at any rate	until
although	alternatively	rather	because
and	on the contrary	then	assuming that
just as	incidentally	thereafter	either
for one thing	on the other hand	for example	consequently
so	indeed	besides	lastly

Table 1
A Selection of Cue Phrases from Knott's 1996 Taxonomy

Fig. 4. A Simple Extract from Knott's 1996 Taxonomy of Cue Phrases

(7) Sally and John are like chalk and cheese. Sally loves reading, $\left\{\begin{array}{l} whereas \\ \checkmark\, but \\ \checkmark\, and \end{array}\right\}$ John hates it.

(8) Sally was sick, $\left\{\begin{array}{l} but \\ \#\ whereas \\ \#\ and \end{array}\right\}$ she went to work as usual.

(9) I don't like John. He's always gloomy $\left\{\begin{array}{l} and \\ \#\ but \\ \#\ whereas \end{array}\right\}$ he never has anything interesting to say.

3.2.2 A Feature-Theoretic Interpretation of Substitutability Relationships

One thing to note immediately about the taxonomy of cue phrases is its hierarchical nature—it contains a number of 'hypernymic' phrases such as *and* and *but*. This finding might seem somewhat at odds with the original argument for the cue phrase methodology, that cue phrases are means for a writer to signal individual relations explicitly. If this is the case, why should ambiguous cue phrases like *and* exist at all?

However, this apparent problem disappears if a slightly more sophisticated conception of cue phrases is adopted. It should be noted that even very general cue phrases like *and* are not substitutable for all other cue phrases. For instance, *and* would never be substituted for *when*:

(10) Martin was six years old $\left\{\begin{array}{l} when \\ \#\ and \end{array}\right\}$ he found out where babies come from.

In other words, *and* does tell the reader *something* about the relation being used. In fact, it is natural to think of the cue phrase as signalling *some features* of a relation, while leaving other features undefined. Indeed, each of the substitutability relationships in the taxonomy readily admits of a feature-theoretic interpretation:

- If X is synonymous with Y, then they signal identical features.
- If X is exclusive with Y, then they signal *different values* of some feature.
- If X is a hyponym of Y, then X signals all the features that Y signals, and *some other* feature in addition, for which Y is undefined.
- If X and Y are contingently substitutable, then X and Y signal some of the same features, but in addition X is defined for a feature for which Y is undefined, and Y is defined for a feature for which X is undefined.

Eight features have so far been motivated from the taxonomy: for instance, PRESUPPOSITIONALITY (used, among other things, to distinguish subordinators and co-ordinators) and MODAL STATUS (used to indicate whether the relation holds between hypothetical or actual propositions). For precise definitions, see Knott (1996), and Knott and Mellish (1996). But for present purposes, the interesting point to note is that a number of these features resemble parameters proposed by Sanders *et al.*: in particular, there seem to be close analogues to the BASIC OPERATION, SOURCE OF COHERENCE and POLARITY parameters. These similarities will be further discussed in the next section.

4 A Comparison Between the Two Methodologies

The studies in Sections 3.1 and 3.2 provide two quite separate motivations for a decomposition of relations into orthogonal features. For Sanders *et al.*, the decomposition is initially advanced as a psychologically plausible hypothesis, and subsequently supported by empirical experiments. For Knott and Dale, it emerges from a purely linguistic study, itself legitimised by means of an argument that the relations people use are likely to be reflected in the set of cue phrases available to mark them. Given that both studies begin from a conception of relations as psychological constructs, this convergence is an interesting result.

Moreover, as just mentioned, there are some apparent similarities in the individual features which are motivated by the two studies. To take a simple example: we have already seen that Sanders *et al.* identify 'prototypical markers' for relations in one of their experiments. Thus *although* is 'prototypical' for NEGATIVE POLARITY CAUSAL SEMANTIC relations, and *because* for POSITIVE POLARITY CAUSAL SEMANTIC relations. *Although* and *because* are exclusive cue phrases; according to Knott and Dale's methodology, this alone is evidence that they differ over the value of one or more features.

It would certainly be a useful result if there were similarities between the definitions reached in the two accounts. On the whole, each approach has its weak points: for instance, Sanders *et al*'s choice about which parameters to

test for could be queried; so can Knott and Dale's argument for using cue phrases as evidence for cognitive constructs. However, a convergence between the results obtained by the two methods would go some way towards meeting these objections. For instance, the choice of parameter definitions is much more constrained in the substitution methodology than in Sanders *et al*'s approach, and so the parameter definitions in Knott and Dale would lend support to those proposed by Sanders *et al*. On the other hand, Sanders *et al*'s experimental findings would support Knott and Dale's decision to use cue phrases as a source of evidence.

However, to find out whether Sanders *et al*'s parameters are most appropriate for describing the relationships between prototypical markers, we need a more systematic study of the relationships between cue phrases. For instance, is the distinction between POSITIVE and NEGATIVE polarity useful in other parts of the taxonomy as well? Furthermore, we are not just interested in the prototypical markers of relations; in Sanders *et al*'s account, these are not always sufficient to identify relations explicitly. We are also interested in classifying the most specific cue phrases, as well as the very general ones, in terms of the hypothesised set of features.

An objection can be raised at this point. It might be asked, if Sanders *et al.* use cue phrases as experimental indicators of the relations 'used by writers', should we really be surprised that a classification of cue phrases yields similar parameters? Even if the correspondences between cue phrases and relations in Sanders *et al* are only approximate, it would surely not be a complete coincidence if the two studies converged on a single set of features.

However, this objection misses two important points. Firstly, it must be reiterated that the two studies investigate cue phrases in different languages. Sanders *et al*'s experiments are on Dutch subjects, while Knott and Dale's taxonomy is of English cue phrases; so if the same features emerged from each study, it would at least indicate some striking similarities between the mechanisms available for signalling relations in the two languages. Secondly, even if the two studies were of the same language, the objection does not hold—the studies are based on completely different classes of data. Motivating parameters in the taxonomy depends on finding hierarchical relationships between different cue phrases. Sanders *et al*'s parameters are motivated from an analysis of the disagreements between subjects over which cue phrase to use in given contexts; this analysis could be carried out even if no hierarchical relationships existed between cue phrases at all.

In this paper, we will adopt the *prima facie* similarities discussed in the previous section as the starting point for a more detailed investigation of the parameters which emerge from a taxonomy of cue phrases. As a first test of the viability of such an approach we will explore cue phrases from two languages: English and Dutch.

The main research question to be addressed is: to what extent can the relationships between cue phrases be described in the same terms that were used by Sanders *et al.* to describe coherence relations? As a question to be addressed by the substitution methodology, this question can be broken down into two components:

  (i) Do the English and Dutch taxonomies of cue phrases resemble each other?
  (ii) Can either taxonomy be described using Sanders *et al*'s parameters?

Both of these questions have a bearing on a more fundamental question, of whether relations model cognitively primitive mechanisms in readers and writers. Question 1 bears directly on it, in that structural similarities between the Dutch and English taxonomies could conceivably point to the operation of linguistic universals. However, it should be borne in mind that Dutch and English are historically close languages.

Question 2 is in fact more important, although its bearing on the issue of psychological reality is indirect. If the same set of features emerges from the taxonomies as emerges from Sanders *et al*'s experiments, then the convergence of the two methodologies is strong support for the existence of the psychological constructs they both claim to be investigating.

4.2   Research Method

To answer the research questions we will start from Sanders *et al*'s primitives and investigate whether the use and meaning of cue phrases can be described in these terms. To that end, we will make use of the substitution methodology and apply this to English and Dutch cue phrases.

4.2.1   Choosing a Set of Connectives to Investigate

A first task is to decide on a group of connectives to study. In this first exploratory study, we have opted for a small group of phrases, whose substitutability relationships are likely to have a bearing on the primitives hypoth-

esised by Sanders *et al*. We will be focussing on obvious candidate areas of the cue phrase lexicons of the two languages, where we expect that clear substitution data can be obtained. In each area we will test to what extent the primitives can account for the occurence of, and differences between certain cue phrases.

We will by no means describe the total set of cue phrases for these two languages; this is a task far beyond the scope of the present paper. This enterprise should be regarded as a first theory-driven systematic and cross-linguistic cue phrase study. Any results obtained should be taken as strong hypotheses which could subsequently be tested for other, less related languages, and with other methods; see Section 6 for some discussion of these points.

4.2.2 *Some Modifications to the Substitutability Test*

Before looking again at the taxonomies, the test for substitutability must first be outlined in a little more detail. The test we will be using in what follows is not quite as strict as that presented in Section 3.2; some differences between cue phrases will be ignored.

For one thing, the position of the candidate phrase in the original text is not constrained—it can appear in the same place as the original cue phrase, or in a different place. For instance, consider text 11:

(11) $\qquad$ Jane liked sailing boats. John, $\left\{ \begin{array}{l} \textit{however,} \\ \checkmark\textit{but} \leftarrow \end{array} \right\}$ did not.

The cue phrase *but* cannot be inserted into the text at the same point as *however*. But if it is placed at the beginning of the second clause, it is substitutable. In cases where the candidate phrase needs to be moved in this way, an arrow will indicate the direction of movement.

Cue phrases can even be moved from one of the related clauses to the other, as in text 12:

(12) $\qquad$ $\left\{ \begin{array}{l} \textit{Because} \\ \checkmark\textit{so} \hookrightarrow \end{array} \right\}$ Jane liked sailing boats, she took a job with a charter company.

Such phrases will be termed 'swap-substitutable', and will be identified by a hooked arrow in the direction of the clause in which they must appear.[5]

While the candidate phrase can be swapped from one clause to another in the test, the order of the clauses themselves cannot be changed. Thus in text 13,

---

[5] The notion of swap-substitutability is a departure from Knott and Dale (1994).

so cannot be substituted for *because*:

(13) Jane took a job with a charter company, $\left\{ \begin{array}{c} because \\ \# \ so \end{array} \right\}$ she liked sailing boats.

The hash sign now denotes a cue phrase that cannot be substituted for the original phrase in any position in either clause.

Another relaxed constraint on substitutability is the amount of background knowledge that a reader can be assumed to have about the context being considered. This is to permit very general cue phrases to be substituted for more specific ones and vice versa. For instance, a writer might not always be prepared to substitute *and* for *whereas*—the former cue phrase conveys less information, and might be confusing for a reader who could not easily infer that a contrast was being signalled. So in the test, the writer is allowed to assume that the reader knows as much or as little as is appropriate for the candidate phrase.

A final modification is needed to allow the substitutability test to be applied in Dutch, where the position of the verb is dependent on the choice of cue phrase. If need be, the position of the verb in the original text can be altered so as to be appropriate for the candidate phrase.

These modifications make the substitutability test a little more complex. However, they are well motivated; and the modified test can still be hypothesized to bear a close resemblance to a task which ordinary writers are faced with.

5 A Cross-Linguistic Investigation of Sanders *et al*'s Primitives using the Substitutability Test

This section is divided into a number of subsections, each relating to a different class of cue phrases, and a different portion of the taxonomy. The aim is to investigate whether the distinctions found in Sanders *et al.* are also found using the substitution methodology in English and/or Dutch.

5.1 Some Preliminary Categorisations

5.1.1 The English Taxonomy

We can begin by considering a set of exclusive cue phrases, illustrated in Figure 5. Motivating examples for this diagram are as follows:

Fig. 5. Four Exclusive English Cue Phrases

(14)   John was sad, $\left\{\begin{array}{l} \textit{because} \\ \# \ \textit{also} \\ \# \ \textit{nevertheless} \\ \# \ \textit{whereas} \end{array}\right\}$ his toy was broken.

(15)   John likes mussels. $\left\{\begin{array}{l} \textit{also} \\ \# \ \textit{because} \\ \# \ \textit{nevertheless} \\ \# \ \textit{whereas} \end{array}\right\}$ likes them.
Bill

(16)   It was John's $\left\{\begin{array}{l} \textit{nevertheless} \\ \# \ \textit{also} \\ \# \ \textit{because} \\ \# \ \textit{whereas} \end{array}\right\}$ he was sad.
birthday;

(17)   Jim's only $\left\{\begin{array}{l} \textit{whereas} \\ \# \ \textit{also} \\ \# \ \textit{because} \\ \# \ \textit{nevertheless} \end{array}\right\}$ Jane's only
interested in   interested in rugby.
cricket;

There is potential for the exclusivity relationships between the four cue phrases in Figure 5 to be described using two of Sanders *et al*'s parameters: POLARITY and BASIC OPERATION. *Because* and *nevertheless* both seem CAUSAL in nature, while *also* and *whereas* both seem ADDITIVE. The POLARITY parameter makes a different cut through the phrases—*also* and *because* are both POSITIVE, while *nevertheless* and *whereas* are NEGATIVE.

Confirmation for this analysis is provided by two hypernymic cue phrases: *but* and *and*. As shown in Figure 6, *and* is a hypernym of both *also* and *whereas*, while being exclusive with both *because* and *nevertheless*. Note that the empty box in the diagram is no more than a notational convenience; it captures the generalisation that *and* is exclusive with both of the cue phrases in its daughter categories.

Motivating examples appear below:

(18)   Jim's only $\left\{\begin{array}{l} \textit{whereas} \\ \checkmark \ \textit{and} \\ \# \ \textit{also} \\ \# \ \textit{because} \\ \# \ \textit{nevertheless} \end{array}\right\}$ Jane's only
interested in   interested in rugby.
cricket;

18

ADDITIVE CAUSAL

```
        ┌──────┐
        └──────┘
    ADDITIVE      CAUSAL
    ┌──────┐    ┌──────┐
    │ and  │    │      │
    └──────┘    └──────┘
+VE POLARITY -VE POLARITY  +VE POLARITY -VE POLARITY
 ┌──────┐   ┌──────┐    ┌──────┐   ┌──────────┐
 │ also │   │whereas│   │because│  │nevertheless│
 └──────┘   └──────┘    └──────┘   └──────────┘
```

Fig. 6. *And* as a Marker of ADDITIVE Relations

(19) John likes mussels. Bill $\left\{\begin{array}{l} also \\ \checkmark\, and\ \leftarrow \\ \#\ because \\ \#\ nevertheless \\ \#\ whereas \end{array}\right\}$ likes them.

And can thus be thought of as signalling only the ADDITIVE feature. In a similar fashion, *but* appears defined only for NEGATIVE POLARITY, as Figure 7 illustrates.

```
            ┌──────┐
            └──────┘
   -VE POLARITY      +VE POLARITY
    ┌──────┐          ┌──────┐
    │ but  │          │      │
    └──────┘          └──────┘
 ADDITIVE   CAUSAL   ADDITIVE   CAUSAL
 ┌──────┐ ┌──────────┐ ┌────┐ ┌────────┐
 │whereas│ │nevertheless│ │also│ │because │
 └──────┘ └──────────┘ └────┘ └────────┘
```

Fig. 7. *But* as a Marker of NEGATIVE POLARITY Relations

Motivating examples for this diagram are as follows:

(20) Jim's only interested in cricket; $\left\{\begin{array}{l} whereas \\ \checkmark\, but \\ \#\ also \\ \#\ because \\ \#\ nevertheless \end{array}\right\}$ Jane's only interested in rugby.

(21) It was John's birthday; $\left\{\begin{array}{l} nevertheless \\ \checkmark\, but \\ \#\ also \\ \#\ because \\ \#\ whereas \end{array}\right\}$ he was sad.

5.1.2 The Dutch Taxonomy

The four cue phrases we have used to introduce the Dutch taxonomy are presented in Figure 8. Wherever possible, we have used direct translations of

Fig. 8. Four Exclusive Dutch Cue Phrases

the English examples to motivate the Dutch diagrams. Motivating examples for this first diagram are given below, along with translations:[6]

(22) Jan was bedroefd, $\left\{\begin{array}{l} omdat \\ \# \ ook \\ \# \ niettemin \\ \# \ daarentegen, \end{array}\right\}$ zijn speeltje gebroken was.
(John was sad, *his toy was broken.)*

(23) Jan houdt van mosselen. $\left\{\begin{array}{l} Ook \\ \# \ omdat \\ \# \ niettemin \\ \# \ daarentegen \end{array}\right\}$ Bill houdt ervan.
(John likes mussels. *Bill likes them.)*

(24) Het was Jan's verjaardag; $\left\{\begin{array}{l} niettemin \\ \# \ ook \\ \# \ omdat \\ \# \ daarentegen \end{array}\right\}$ was hij bedroefd.
(It was John's birthday; *he was sad.)*

(25) Jim heeft alleen maar belangstelling voor cricket; $\left\{\begin{array}{l} daarentegen \\ \# \ ook \\ \# \ omdat \\ \# \ niettemin \end{array}\right\}$ heeft Jane alleen belangstelling voor rugby.
(Jim's only interested in cricket; *Jane's only interested in rugby.)*

The distinction between ADDITIVE and CAUSAL relations in Dutch can be motivated by the cue phrase *en*, the translation of *and*. Figure 9 is the Dutch analogue of Figure 6. Motivating examples appear below:

(26) Jim heeft alleen maar belangstelling voor cricket; $\left\{\begin{array}{l} daarentegen \\ \checkmark \ en \\ \# \ ook \\ \# \ omdat \\ \# \ niettemin \end{array}\right\}$ heeft Jane alleen belangstelling voor rugby.
(Jim's only interested in cricket; *Jane's only interested in rugby.)*

---

[6] Note that the position of the verb within a clause is allowed to vary in the Dutch examples, as discussed in Section 4.2.2.

Fig. 9. *En* as a Marker of ADDITIVE Relations

$$
(27) \quad
\begin{array}{l}
\text{Jan houdt van} \\
\text{mosselen. Bill} \\
\text{houdt} \\
\textit{(John likes} \\
\textit{mussels. Bill likes}
\end{array}
\left\{
\begin{array}{l}
ook \\
\checkmark\ en\ \leftarrow \\
\#\ omdat \\
\#\ niettemin \\
\#\ daarentegen
\end{array}
\right\}
\begin{array}{l}
\text{van mosselen.} \\
\textit{mussels.)}
\end{array}
$$

Likewise, the POSITIVE/NEGATIVE POLARITY distinction can be motivated by *maar*, the Dutch translation of *but*, as illustrated in Figure 10, the Dutch analogue of Figure 7. Again, motivating examples are given below:
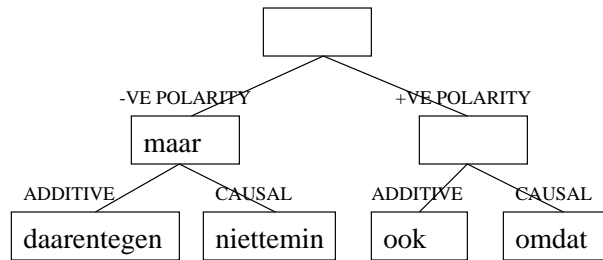

Fig. 10. *Maar* as a Marker of NEGATIVE POLARITY Relations

$$
(28) \quad
\begin{array}{l}
\text{Jim heeft alleen} \\
\text{maar belangstelling} \\
\text{voor cricket;} \\
\textit{(Jim's only} \\
\textit{interested in} \\
\textit{cricket;}
\end{array}
\left\{
\begin{array}{l}
daarentegen \\
\checkmark\ maar \\
\#\ ook \\
\#\ omdat \\
\#\ niettemin
\end{array}
\right\}
\begin{array}{l}
\text{Jane alleen} \\
\text{belangstelling heeft} \\
\text{voor rugby.} \\
\textit{Jane's only} \\
\textit{interested in} \\
\textit{rugby.)}
\end{array}
$$

$$
(29) \quad
\begin{array}{l}
\text{Het was Jan's} \\
\text{verjaardag;} \\
\textit{(It was John's} \\
\textit{birthday;}
\end{array}
\left\{
\begin{array}{l}
niettemin \\
\checkmark\ maar \\
\#\ ook \\
\#\ want \\
\#\ daarentegen
\end{array}
\right\}
\begin{array}{l}
\text{was hij bedroefd.} \\
\textit{he was sad.)}
\end{array}
$$

Four broad categories of cue phrase have now been identified in both English and Dutch. In the rest of Section 5, each category will be considered individually. Within each category, we will concentrate on looking for evidence for the other two of Sanders *et al*'s parameters—SOURCE OF COHERENCE and ORDER OF SPANS, which seem to emerge mainly at the lower levels of the taxonomy. We begin in this section with the POSITIVE CAUSAL relations.

5.2.1 The English Taxonomy

In English, a prototypical example of a POSITIVE CAUSAL cue phrase is *because*. A diagram showing some of the hyponyms of *because* is given in Figure 11.


Fig. 11. Some POSITIVE CAUSAL Cue Phrases in English

As the following examples illustrate, *on the grounds that* is contingently substitutable for *therefore* and *so*:

(30) The bouncers refused us access to the bar, $\left\{ \begin{array}{l} \textit{because} \\ \checkmark \textit{ on the grounds} \\ \quad \textit{that} \\ \# \textit{ therefore} \\ \# \textit{ so} \end{array} \right\}$ we were wearing jeans.

(31) $\left\{ \begin{array}{l} \textit{Because} \\ \checkmark \textit{ on the grounds} \\ \quad \textit{that} \\ \checkmark \textit{ therefore} \hookrightarrow \\ \checkmark \textit{ so} \hookrightarrow \end{array} \right\}$ we were wearing jeans, the bouncers refused us access.

(32) $\left\{ \begin{array}{l} \textit{Because} \\ \checkmark \textit{ therefore} \hookrightarrow \\ \checkmark \textit{ so} \hookrightarrow \\ \# \textit{ on the grounds} \\ \quad \textit{that} \end{array} \right\}$ it was a bird with an orange chest and light belly, it must have been a robin.

These patterns of substitutability can be used to motivate both of Sanders *et*

al's remaining parameters: ORDER OF SPANS and SOURCE OF COHERENCE. Note that *therefore* and *so* can only be used for BASIC ORDER relations, where the cause is presented before the effect. They cannot, for example, be used in the NON-BASIC context of text 30. In addition, texts 30 and 32 show that *on the grounds that* can only signal SEMANTIC relations. It cannot be used in the argumentative context of text 32. In text 31, which is both BASIC and SEMANTIC, all three cue phrases can be used.

The picture is in fact more complex than this. *On the grounds that* seems to be specialised not only for SEMANTIC relations, but also for relations expressing a volitional cause. Consider text 33:

$$
(33) \quad \begin{array}{c} \text{Yesterday, some} \\ \text{tiles fell off the} \\ \text{roof,} \end{array} \left\{ \begin{array}{l} because \\ \# \ on \ the \ grounds \\ \quad that \end{array} \right\} \begin{array}{l} \text{there was a strong} \\ \text{wind.} \end{array}
$$

This text descibes a cause occuring in the world, and is hence SEMANTIC—but *on the grounds that* is clearly inappropriate. We can therefore hypothesise a new feature, VOLITIONALITY, with alternative values VOLITIONAL and NON-VOLITIONAL, to describe the cue phrase. Note that *on the grounds that* cannot be explained solely in terms of this new feature; it is inappropriate in volitional contexts if they are PRAGMATIC:

$$
(34) \quad \begin{array}{c} \text{Bill must be ill} \\ \text{today,} \end{array} \left\{ \begin{array}{l} because \\ \# \ on \ the \ grounds \\ \quad that \end{array} \right\} \begin{array}{l} \text{he has stayed in} \\ \text{bed all morning.} \end{array}
$$

Further grounds for the SEMANTIC-PRAGMATIC distinction come from two other hyponyms of *because*; *it follows that* and *for this reason*. Note that both of these cue phrases are also hyponyms of *therefore* and *so*; and thus specific to BASIC ORDER relations. Their support for the distinction can be seen in the following examples:

$$
(35) \quad \begin{array}{c} \text{The lights in the} \\ \text{living room are out.} \end{array} \left\{ \begin{array}{l} So \\ \surd \ It \ follows \ that \\ \# \ For \ this \ reason \\ \surd \ Because \ \hookleftarrow \end{array} \right\} \begin{array}{l} \text{the neighbours are} \\ \text{not at home.} \end{array}
$$

$$
(36) \quad \begin{array}{c} \text{Last Saturday} \\ \text{John's car ran out} \\ \text{of petrol on his} \\ \text{way back from the} \\ \text{office;} \end{array} \left\{ \begin{array}{l} so \\ \surd \ for \ this \ reason \\ \# \ it \ follows \ that \\ \surd \ because \ \hookleftarrow \end{array} \right\} \begin{array}{l} \text{he had to walk all} \\ \text{the way home.} \end{array}
$$

In text 36, *it follows that* is ruled as exclusive because it imparts a distinctively argumentative flavour to the text, which is originally purely descriptive.

A more complete taxonomy of the POSITIVE CAUSAL cue phrases is given in
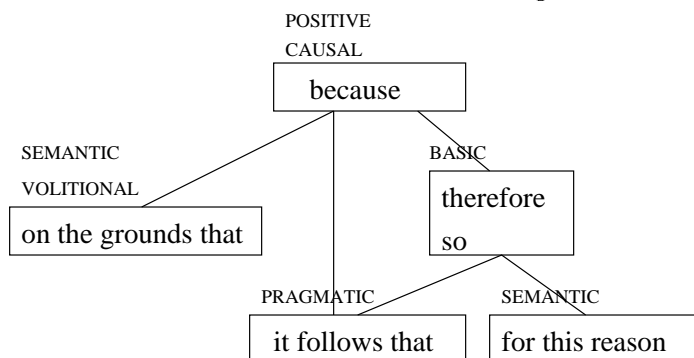Figure 12, which also includes labels for the different categories.

Fig. 12. More POSITIVE CAUSAL Cue Phrases in English

5.2.2 The Dutch Taxonomy

The Dutch taxonomy is slightly different from the English one, as Figure 13
shows. There is no single cue phrase dominating all the others. *Omdat*, being
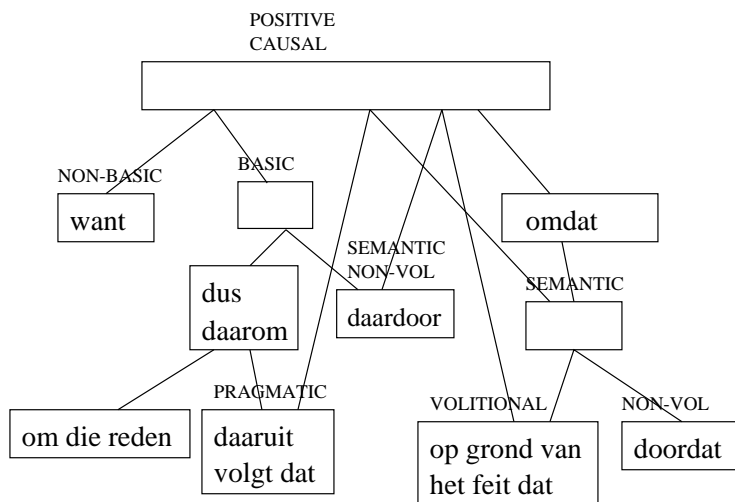
Fig. 13. Some POSITIVE CAUSAL Cue Phrases in Dutch

a subordinator, can be used for both BASIC and NON-BASIC order relations;
but it can't always be used for PRAGMATIC relations:

(37) De buren zijn niet
       thuis, { *want* } het licht in hun
       (The neighbours { # *omdat*[7] } woonkamer is uit.
       are not at home, *the light in their*
       living-room is out.)

---

[7] There is some discussion about the acceptability of *omdat* in this case. In our
view, *omdat* cannot replace *want*, unless it is preceded by a very atypical long

On the other hand, *want* can be used to express both SEMANTIC and PRAG-
MATIC POSITIVE CAUSAL relations, but only those with NON-BASIC ORDER:

(38) Het moet wel een
roodborstje
geweest zijn
*(It must have been
a robin,*
$\left\{ \begin{array}{l} want \\ \checkmark\,omdat \\ \#\ dus \end{array} \right\}$
het was een vogel
met een oranje
borst en een lichte
buik.
*it was a bird with
an orange chest
and a light belly.)*

(39) $\left\{ \begin{array}{l} Omdat \\ \checkmark\,dus \hookrightarrow \\ \checkmark\,daaruit\ volgt\ dat \hookrightarrow \\ \checkmark\,om\ die\ reden \\ \#\ Want \end{array} \right\}$
het een vogel was met een oranje borst
en een lichte buik, moet het wel een
roodborstje geweest zijn.
*it was a bird with an orange chest and
a light belly, it must have been a robin.)*

These examples also show that *dus* can only be used in relations with BASIC
ORDER.

The SEMANTIC/PRAGMATIC distinction for Dutch POSITIVE CAUSALs can be
motivated from the exclusive relationship between *daaruit volgt dat* and both
doordat and *op grond van het feit dat*. The former cue phrase is typically
PRAGMATIC, while the latter two are typically SEMANTIC. For instance, note
the exclusivity between the phrases in the following argumentative context,
even if *doordat* and *op grond van het feit dat* are moved to the first clause:

(40) Het licht in hun
woonkamer is uit.
*(The light in their
living room is out.*
$\left\{ \begin{array}{l} Daaruit\ volgt\ dat \\ \checkmark\,dus \\ \#\ doordat \\ \#\ op\ grond\ van \\ \quad het\ feit\ dat \\ \#\ om\ die\ reden \\ \#\ daarom \end{array} \right\}$
de buren zijn niet
thuis.
*the neighbours are
not at home.)*

It might also be thought that the SEMANTIC/PRAGMATIC distinction can be
motivated from the relationship between *daaruit volgt dat* and *om die reden*.
Om die reden also seems to be prototypically SEMANTIC, and is often not

---

pause. Hence a writer would never substitute omdat for want in this example. And
that is the criterion in the substitutability test.

substitutable by *daaruit volgt dat*:

(41) Afgelopen zaterdag kwam Jan zonder benzine te staan op de terugweg van kantoor. *(Last Saturday John's car ran out of petrol on his way back from the office.)* $\left\{ \begin{array}{l} dus \\ \checkmark\ om\ die\ reden \\ \#\ \ daaruit\ volgt\ dat \\ \checkmark\ daarom \end{array} \right\}$ heeft hij drie kilometer naar huis moeten lopen. *he had to walk three kilometers home.)*

However, note that in text 39, the two phrases do seem to be substitutable. The status of *om die reden* is therefore still unclear.

The distinction between VOLITIONAL and NON-VOLITIONAL relations also follows from the diagram in Figure 13: *doordat* can only be used in NON-VOLITIONAL relations, while *op grond van het feit dat* can only be used in VOLITIONAL ones. Note that these two cue phrases have a common hyponym, *omdat*.

(42) Er zijn gisteren een paar pannen van het dak gevallen, *(Some tiles fell off the roof yesterday,* $\left\{ \begin{array}{l} omdat \\ \checkmark\ doordat \\ \#\ \ op\ grond\ van \\ \quad het\ feit\ dat \end{array} \right\}$ er een harde wind stond. *there was a strong wind.)*

(43) De uitsmijters weigerden ons de toegang, *(The bouncers refused us access to the bar,* $\left\{ \begin{array}{l} omdat \\ \checkmark\ op\ grond\ van \\ \quad het\ feit\ dat \\ \#\ \ doordat \end{array} \right\}$ we een spijkerbroek droegen. *we were wearing jeans.)*

Dutch differs from English in that connectives exist specifically for identifying SEMANTIC NON-VOLITIONAL relations—*doordat* and *daardoor*.[8] *Daardoor* is less widely applicable than *doordat*, in that it is restricted to BASIC order

---

[8] The preposition *through* can sometimes be used to indicate a non-volitional cause, but it requires the cause to be nominalised, and often works only for causes expressed as abstract nouns (e.g. *through sheer hard work, through negligence*). *Doordat* is much more generally applicable.

relations (first antecedent, then consequent).

(44)
$$\left\{\begin{array}{l} \textit{Omdat} \\ \checkmark \textit{Doordat} \\ \checkmark \textit{daardoor} \rightarrow \end{array}\right\}$$
er gisteren een harde wind stond, zijn
er een paar pannen van het dak
gevallen.
there was a strong wind yesterday,
some tiles fell off the roof.)

Consider next the pattern of substitutability for the phrase *dus*. We have
already seen that it can only appear in BASIC ORDER. We have also seen
that it can appear in PRAGMATIC contexts; for instance, text 40 above. It can
also be used in VOLITIONAL SEMANTIC contexts, as text 44 shows. However,
it can *not* be used to signal NON-VOLITIONAL SEMANTIC POSITIVE CAUSAL
relations. Consider the following case:

(45)
Er stond gisteren
een harde wind,
(There was a
strong wind
yesterday,
$$\left\{\begin{array}{l} \textit{daardoor} \\ \checkmark \textit{doordat} \leftarrow \\ \checkmark \textit{omdat} \leftarrow \\ \# \ \textit{dus} \end{array}\right\}$$
zijn er een paar
pannen van het dak
gevallen.
some tiles fell off
the roof.)

The substitution of *daardoor* by *dus* leads to an odd interpretation; as though
a conclusion is being drawn that the tiles have fallen, or as though the tiles fell
of their own volition. In summary, *dus* cannot be used to express SEMANTIC
NON-VOLITIONAL relations, while with *daardoor* and *doordat* the situation
is exactly the reverse: they can only express SEMANTIC NON-VOLITIONAL
relations (Pander Maat and Sanders, 1995).

A final puzzle concerns the relationship between *dus* and *daarom*. For the
moment, they appear in our diagram as always substitutable for one another.
A typical context where both phrases are appropriate is given in text 41.
However, text 40 provides a context where *dus* is not substitutable by *daarom*,
and text 46 provides a context where *daarom* is substitutable, but not *dus*.

(46)
$$\left\{\begin{array}{l} \textit{omdat} \\ \checkmark \textit{daarom} \hookrightarrow \\ \# \ \textit{dus} \end{array}\right\}$$
we een spijkerbroek droegen, weigerden
de uitsmijters ons de toegang.
we were wearing jeans, the bouncers
refused us access to the bar.)

Further dimensions of variation are required to account for the distribution of
these two closely related connectives in these examples. The difference might
have to do with the amount of "certainty" the speaker shows, or with the per-
spective (J. Sanders, 1994) from which the utterance is produced (see Pander
Maat and Sanders, 1995).

27

In comparing the Dutch and English taxonomies for POSITIVE CAUSAL relations, the most interesting points to note are the subtle differences between them. For instance, as we have seen, there is no cue phrase corresponding precisely to *doordat* or *daardoor* in the English taxonomy. Mismatches of this sort between the two taxonomies are interesting: the combination of SEMANTIC and NON-VOLITIONAL emerges as a legitimate possibility from analysis of the English taxonomy, but no actual English cue phrase exists with these characteristics. Finding a cue phrase in Dutch which signals just these features raises an interesting possibility: that cue phrases in both languages are built up from the same primitives, even though they might not group these primitives in exactly the same way.

Another example of gaps in one taxonomy being filled by phrases in the other is provided by the English cue phrase *because*. This phrase sits at the top of the taxonomy of English POSITIVE CAUSAL cue phrases, and is undefined for all of the other parameters discussed so far. There is no such cue phrase in Dutch—however, the distinctions made elsewhere in the Dutch taxonomy do allow for the possibility of such a phrase.

A final interesting conclusion concerns the translation of English *for this reason* with Dutch *daarom*. This is not always adequate; the use of *daarom* seems to be more restricted.

We know already that there is no one-to-one mapping between the sets of cue phrases in English and Dutch. But it would be interesting to discover that both taxonomies exhibit the same *dimensions of variation*. Such a finding would enable the hypothesis that 'the cue phrases in a language mirror a set of basic cognitive mechanisms' to be maintained at a deeper level.

5.3 NEGATIVE CAUSAL *Relations*

5.3.1 *The English Taxonomy*

Figure 14 shows some of the hyponyms of the cue phrase *nevertheless*.

Fig. 14. NEGATIVE CAUSAL Cue Phrases in English

Relations signalled by *nevertheless* are similar to those discussed in the previous section, in that a causal operation can be identified which links the propositions in the two related spans. However, rather than linking a cause directly with a result, *nevertheless* presents a proposition along with the *negation* of its expected result. For instance:

(47)   Bill was very brave. He was badly injured; *nevertheless* he finished the game.

Such relations are termed NEGATIVE POLARITY CAUSAL. The BASIC OPERATION is a causal relation, in this case one whose antecedent is 'X is badly injured' and whose consequent is 'X does not finish the game'. The relation in text 47 has NEGATIVE POLARITY, because the consequent of the relation is negated.

A second example of a NEGATIVE CAUSAL relation is given in text 48.

(48)   A: Do you think Scotland will win tonight?
         B: I know everyone's betting against them. *Nevertheless*, they should win. They have a great team.

These two examples illustrate the SEMANTIC/PRAGMATIC distinction seen in the previous section. While the basic operation in text 47 holds between two propositions describing events in the world, we cannot say the same for text 48: we do not everyone betting against a team to *cause* them to lose. The expected cause in this case must be understood at the level of arguments about what is true in the world: the *belief* that everyone's betting against Scotland would normally cause the *inference* that they will lose.

For our present purposes, the interesting thing is to note that while a phrase like *nevertheless* can be used for both SEMANTIC and PRAGMATIC NEGATIVE CAUSAL relations, other cue phrases seem specialised for one or the other reading. For instance, the phrase *admittedly...but* can only be used in PRAGMATIC contexts, as the following examples show:

(49)   Bill was very brave. $\left\{ \begin{array}{l} despite\ this \\ \checkmark\ nevertheless \\ \#\ Admittedly...but^9 \end{array} \right\}$ he finished the
         He was badly injured; game.

(50) A: Do you think
Scotland will win
tonight?
B: Everyone's
betting against
them;
$\left\{\begin{array}{l} Admittedly\ldots but \\ \checkmark\ nevertheless, \\ \#\ despite\ this, \end{array}\right\}$
they should win.
They have a great
team.

Admittedly. . . but seems quite wrong in text 49—it suggests that Bill's being badly injured actually *detracts* from the claim that Bill was brave. However, the phrase is a good signaller of the argumentative relation in text 50.

The examples also indicate patterns of substitutability for the phrase *despite this*, which shows a preference for SEMANTIC NEGATIVE CAUSAL relations. In the PRAGMATIC context of text 50, the phrase is quite odd; it seems to presuppose that everyone's betting against Scotland would normally *ensure* their losing. Although *despite this* is not an impossible cue phrase here, it is certainly strange in the argumentative context set up by *admittedly. . . but*. For this reason, *despite this* and *admittedly. . . but* are shown as exclusive in the extract from the taxonomy in figure 14.

5.3.2 The Dutch Taxonomy

Here the English taxonomy does seem to have its equivalent in Dutch—see Figure 15. The motivating examples are given below:

Fig. 15. NEGATIVE CAUSAL Cue Phrases in Dutch

(51) Bill was erg
dapper. Hij was
zwaar geblesseerd,
*(Bill was very
brave. He was badly
injured,*
$\left\{\begin{array}{l} desondanks \\ \checkmark\ niettemin \\ \#\ Toegegeven\ldots maar \end{array}\right\}$
speelde hij de
wedstrijd uit.
*he finished the
game.)*

---

[9] The cue phrase *admittedly. . . but* is to be read as distributed between the two clauses of a relation. The version of text 49 using this phrase would thus read as follows: # *Bill was very brave.* Admittedly, *he was badly injured;* but *he finished the game.*

(52)

A: Denk je dat
Schotland
vanavond wint?
(A: Do you think
Scotland will win
tonight?)
B: Iederen denkt
van niet;
(B: Everyone's
betting against
them;

$\left\{\begin{array}{l} \textit{Toegegeven}\ldots\textit{maar} \\ \checkmark\,\textit{niettemin} \\ \#\ \textit{desondanks} \end{array}\right\}$

ze zouden moeten
winnen. Het is een
sterk team.
they should win.
They have a great
team.)

5.4 NEGATIVE ADDITIVE *Relations*

5.4.1 *The English Taxonomy*

The substitutability relationships exhibited by the contrastive marker *whereas* can also be described in terms of the SEMANTIC-PRAGMATIC distinction, as Figure 16 shows.

Fig. 16. NEGATIVE ADDITIVE Cue Phrases in English

The exclusivity between *whereas* and *then again* is illustrated in the following examples.

(53)

The Dalton brothers have
come to live next door. They
are very different. Averell is
tall;

$\left\{\begin{array}{l} \textit{whereas} \\ \checkmark\,\textit{on the other hand} \\ \#\ \textit{then again,} \\ \#\ \textit{having said that,} \end{array}\right\}$

Joe is
short.

(54)

I just cannot decide
whether to bet for
or against the Scots
tonight. They're on
good form;

$\left\{\begin{array}{l} \textit{then again,} \\ \checkmark\,\textit{on the other hand} \\ \#\ \textit{whereas} \end{array}\right\}$

Holland haven't
lost at home all
season.

Whereas is used to signal a contrast between two contradictory predicates,

such as 'tall' and *short* in text 53. In such cases, the contrast is between states of affairs described in the text. *Then again*, on the other hand, seems to signal a contrast between the argumentative force of two spans. In text 54, the two related spans 'Scotland are on good form' and 'Holland haven't lost at home all season' suggest contradictory conclusions about the outcome of a football match. There is no sense in which the two spans in text 53 argue in opposite directions—quite the contrary. So here, *then again* is not appropriate. Note that *on the other hand* is appropriate for both kinds of contexts. We can therefore think of it as being defined as NEGATIVE and ADDITIVE, but undefined for SOURCE OF COHERENCE.

A final question relates to the difference between *then again*, here described as PRAGMATIC NEGATIVE ADDITIVE, and the phrase *admittedly. . . but*, which was described in Section 5.3.1 as PRAGMATIC NEGATIVE CAUSAL. The feature combinations are incompatible, which suggests that the phrases are exclusive; but are they really so different? *Admittedly. . . but* can certainly appear in similar contexts to those in which *then again* is appropriate:

(55)   *Admittedly,* Scotland are on good form; *but* Holland haven't lost at home all season.

However, this text is significantly different from the preceding text 54. Both texts present premises supporting opposite conclusions; but while *then again* is explicitly neutral as to which premise wins, *admittedly. . . but* explicitly favours the second premise. This becomes even clearer when a conclusion is presented in the text:

(56)   Holland should win tonight. Scotland are on good form; $\left\{ \begin{array}{l} admittedly. . . but \\ \# \; then \; again \end{array} \right\}$ Holland haven't lost at home all season.

Then again sits oddly here with the speaker's obvious inclination towards the conclusion supported by the second premise. In other words, while on the face of it *admittedly. . . but* can link two conflicting premises, the second premise is better analysed as an indirect statement of the conclusion which has been reached. Such cases can thus be grouped along with the PRAGMATIC NEGATIVE CAUSAL examples discussed in section 5.3.1.

5.4.2   The Dutch Taxonomy

A similar story can be told for Dutch NEGATIVE ADDITIVE cue phrases. The relevant area of the taxonomy is given in Figure 17. Motivating examples are

```
              NEGATIVE
              ADDITIVE
      ┌─────────────────────────────┐
      │   daar staat tegenover dat  │
      └─────────────────────────────┘
   SEMANTIC  ╱                    ╲  PRAGMATIC
  ┌──────────────┐        ┌─────────────────────┐
  │  daarentegen │        │  aan de andere kant │
  └──────────────┘        └─────────────────────┘
```

Fig. 17. NEGATIVE ADDITIVE Cue Phrases in Dutch

given below:

(57) De gebroeders Dalton zijn hiernaast komen wonen. Ze zijn heel verschillend. Averell is groot.
(*The Dalton brothers have come to live next door. They are very different. Averell is tall.*)
$\left\{\begin{array}{l} Daarentegen \\ \checkmark\, Daar\ staat \\ \quad tegenover\ dat \\ \#\ Aan\ de \\ \quad andere\ kant \end{array}\right\}$ is Joe klein. *Joe is short.*)

(58) Ik kan maar niet beslissen of ik voor of tegen de Schotten zal wedden vanavond. Ze zijn goed in vorm.
(*I just cannot decide whether to bet for or against the Scots tonight. They're on good form.*)
$\left\{\begin{array}{l} Aan\ de \\ andere\ kant, \\ \checkmark\, Daar\ staat \\ \quad tegenover\ dat \\ \#\ Daarentegen \end{array}\right\}$ Nederland heeft het hele seizoen nog geen thuiswedstrijd verloren. *Holland haven't lost at home all season.*)

Daar staat tegenover dat and *daarentegen* seem to behave very differently. The former phrase can be used to express both semantic and pragmatic relations, whereas *daarentegen* is restricted to semantic relations. From the point of view of the English-Dutch comparison it should be noted that the closest literal translation of English *On the other hand*, Dutch *aan de andere kant*, really has another meaning: it is only used to express PRAGMATIC relations.

Having said that, the distinction between *toegegeven... maar* and *aan de andere kant* seems similar to that discussed in the previous section between *admittedly... but* and *then again*. Like *admittedly... but, toegegeven... maar* either presents a conclusion, or a premise whose conclusion is explicitly adopted

by the writer. *Aan de andere kant*, like *then again*, is constrained to present two opposing premises, and remains neutral as to which premise wins out. A final example illustrates this:

(59) Nederland gaat vanavond winnen. Schotland is goed in vorm (*Holland should win tonight. Scotland are on good form* $\left\{ \begin{array}{c} toegeven\dots maar \\ \#\ aan\ de \\ andere\ kant \end{array} \right\}$ Nederland heeft het hele seizoen nog niet thuis verloren. *Holland haven't lost at home all season.*)

5.5 POSITIVE ADDITIVE *Relations*

5.5.1 *The English Taxonomy*

Finally, consider some of the hyponyms of *also*, illustrated in Figure 18. Here
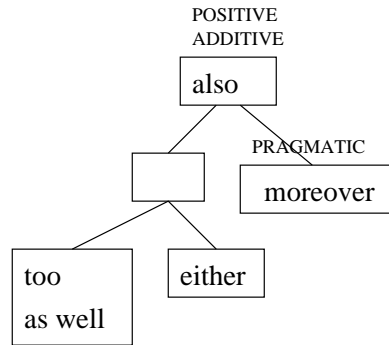
Fig. 18. POSITIVE ADDITIVE Cue Phrases in English

it seems slightly harder to motivate a straightforward SEMANTIC-PRAGMATIC distinction. *Moreover* seems clearly to be specific for PRAGMATIC ADDITIVE relations; it is appropriate in text 60, but not in text 61:

(60) I don't like Bill. He's always bothering me; $\left\{ \begin{array}{c} moreover, \\ \checkmark\ also, \end{array} \right\}$ my little sister is frightened of him.

(61) After he had read the morning paper, Bob set about cleaning the house. He swept the floors and washed them; he $\left\{ \begin{array}{c} also \\ \#\ moreover \end{array} \right\}$ tidied the cupboards. He had lunch at about twelve o'clock.

34

The problem with substituting *moreover* in text 61 is that it forces an argumentative reading that was not expressed in the original text. This becomes especially clear when the closing sentence of the text is taken into account. *Moreover* sits oddly with the narrative style of the opening sentence, and also with the neutral concluding sentence; for these reasons it is ruled as exclusive. In a coherent text with *moreover*, the preceding segments would have to be interpreted as arguments towards some kind of conclusion. It is, of course, possible to get this reading; see text 62.

(62) Bob has really been working hard today. He fetched the groceries. He $\left\{ \begin{array}{l} \textit{also} \\ \checkmark \textit{moreover} \leftarrow \end{array} \right\}$ swept the floors and washed them. He tidied the cupboards. I have never seen him working so hard on a Saturday.

However, it is less easy to find cue phrases specific to SEMANTIC ADDITIVE relations. *Too* and *as well* might seem to be plausible candidates; for instance, they are slightly odd in the following argumentative context:

(63) I don't like Bill. He's always bothering me. $\left\{ \begin{array}{l} \textit{Moreover,} \\ \#\ ?\ \textit{as well} \\ \#\ ?\ \textit{too} \\ \checkmark \textit{also} \end{array} \right\}$ my little sister is frightened of him.

Too and *as well* in this context are confusing. It is no longer clear that the final clause is to count as a second reason for my disliking Bill; a second interpretation, where the final clause is simply an additional fact, is also available.

There are nevertheless some clearly argumentative contexts in which *too* and *as well* can appear, namely when they share the same topic. For instance, they are perfectly acceptable as substitutes in text 62.

It thus seems hard to find a cue phrase tailored specifically for SEMANTIC POSITIVE ADDITIVE relations.

Note that the phrase *also* shows different patterns of substitutability as a clause modifer (i.e. in clause-initial position) and as a VP modifier. As a clause modifier, it patterns with *moreover*, while as a VP modifier it patterns with *too* and *as well*. (It is this which allows us to treat *also* as a common hypernym of all of these phrases.) The similarities and differences between the two types of *also* have not yet been fully examined; for the time being, we have grouped them together.

Finally, the exclusive relationship between *too/as well* and *either* should be briefly noted. Texts 64 and 65 illustrate this exclusivity:

(64)   John likes football. Bill likes football $\left\{ \begin{array}{l} too. \\ \checkmark\ as\ well. \\ \#\ either. \end{array} \right\}$

(65)   John doesn't like football. Bill doesn't like football $\left\{ \begin{array}{l} either. \\ \#\ too. \\ \#\ as\ well. \end{array} \right\}$

The difference between these examples seems quite straightforward: text 64 involves clauses with no explicit negations, and text 65 uses clauses with explicit negations. We thus need a feature, which we might call SURFACE POLARITY, which takes alternative values of POSITIVE and NEGATIVE. Clearly, most of the cue phrases in the taxonomy will be undefined for this feature.[10]

5.5.2   The Dutch Taxonomy

In the area of POSITIVE ADDITIVE relations, the Dutch taxonomy is very similar to the English one—see Figure 19. As in the English taxonomy there

Fig. 19. Some POSITIVE ADDITIVE Cue Phrases in Dutch

is a SEMANTIC/PRAGMATIC distinction in Dutch, reflected in the difference between *bovendien* and *ook*, as illustrated in the first two texts below.

(66)   Ik mag Bill niet. Hij valt me altijd lastig; $\left\{ \begin{array}{l} Bovendien \\ \checkmark\ Ook \end{array} \right\}$ is mijn kleine zusje bang van hem.
(I don't like Bill. He is always bothering me;   *my little sister is afraid of him.)*

---

[10] It will be required for some phrases, though; for instance *rather*, *instead* and *on the contrary* all require the first clause to have negative (surface) polarity, and the second clause to have positive (surface) polarity. (The taxonomy in Knott (1996) includes these phrases, but their feature-theoretic interpretation has not yet been fully worked out.)

(67) Nadat hij de ochtendkrant had gelezen, ging Bob het huis poetsen. Hij veegde de vloeren en dweilde ze. *(After he had read the morning paper, Bob set about cleaning the house. He swept the floors and washed them.* $\left\{\begin{array}{l} Ook \\ \# \ Bovendien \end{array}\right\}$ maakte hij de kasten schoon. Om 12 uur ging hij lunchen. *he tidied the cupboards...)*

In text 67 *bovendien* is not absolutely impossible, but (as with English *moreover*) it forces an argumentative reading of the text which is not present with the original phrase. In an explicitly argumentative context, though, both *bovendien* and *ook* are acceptable:

(68) Bob is werkelijk flink aan het werk geweest vandaag. Hij deed boodschappen. Hij veegde de vloeren en dweilde ze. *(Bob has really been working hard today. He fetched the groceries...* $\left\{\begin{array}{l} Ook \\ \checkmark Bovendien \end{array}\right\}$ maakte hij de kasten schoon. Ik heb hem nog nooit zo hard zien werken op een zaterdag. *he tidied the cupboards...)*

As with English *also*, the Dutch phrases *ook* and *eveneens* can appear as clause or VP modifiers, and pattern differently in the two cases. For instance, in text 69 *ook* is acceptable as a clause modifier (text 69 a), but not as a VP modifier (text 69 b):

(69) Ik mag Bill niet. Hij valt me altijd lastig.
(a) Ook is mijn kleine zusje bang van hem.
(b) # Mijn kleine zusje is ook bang van hem.

Note finally that the distinction between *eveneens* and *evenmin* to some extent parallels the distinction between English *too/as well* and *either*, and similarly suggests the existence of a parameter sensitive to 'surface polarity'.

(70) Jan houdt van voetbal. Bill houdt *(John likes football. Bill likes* $\left\{\begin{array}{l} ook \\ \checkmark \ eveneens \\ \# \ evenmin \end{array}\right\}$ van voetbal. *football)*

37

(71)
Jan houdt niet van voetbal. Bill houdt $\left\{ \begin{array}{l} \textit{evenmin} \\ \#\ \textit{eveneens} \\ \#\ \textit{ook} \end{array} \right\}$ van voetbal.
(John doesn't like football. Bill likes football.)

However, unlike English *either*, *evenmin* can only be combined with the drop of the negation in the second segment (In this respect it is more like English *neither* or *nor*.)

5.6 Summary: The Set of Parameters

Figures 20 and 21 provide taxonomies of all the cue phrases considered so far in English and Dutch, respectively. To conclude this section: it does seem that both taxonomies lend themselves to description in terms of Sanders *et al*'s parameters: BASIC OPERATION ORDER OF SPANS, POLARITY and SOURCE OF COHERENCE. When we compare the two languages, we can identify many systematic similarities, but also differences in the exact meaning of apparent literal translations.
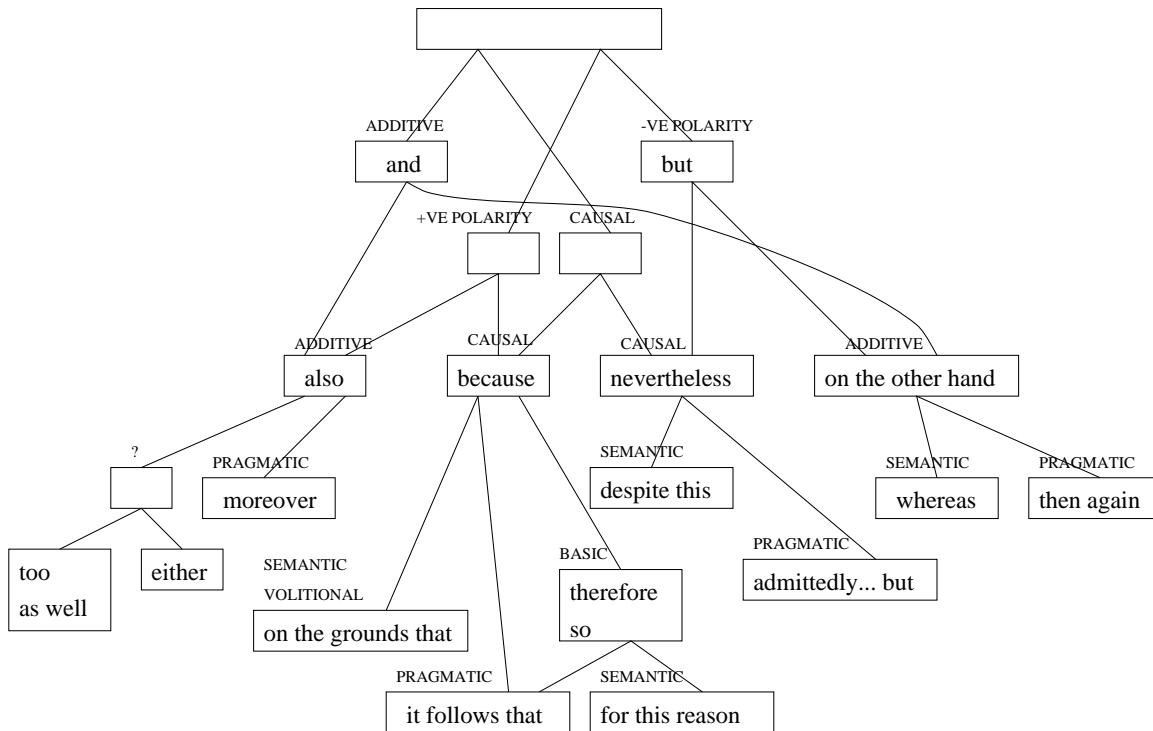

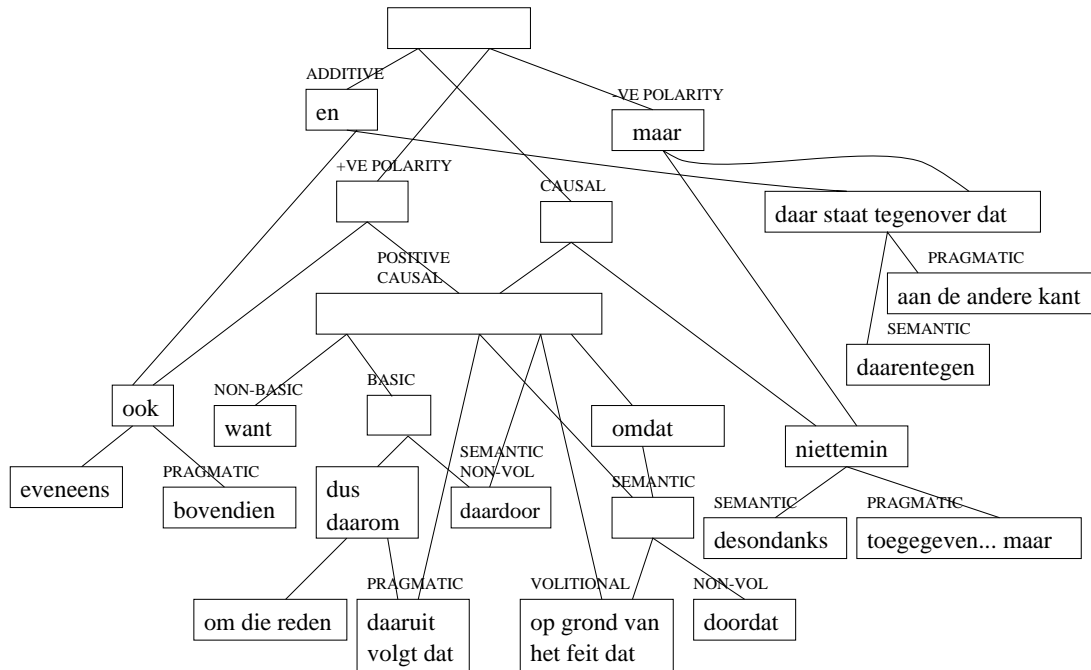Fig. 20. Summary of the English Taxonomy

Fig. 21. Summary of the Dutch Taxonomy

6 Discussion of the Substitutability Study

6.1 Extensions and Modifications to the Taxonomy

If Figure 20 is studied in depth, a number of exceptions to the proposed generalisations can in fact be found. For instance, it will be noted that there are contexts in which *and* can apparently be substituted for *so*. Consider text 72:

(72) John was careful with his boat. He had just finished painting it,
$\left\{\begin{array}{l} so \\ \checkmark\ and \\ \#\ but \\ \#\ whereas \end{array}\right\}$
he didn't want it to get scratched.

This calls into question our classification of *and* as a purely ADDITIVE cue phrase.

39

Problems also arise with the cue phrases *also* and *and*. There are contexts where *and* cannot be substituted for *also* (or some of its hyponyms):

(73) Sue planned to go to the party. But when she heard that Bob was $\left\{ \begin{array}{l} also \\ \checkmark\ too \rightarrow \\ \checkmark\ as\ well \rightarrow \\ \#\ and \end{array} \right\}$ going, she thought twice about it.

Clearly, the taxonomy must be altered to allow for examples such as these; the revised taxonomy will be considerably more complex, and some new parameters will be needed to describe it. [11] However, this is not necessarily a drawback for the approach we have taken here. Again, it must be recalled that Sanders *et al*'s parameters are only claimed to identify *classes* of relations—within these classes, further subdivisions are postulated which remain to be justified.

6.2   Equivocal Conclusions from the Substitutability Test

A further remark concerns our decision to use cue phrases and contexts where the substitutability relationships between phrases are relatively clear-cut. Many examples in Section 5 were chosen so that they differ optimally with respect to one of the four parameters identified by Sanders *et al.* By using this methodology we were able to identify 'exclusively substitute cue phrases' like *for this reason* versus *it follows that*. However, it is not hard to find contexts where the relationships are less clear. This is particularly true for the parameter SOURCE OF COHERENCE. For instance, consider the following cases (Sanders, 1997):

(74)   The lights are out. *So* the Carsons are not at home.
(75)   The Carsons had to go to Albuquerque early in the morning. *That's why* they were not at home yesterday evening.
(76)   There is a party downtown. *So* the Carsons are not at home.

Text 74 is clearly a PRAGMATIC CAUSAL relation, between a state-of-affairs and the speaker's claim or conclusion. The lights being out certainly does not cause the fact that the neighbours are not at home. The fact that they had to leave for Albuquerque does; therefore in text 75 a SEMANTIC relation exists; it might be an answer to the question "Why were the neigbours not at home yesterday evening?" Text 76 is a less clear case, which is ambiguous as regards the SOURCE OF COHERENCE parameter. It can be paraphrased either as expressing a *semantic* relation ('The fact that there is a party causes the fact that the neighbours are not at home') or a PRAGMATIC relation ('The fact that there is a party leads to my claim that the neighbours are not at home').

---

[11] Many of these changes have been incorporated into the taxonomy given by Knott (1996), which is much more complex, and motivates a larger set of features.

In context it would probably become clear which relation is intended, but the point is that in this paper we have typically discussed cases like texts 74 and 75, and not cases like text 76.

In summary, we should stress that our conclusions are based on a study of the behaviour of a group of particularly precise cue phrases in a number of particularly clear-cut contexts—as are most examples in the literature, in fact. There are, however, many less clear-cut contexts; and there also seem to be many cue phrases which are less specific and whose behaviour is harder to identify. It is of crucial importance to challenge existing text analytical conceptions and apply them to the indeterminate examples from corpora of naturally-occurring text.

Nonetheless, by beginning with the simple cases, we have developed a powerful method for the investigation of cue phrases, which should prove very useful in considering more complex cases. Now that the substitution methodology has been applied and examined in a simple domain, it should be possible to apply it to text corpora; see for instance Pander Maat (1997) and Pander Maat and Sanders (1995).

6.3   The Issue of 'Maximally Specific' Cue Phrases

It is interesting to examine the phrases at the leaves of the taxonomies we have constructed, i.e. the most specific phrases for the relations we have been considering. Almost all these cue phrases are lexical signals or idiom chunks, rather than simple connectives; and they tend not to be the most frequently used cue phrases, for either language. These two observations raise very interesting questions. After all, given a cognitive approach to language, and in particular to coherence relations, it would seem profitable for readers and listeners if speakers and writers would simply use the most specific cue phrase to express a given relation. This would be a theory of listener economy. However, this does not seem to be happening in the two languages we have looked at. Instead the most specific cue phrases are rarely grammaticalized, and the simplest cue phrases tend to be very general. Such findings lend support to a theory of *speaker* economy, where speakers choose the most general connectives to express a relation, and readers are required to draw implicatures to disambiguate. See Oberlander and Knott (1996) and Spooren (1997) for further discussion of this idea.

41

The present study has noted some interesting similarities and differences between the cue phrases in English and Dutch. Given that these two languages are very closely related—as mentioned previously—the two taxonomies can hardly count by themselves as substantial evidence for the cognitive primitives we have been hypothesising. Nevertheless, the present study is very useful as a model for further comparative work undertaken on a wider range of languages. Many studies have already revealed distinctions between cue phrases which seem on the face of it to be analysable in terms of the primitives we have been proposing here. For instance, a distinction between what we have been calling SEMANTIC and PRAGMATIC phrases has been proposed for several different languages: between Japanese *kara* and *node* (Takahara, 1990), French *car* and *puisque*, French *donc* and *à cause de ça* (Bentolila, 1986; Lamiroy, 1994), and German *denn* and *weil* (Keller, 1995), to mention only a few. The substitution methodology provides a good framework for more detailed comparative studies investigating similarities such as these. By finding clear-cut contexts in one language, making literal translations into another, and then investigating patterns of substitutability in both languages, we have been able to make systematic comparisons along several dimensions.

The methodology is useful particularly in that it allows comparisons to be made not just between individual phrases in two languages, but between *regions of taxonomies* in two languages. As we have seen in the present study, this means that we can detect in a principled way similarities between *components* of the semantics of phrases in two languages, even when there is no one-to-one translation between the phrases. This ability would be particularly useful in comparing languages whose sets of cue phrases are very different. And it is in just such cases that stronger evidence for the cognitive basis of the parameters must be sought. Again, this is an issue for further work.

7   Summary and Conclusions

This paper has presented two accounts for determining a set of coherence relations. They are independent, although they both rest on a conception of relations as modelling cognitive constructs operative in readers and writers. The relational account proposed to classify all coherence relations in terms of four cognitively plausible primitives, common to all relations, and to show how coherence relations themselves are related. It involves experiments on Dutch speakers, and rests on an analysis of disagreements between them over the relations present in various sample texts. The cue phrase account starts from the linguistic realisation of relations and rests on an argument that the

distribution of cue phrases in a language provides insight into the cognitive concepts underlying language processing. A taxonomy of English cue phrases is created which lends itself to a feature-theoretic interpretation.

The main question addressed in this paper is whether the two accounts and methodologies converge over the set of features they motivate. To test this, English and Dutch taxonomies were examined from the point of view of Sanders *et al*'s parameters. It was found that they provide a good framework for describing the relationships in both taxonomies.

The two methodologies also complement each other neatly. Sanders *et al.* argue for a parameter-based account of relations from psychological first principles, questioning the psychological plausibility of an unordered set of relations, and providing precedents for the parameters they propose. In Knott and Dale's study, the paramaterisation of relations falls directly out of the taxonomy of cue phrases as a fact to be explained. On the other hand, the link between relations and cue phrases argued for by Knott and Dale gives Sanders *et al.* a means of justifying the use of cue phrases as experimental indicators of the relations a writer is using; and an investigation of the taxonomy provides a systematic way of generating experimental hypotheses to test. The two methodologies thus fill in missing arguments for each other.

Most importantly, the convergence of the methodologies provides support for the hypothesis on which they are both based: that the set of coherence relations can be taken to represent a finite, ordered set of cognitive concepts underlying the construction and interpretation of text. It also provides support for the particular parameters proposed—POLARITY, BASIC OPERATION, SOURCE OF COHERENCE, and ORDER OF SPANS.

However, this support is qualified by a number of considerations. Firstly, we have only examined very restricted sets of cue phrases in English and Dutch. A larger taxonomy would provide a much more complicated picture, and no doubt require the postulation of additional primitives in order to be interpreted. Secondly, we have tended to concentrate on phrases and examples where the patterns of substitutability are relatively clear-cut; this raises problems in several areas which can only properly be addressed in corpus-based studies, using a more quantitative conception of substitutability relationships. Finally, we have considered a pair of languages which are closely related, and whose taxonomies are therefore relatively easy to compare.

For all these reasons, the present study should be seen primarily as a first experiment in the application of the substitution methodology in a cross-linguistic domain. Nonetheless, it demonstrates the usefulness of the methodology, in generating interesting hypotheses, and in providing a systematic framework for the exploration of these hypotheses in further theoretical and

empirical work.

References

Bentolila, Fernand, 1986. CAR en français écrit. *La Linguistique*, Vol. 22, 95–115.

Bereiter, Carl and Scardamalia, Marlene, 1987. *The Psychology of Written Composition*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Black, John and Bern, Hyman, 1981. Causal coherence and memory for events. *Journal of Verbal Learning and Verbal Behavior*, Vol. 20, 267–275.

Clark, Herbert H., 1974. Semantics and comprehension. In T. Sebeok, editor, *Current trends in linguistics*, Vol. 12, pages 1291–1498. Mouton, The Hague.

Flower, Linda L. and Hayes, John R., 1980. The dynamics of composing: Making plans and juggling constraints. In L. Gregg and E. Steinberg, editors, *Cognitive Processes in Writing*, pages 31–50. Lawrence Erlbaum Associates, Hillsdale, NJ.

Gernsbacher, Morton-Ann and Givòn, Talmy, editors, 1995. *Coherence in spontaneous text*. John Benjamins, Amsterdam/Philadelphia.

Grosz, Barbara J. and Sidner, Candace L, 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics*, pages 175–203.

Haberlandt, Karl, 1982. Reader expectations in text comprehension. In J. F. Le Ny and W. Kintsch, editors, *Language and Language Comprehension*, pages 239–249. North-Holland, Amsterdam.

Haberlandt, Karl and Bingham, George, 1978. Verbs contribute to the coherence of brief narratives: Reading related and unrelated sentence triples. *Journal of Verbal Learning and Verbal Behavior*, Vol. 17, 419–425.

Hobbs, Jerry R., 1979. Coherence and coreference. *Cognitive Science*, Vol. 3, 69–90.

Hobbs, Jerry R., 1983. Why is discourse coherent? In F. Neubauer, editor, *Coherence in Natural Language Texts*, pages 29–70. Buske, Hamburg.

Hobbs, Jerry R., 1990. *Literature and Cognition*. CSLI, Menlo Park, CA.

Hobbs, Jerry R., 1985. On the coherence and structure of discourse. Technical Report CSLI-85-37, Center for the Study of Language and Information, Stanford University.

Hovy, Eduard H., 1993. Automated discourse generation using discourse structure relations. *Artificial Intelligence*, Vol. 63, 341–385.

Hovy, Eduard H., 1990. Parsimonious and profligate approaches to the question of discourse structure relations. In *Proceedings of the 5th International Workshop on Natural Language Generation*, Pittsburgh.

Keller, Rudi, 1995. Das epistemische weil. In R. Keller, editor, *Zeichentheorie. Zu einer Theorie semiotischen Wissens*. Francke Verlag, Tübingen und Basel.

Knott, Alistair, 1996. *A Data-Driven Methodology for Motivating a Set of*

Coherence Relations. Ph.D. thesis, Department of Artificial Intelligence, University of Edinburgh.

Knott, Alistair and Dale, Robert, 1994. Using linguistic phenomena to motivate a set of coherence relations. *Discourse Processes*, Vol. 18(1), 35–62. Also available as Technical Report RP-34, Human Communication Research Centre, University of Edinburgh, 1992.

Knott, Alistair and Mellish, Chris, 1996. A feature-based account of the relations signalled by sentence and clause connectives. *Language and Speech*, Vol. 39(2–3), 143–183.

Lamiroy, Béatrice, 1994. Pragmatic connectives and l2 acquisition. the case of french and dutch. *Pragmatics*, Vol. 4, 183–201.

Longacre, Robert E., 1983. *The Grammar of Discourse: Notional and Surface Structures*. Plenum Press, New York.

Mann, William C. and Thompson, Sandra A., 1986. Relational propositions in discourse. *Discourse Processes*, Vol. 9, 57–90.

Mann, William C. and Thompson, Sandra A., 1988. Rhetorical structure theory: A theory of text organization. *Text*, Vol. 8(3), 243–281. Also available as Tech Report RR-87-190, USC Information Sciences Institute, Marina del Rey, CA.

Martin, James, 1992. *English Text: System and Structure*. Benjamin, Amsterdam.

Moore, Johanna D. and Paris, Cécile L., 1993. Planning text for advisory dialogues: Capturing intentional and rhetorical information. *Computational Linguistics*, Vol. 19, 651–694.

Noordman, Leo G. M. and Vonk, Wietske, 1993. Readers' knowledge and the control of inferences in reading. *Language and Cognitive Processes*, Vol. 7, 373–392.

Oberlander, Jon and Knott, Alistair, 1996. Issues in cue phrase implicature. In *Proceedings of the AAAI Spring Symposium on Computational Implicature*, pages 78–85, Stanford, CA.

Pander Maat, Henk, 1997. The classification of negative coherence relations and negative connectives. *Journal of Pragmatics*, (this issue).

Pander Maat, Henk. and Sanders, Ted J. M., 1995. Nederlandse causale connectieven en het onderscheid tussen inhoudelijke en epistemische coherentie-relaties (Dutch causal connectives and the distinction between content and epistemic coherence relations). In *Leuvense Bijdragen*, pages 349–379.

Sanders, José M., 1994. *Perspective in Narrative Discourse*. Ph.D. thesis, Tilburg University.

Sanders, Ted J. M., 1997. Semantic and Pragmatic Sources of Coherence: On the categorization of coherence relations in context. *Discourse Processes*, Vol. 20.

Sanders, Ted J. M. and Noordman, Leo G. M., submitted for publication. The role of coherence relations and their linguistic markers in text processing. Manuscript.

Sanders, Ted J. M., Spooren, Wilbert P. M., and Noordman, Leo G. M., 1992.

Towards a taxonomy of coherence relations. *Discourse Processes*, Vol. 15, 1–35.

Sanders, Ted J. M., Spooren, Wilbert P. M., and Noordman, Leo G. M., 1993. Coherence relations in a cognitive theory of discourse representation. *Cognitive Linguistics*, Vol. 4, 93–133.

Segal, Erwin. M., Duchan, Judith F., and Scott, Paula J., 1991. The role of interclausal connectives in narrative structuring: Evidence from adults' interpretation of simple stories. *Discourse Processes*, Vol. 14, 27–54.

Spooren, Wilbert P. M., 1997. The processing of underspecified coherence relations. *Discourse Processes*, (to appear).

Sweetser, Eve, 1990. *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Cambridge University Press.

Takahara, Paul O., 1990. Semantic and pragmatic connectives in English and Japanese. Paper presented at the IPrA conference, Barcelona, July.

Trabasso, Tom. and Sperry, Linda, 1985. Causal relatedness and importance of story events. *Journal of Memory and Language*, Vol. 24, 595–611.

Trabasso, Tom. and van den Broek, Paul, 1985. Causal thinking and the representation of narrative events. *Journal of Memory and Language*, Vol. 24, 612–630.

Traxler, Matthew, Sanford, Tony, Aked, Joy and Moxey, Linda, 1996. Processing causal and diagnostic statements in discourse. *Journal of Experimental Psychology: Learning, Memory and Cognition*.

Wason, Peter C. and Johnson-Laird, Philip N., 1972. *Psychology of reasoning: structure and content*. Batsford, London.