<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper published in *Philosophical Studies*. This paper has been peer-reviewed but does not include the final publisher proof-corrections or journal pagination.

Citation for the original published paper (version of record):

Pagin, P. (2014)

Pragmatic Enrichment as Coherence Raising.

Philosophical Studies, 168(1): 59-100

<http://dx.doi.org/10.1007/s11098-013-0221-8>

Access to the published version may require subscription.

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:su:diva-111592>

Pragmatic Enrichment as Coherence Raising*

Peter Pagin

October 11, 2014, 19:17

1. Introduction

This paper is concerned with the phenomenon of *pragmatic enrichments* and with the connection between pragmatic enrichments and *subject matter coherence*. Its main thesis is that free enrichments (not linguistically mandated) occur when they *raise* the degree of coherence. It is maintained that, as a consequence, pragmatic enrichments can to a large extent be predicted. Such predictions are systematic in the sense of relying on just a few principles, but unsystematic in the sense of also relying heavily on assumed general background theory of language users (“world knowledge”), especially as concerns causal regularities.

The term ‘enrichment’ has come to be well established in the pragmatics literature, and largely accepted across different theoretical standpoints, even though it is often employed as an alternative to more technical vocabulary that *does* differ between the positions. The phenomenon the term is used for describing is pretty much the same. A typical example is the following, from Carston 2002, p. 71:

- (1) a. He handed her the key and she opened the door.

*Earlier versions of this work have been presented at a conference on Peter Strawson’s work at the Academy of Science in Prague in April 2009, a conference on systematic pragmatics at Stockholm University in May 2009, at the Euro-XPrag conference in Leuven in June 2010, at the *Content, Context and Conversation* Lichtenberg-Kolleg Workshop at the Georg-August Universität Göttingen in June 2011, and at the *Language and Linguistics* workshop at the Inter-University Center of Dubrovnik in September 2011.

I am grateful for comments from the audiences on those occasions, especially Nicholas Asher, Christian Beyer, Stephen Butterfill, Herman Cappelen, Robyn Carston, Mari Duží, Bart Geurts, Kathrin Glüer, Julie Hunter, Jeff King, Luisa Martí, Petr Kotatko, Stephen Neale, Ira Noveck, Stefano Predelli, Francois Recanati, Petra Schumacher, Mandy Simons, Barry Smith, Paul Snowdon, Dag Westerståhl, and Deirdre Wilson. As always, I owe much to Kathrin Glüer’s inputs in discussions over the years of working on this theme.

I am also extremely indebted to comments from two anonymous referees. Their remarks led to improvements not only of the presentation and the discussion, but also of the theory itself.

This work has been made possible by a grant from the Tercentenary Foundation of the Swedish National Bank, project *Interpretational Complexity*, as well as by support from the ESF Eurocores project *CCCOM, Communication in Context*, led by Professor Åsa Wikforss,

- b. He handed her the key and she opened the door [*with the key that he had handed her*].

A normal and typical interpretation of (1a) associates with it a content that is more completely articulated in (1b), which includes additional linguistic material in brackets. The semantic contribution of this material to (1b) is the *pragmatic* enrichment of (1a). That is, when processing (1a) the hearer/reader tends to interpret it as representing a type of situation which is more completely represented by (1b). That the referent of ‘she’ used the key handed to her by the referent of ‘he’ is not semantically represented in (1a), but is “read into” it, i.e. pragmatically added during interpretation. It is semantically represented in (1b).

The idea is *not* that the hearer tacitly adds the bracketed *expression* during interpretation. It is also not the case, as in normal examples of ellipsis recovery, that there is a particular expression that *would be* recovered in any effort of making the enrichment explicit. Rather, it is the semantic content (in context) of the added phrase that matters, and often there are alternative possible linguistic additions that are semantically equivalent as far as the linguistic and extra-linguistic context goes. I’ll come back to the idea of pragmatic enrichment in section 2.

The idea of *coherence*, as developed in the theory of so-called *coherence relations*, or *discourse relations*, or *rhetorical relations*, is primarily an idea of what makes a text, or a discourse, hang together, to form an organized whole rather than just a sequence of unrelated or loosely related sentences. For instance, the relation between the first and the second conjunct of (1b) would be that of *result* in one established schema of coherence relations (section 3). There is a causal connection between the two events as they are described. By contrast, the relation between the first and second conjunct of (1a), without enrichment but *with* interpreting ‘she’ in the second as anaphoric on ‘her’ in the first, is just that of describing two events involving the same person. On common views, from semantics alone we don’t even get the temporal order between the events.¹

The example illustrates the main idea: the level of subject matter coherence is *higher* in (1b) than in (1a), both considered as unenriched. The conjuncts in (1b) are semantically more tightly connected. Hence, the enrichment that modifies the semantic content of (1a) into that of (1b) raises the level of coherence.² We tend to read a stronger connection, and thereby an enriched content, into the sentence. This requires a *strength ordering* between types of coherence relation. That will be the topic of section 4.

¹It is a further question whether semantics alone provides us with the interpretation that the two events are contiguous or whether that is a pragmatic choice.

²When formulating the first versions of these ideas, in early 2009, I was basically unaware of the coherence theory literature, and only alerted to the connection the following year.

Standard theories of coherence relations are concerned with relations between sentences, or whole clauses within a sentence. In the present proposal, the idea of coherence relations will be applied to a domain that is expanded in two respects. Firstly, it will be applied also to relations between contents of parts of speech, for instance between NP (noun phrase) and VP (verb phrase) interpretations. The way the NP referent is specified often indicates a certain modification of the property ascribed in the VP. Examples of this will be provided. Secondly, coherence will also be applied to relations between sentences in a context of utterance and contextual attitudes, i.e. attitude contents that are salient to speaker and hearer in the context, for instance because of a shared observation.

The rest of the paper will be organized as follows. In section 2, I shall characterize the phenomenon of enrichment according to current pragmatic frameworks. In section 3, I shall present basic ideas of coherence theory according to some standard frameworks. In section 4, I shall propose a schema of such relations with a strength ordering. In section 5, I spell out the present account, under the name ‘Enrichment Theory’ (‘ET’). In section 6, I provide applications of ET, and in section 7, I discuss some related phenomena that are treated as enrichments in some accounts but not in ET. In sections 8 and 9, I shall compare the present approach to those of Relevance Theory and Levinson’s Neo-Gricean Theory, respectively. In the concluding section 10, I discuss the account with respect to prediction and the modelling of cognitive processes.

2. Pragmatic enrichment

In a very wide sense of ‘pragmatics’, the term denotes everything that a speaker intends and a hearer interprets a linguistic utterance to convey that is not fully determined by standard (morphology and) syntax and semantic. Pragmatics, in this very general sense, covers syntactic and lexical disambiguation, anaphora resolution, ellipsis recovery, presupposition detection and accommodation, saturation, modulation, the understanding of conversational implicature, metaphor, irony, any form of indirect speech act, and more. Some of these, such as disambiguation, consist in choosing between alternatives that are made available by the semantics. Others involve adding something new. Perhaps the most basic kind is saturation.

Saturation (the term is from Recanati 2004, p. 7) is the process of providing values to indexicals and other context dependent expressions, including individuals to ‘I’ and ‘her’ in (2a), a quantifier domain to the quantifier ‘everyone’ in (2b), and a standard of height to the adjective ‘tall’ in (2c):

- (2) a. I like her.
- b. Everyone cheered.

c. He is tall.

These sentences also have tensed verbs, and the hearer typically assigns a time as part of interpreting their contributions to the truth conditions.

Saturation is characterized as being “linguistically mandated”, as Recanati puts it (2004, pp. 7-10). That is, there is some expression that triggers the saturation step. It is triggered because the expression needs a value in order for the sentence to express a proposition, i.e. a content that has a truth value with respect to a possible world. In case of (2a), no proposition is expressed by an utterance of the sentence except if values have been given to ‘I’, and ‘her’, and likewise a time is assigned. So, saturation is both *needed* for a proposition to be expressed, and *triggered* by a context dependent expression that is to be assigned a value.

Saturation, together with disambiguation, anaphora resolution etc, delivers “what is said”, in Grice’s original (locutionary/illocutionary) sense (Grice 1975, p. 25). In Grice’s theory of implicature, *what is said* is the immediate content of the utterance, assuming that the utterance has a particular force (typically assertoric), which provides the input to an implicature. In one of Grice’s classical examples (1975, p. 32),

- (3) A: I am out of petrol.
B: There is a garage around the corner.

B *says that* there is a garage around the corner and *implicates* that the garage is open and has petrol for sale. The conversational implicature is made *by means of* what is said. So *what is said* is supposed to be in place before implicature.

This does not mean that the hearer can always determine what is said independently of calculating an implicature. Sometimes, on the contrary, the hearer fixes an interpretation of an assertion because its implicature satisfies a conversational expectation:

- (4) A: What shall we do about the leakage?
B: I have talked with the insurance company.

Here, B’s answer generates a relevance implicature: B knows something relevant to taking care of the leakage in question, because she has talked with the insurance company *recently*. Only a recent conversation would provide such knowledge. In order that the answer satisfy Grice’s relevance maxim, it must therefore be taken to concern the recent past. On my preferred account, this is effected by means of a *domain restriction* on the present perfect tense. Domain restriction is a kind of saturation, and therefore is part of what delivers *what is said* (for further discussion, see section 5). In this case, then, the domain restriction is applied because of the resulting implicature.

This exemplifies a general pattern of partly bottom-up, partly top-down interpretation. Stephen Levinson names it “Grice’s Circle”:

Grice’s account makes implicature dependent on a prior determination of “the said.” The said in turn depends on disambiguation, indexical resolution, reference fixing, not to mention ellipsis unpacking and generality narrowing. But each of these processes, which are prerequisites to determining the proposition expressed, may themselves depend crucially on processes that look indistinguishable from implicatures. Thus what is said seems both to determine and to be determined by implicature. Let us call this Grice’s circle (Levinson 2000, p. 186).

Similarly, Robyn Carston describes the process (within Relevance Theory) as “[...] being one of mutual adjustment until the propositional forms stabilize into an inferentially sound configuration which meets the expectation of relevance” (Carston 2004a, p. 648).

Top-down effects are common in saturation and can apply even to the interpretation of so-called “automatic” indexicals, like ‘I’. Its associated *character* (cf. Kaplan 1989) is a function from contexts to their speakers, but this is only a default assignment and can be overridden. The first person singular pronoun can be used both impersonally and descriptively, as well as with deferred reference.³

Top-down effects do not derive only from implicature, but can be an effect of simply applying *charity* (in the sense of Davidson 1984) in interpretation. Here is a simple example. Indexical uses of the pronoun ‘he’ by default pick out the most salient male of the context, but this can be overridden by charity considerations, given background knowledge, as in:

(5) He is back,

where the intended referent is not the most salient male of the context *before* the utterance, but *is* the most salient male individual who is commonly known to have been recently absent. When fixing a referent that makes the utterance true, the chosen referent becomes the most salient male *after* the utterance.

Even though the interpretation process may move back and forth, implicature *does* build on what is said. That we sometimes reject a hypothesis about assertoric content because of its implicature effects, or lack of such, does not alter this basic relation. As an analogy, in science a theory is accepted because of its observational consequences. That we accept a theory because of its predictions does not alter the fact that the predictions

³Cf. Nunberg 1993, pp. 20-21, for descriptive examples and, for German examples of impersonal use, Zobel 2010. A referee suggested an actors’ reference to his role as a further example (which I would call deferred).

are derived from the theory, not the other way around. Similarly, in demonstrative sciences, an axiom may be accepted because of the theorems that can be derived from it (as the Axiom of Choice in Set Theory). The order of inferential dependence is sometimes opposite to the order of knowledge acquisition. In terms of inferential dependence, (particularized conversational) implicature depends on assertoric content, on what is said, and belongs to a different interpretational layer (see below for complications).

This paper is concerned with a pragmatic layer of interpretation that is characterized as being located *between* saturation and implicature. This layer of pragmatics has been named “modulations” by Francois Recanati (2004, p. 74; 2010, pp. 5-7). It has been called “explicatures” in Relevance Theory, primarily by Dan Sperber and Deirdre Wilson (1995, p. 182), and Robyn Carston (2002, p. 116), “implicatures” by Kent Bach (1994, p. 126).⁴ Saturation and modulation together has been called “primary pragmatic processes” by Recanati (2004, p. 17), intending to highlight that they occur *before* implicature.

There are important differences between these authors. Relevance theorists are contextualists; they think that the meaning of a sentence in principle underdetermines the content of an utterance made by means of it. In order to express a proposition, a contextually determined inferential process is always needed. Bach (2010, pp. 128-29) rejects contextualism. Recanati (2004, pp. 81-82) takes a kind of middle position. What he calls *the minimal proposition*, the result of saturation alone (after disambiguation etc.), is never, or hardly ever computed and never, or hardly ever, plays any role in utterance interpretation. On his view (Recanati (2010, pp. 39-47), pragmatic modulations are *intertwined* with semantic interpretation: for instance, the semantic interpretation of a predicate may take as argument the *modulated* content of the interpretation of a singular term.

These differences correspond to differences about *what is said*. Bach (1994) takes what is said to be what Grice meant by “what is said”, the locutionary content of an utterance. This is in many cases exactly the minimal proposition, but not in all. For both Relevance theorists (always) and Recanati (virtually always), *what is said* is already the result partly of pragmatic operations, i.e. explicatures and modulations, respectively. This means that we cannot be neutral between these theories in characterizing primary pragmatic effects as something that contributes to or as something that adds to what is said.

The task of neutrally characterizing the domain of primary pragmatics is even more

⁴Terminological differences are a little more complex. Officially, *explicatures* are the *results* of pragmatic operations, whereas *saturation*s and *modulation*s are the operations *themselves*, and *implicatures* are either the operations or the contents *added* by means of them.

complicated, because there are also differences of *extension* between the respective concepts. For Bach (1994, pp. 125-26), implicatures come in two varieties, *completion* and *expansion*. Completion takes place when the sentence uttered fails to express a complete proposition even after values have been given to indexicals, as in

(6) Steel isn't strong enough.

(Bach 1994, p. 127). Before it has been specified what steel is or isn't strong enough *for*, (6) cannot obtain a truth value. Syntactically, (6) is a complete sentence, but semantically it does not express a full proposition. However, by Recanati's conception, this completion is linguistically mandated, since semantically, compounds of the form Adjective+'enough', or Adjective+'enough for', are dyadic relational expressions, taking a subject argument and a purpose argument. The completion step of providing a purpose (e.g. for withstanding a laser beam of such and such a power) would therefore, by Recanati's standards, be classified as *saturation*, not as modulation.

An additional complication is that some of the phenomena that will be in focus here *are* treated as implicatures by Levinson, in the neo-Gricean tradition. More precisely, they are taken to be *generalized conversational implicatures*–GCIs. Generalized conversational implicatures were introduced by Grice (1975, p. 37) and contrasted with *particularized* conversational implicature. The difference was that while particularized implicatures depend on features of the context of utterance, the generalized implicatures are generated normally and by default, without relying on features of the context, although they might be blocked by such features. Grice's first example is

(7) X is meeting with a woman this evening

which normally, according to Grice, carries the implicature that the woman is not X's wife. The most well-known type is *scalar implicatures*, as in

(8) Some students passed the test

which would normally carry the implicature that not all students passed the test. The theory of scalar implicatures has been developed above all by Larry Horn (see Horn 2004 for an overview).

In Levinson's (2000) scheme, GCIs are of three varieties, Q, I, and M implicatures. The Q-implicatures include the scalar ones and are related to Grice's maxim of Quantity. The M-implicatures are related to Grice's maxims of Manner, and are taken to generate implicatures by way of the expression chosen among semantically equivalent ones. The category of I-implicatures is the one that will be relevant to the present concerns, since

many and perhaps all of the phenomena in focus here will count as I-implicatures by Levinson. In particular, Levinson characterizes the hearer's (recipient's) side of the *I-principle* (2000, pp. 114-15) as an *Enrichment Rule*, and spells it out partly in accordance with the present account (cf. section 9).

It might therefore seem to be an immediate clash with the neo-Gricean account that enrichments belong to a layer between saturation and implicatures. However, differences about the pragmatic *location* of enrichment phenomena are not that sharp. Levinson does think that GCIs in general play a role in determining assertoric content, i.e. what is said: "The crucial fact that I will try to establish is that generalized conversational implicatures seem to play a role in the assignment of truth-conditional content" (Levinson 2000, p. 166). According to Levinson, on some versions of the saying-implicating distinction, GCIs are explicatures in the Relevance Theoretic sense (2000, p. 194), and on others they are in many cases implicatures in Bach's sense (2000, p. 197). In the end, Levinson concludes, "[...] there is no consistent way of cutting up the semiotic pie such that 'what is said' excludes 'what is implicated'" (2000, p. 198).

I think, therefore, that there is a fairly wide agreement that the phenomena I shall focus on do belong to a pragmatic layer that comes before *particulated* conversational implicature, and therefore, in some employed senses of the expression, contribute to "what is said". It should be noted, however, that this is *not* equivalent to the idea that semantics and pragmatics are *intertwined*. In a technical sense, semantics is intertwined with pragmatics if the semantic *value* of an expression in a context is provided by pragmatic operations. This goes beyond using pragmatic operations to *select* a value, as e.g. in lexical disambiguation (in choosing between *riverbed* and *monetary institution* in the interpretation of an occurrence of 'bank', pragmatics does not provide the meaning). Clearly, in this sense semantics is intertwined with *saturation*. The most obvious example is the semantic value of an indexical in a context, which depends on the pragmatic operation of selecting a value. By contrast, disambiguation is not an example of intertwinement in this sense.⁵

In the case of enrichments (and modulations generally), as noted above, it is the explicit view of Recanati (2010, pp. 39-47) that they are intertwined with semantic interpretation. This is also Levinson's view. In what he calls *intrusive constructions*, "the truth conditions of the whole depend in part on the implicatures of the parts" (Levinson 2000, p. 198). By contrast, in the account of Pagin and Pelletier 2007, the semantic output depends only on saturation. Enrichments and other modulations occur at a post-semantic stage. Since enrichments can occur at a constituent level (e.g. extending the application

⁵Anaphora resolution, again, is a borderline case; to the extent that coreference is represented syntactically as part of the syntactic input to semantics, pragmatics does not provide a (functional) value.

conditions of a verb), modulations can occur at a post-semantic stage only if the semantic output is a *structured* meaning, and this is precisely what Pagin and Pelletier propose. The main reason here is that *these* pragmatic operations only depend on contents in context, not on which linguistic expressions have generated the content. So, on this account, pragmatic intrusion is limited to saturation, even though enrichments contribute to what is said.⁶

Recanati's (2004) generic term for the pragmatic layer between saturation and implicature is 'modulation'. He distinguishes between three kinds of modulation: *loosening*, *semantic shift* (or *semantic transfer*) and *free enrichment*. An example of loosening is

(9) The ATM swallowed my credit card.

(Recanati 2004, p. 24) where the verb 'to swallow' has its application conditions *extended* to include the cash machine process referred to.

Classic examples of semantic shift come from Geoffrey Nunberg, including

(10) The ham sandwich is sitting at table 20.

as said by one waiter to another in a restaurant (Nunberg 1979; Nunberg 1995). The content of the utterance is that

(10') The ham sandwich *orderer* is sitting at table 20.

What Recanati calls *free enrichment* is any *addition* of content material that is not needed for the content to be proposition, i.e. to have truth conditions; it is freely added in this sense. This is exemplified in (1). One of Recanati's own official examples is given in

(11) A: Do you want something to eat?
B: I have had breakfast.

According to Recanati, it is exemplified in B's utterance, where *today* is freely added to the content *I have had breakfast*. The idea is that without the enrichment, there is still a proposition, true if B has had breakfast at some time or other during his life. However, as in example (4), I prefer to see the restriction to breakfast on the day of utterance as

⁶This is again not the same view as e.g. in Stanley 2000. According to Stanley (2000, pp. 393-94), semantics is concerned with context invariant meaning *and* the assignment of contextual values to variables in the logical form. Stanley further thinks that semantics in *his* sense does provide truth conditions of utterances (assertoric content), hence, unlike the present view, that enrichments do not contribute to assertoric content. Moreover, as is demonstrated in Pagin 2005, semantic context sensitivity (in a more neutral sense than Stanley's) can be encoded by parameters in the *meta-language*, without relying on variables in the logical form of the *object language*.

a quantifier domain restriction, and therefore an example of saturation rather than free enrichment.

Still, I follow Recanati in focusing on *free* enrichments. The current idea is that enrichments occur in order that they raise the level of coherence. In order that the level be raised, there must be a level to begin with, and this already requires propositional content. Without propositional content, some addition of content is needed anyway, regardless of coherence. An example is Bach's (6). Out of context there is no chance for the hearer to find a plausible completion. A small addition to the linguistic context changes the situation:

(6') The material must withstand a beam for 30 minutes. Steel isn't strong enough.

The understanding of the second sentence is clearly *that steel isn't strong enough for withstanding a beam for 30 minutes*. This certainly is not a case of free enrichment. Also, it is not so clear that it is a case of enrichment at all, in the sense of an example where it is primarily *content* that is added. Rather, (6') looks like a case of ellipsis, where it is linguistic material from the preceding sentence that is tacitly recovered for providing the missing content. No doubt other examples of what Bach calls *completion* are not analyzable as cases of ellipsis, but although there may be cases where one candidate completion is preferred over another for reasons of coherence, the completion itself is triggered without regard of coherence. Rather, the phenomenon mostly in focus here is the one that Bach calls *expansion* (1994, pp. 125-26).

Within Relevance Theory, the notion of *explicature* is pretty much equivalent to *what is said* (cf. Carston 2004b). It is tightly connected with the ideas of *development* and *logical form*. The official definition in Sperber and Wilson 1995, p. 182 (where the explicature is precisely what is *explicit*):

An assumption communicated by an utterance U is *explicit* if and only if it is a development of a logical form encoded by U.

One aspect of developing logical form is enrichment. Carston says:

It is the enriched propositions that are communicated as explicatures and which function as premises in the derivation of implicatures; the uninformative, irrelevant, and sometimes truistic or patently false minimal propositions appear to play no role in the process of utterance understanding, which is geared to the recovery of just those propositional forms which the speaker intends to communicate. The pragmatic process at work here is known as free enrichment; it is "free" in that it is not under linguistic control. So, unlike saturation, it is an optional process, in the

sense that there can be contexts in which it does not take place, though these tend to be somewhat unusual (Carston 2004a, p. 639).

According to Relevance Theory, explicature is always underdetermined by the linguistic expression used (hence by what they call “logical form”), and so *development* is mandatory. But not everything that may go into development is mandatory, and free enrichment is a case, or maybe *the* case, in point.

3. Coherence relations

Discourse semantics and discourse theory emerged in the late 1970s and early 1980s. Whereas discourse *semantics*, for instance in the DRT (*Discourse Representation Theory*) framework (cf. Kamp and Reyle 1993) was concerned with giving a formal theory of how semantic content is incrementally added across fairly simple discourse, e.g. in order to handle cross sentence anaphora and donkey anaphora, discourse relations theory was designed to give an informal, high-level description of how parts of discourse are interrelated. The two projects were eventually combined with the emergence of SDRT, Segmented DRT, developed by Nicholas Asher and Alex Lascarides, especially in Asher and Lascarides 2003.

Jerry Hobbs is one of the leading pioneers in discourse relations theory, and his theory is succinctly presented in Hobbs 1985. I shall more primarily use a later work by Andrew Kehler (2002), who follows and develops Hobbs’s theory. After a suggestion by Hobbs, Kehler uses the categories of *connections between ideas* of David Hume (1748) as his basic categories of discourse relations: *Resemblance*, *Cause-Effect*, and *Contiguity*.

Kehler’s resemblance relations are *Parallel*, *Contrast*, *Exemplification*, *Generalization*, *Exception*, and *Elaboration*. As an example, *Parallel* is explained as follows (Kehler 2002, p. 16):

(Parallel) Infer $p(a_1, a_2, \dots)$ from the assertion of S_1 and $p(b_1, b_2, \dots)$ from the assertion of S_2 , where from some property vector $\vec{q}$, $q_i(a_i)$ and $q_i(b_i)$ for all i .

This quasi-formal style was introduced by Hobbs. Here S_1 is the first sentence of a two-sentence discourse. The illustrating example is

- (12) Dick Gephardt organized rallies for Gore, and Tom Daschle distributed pamphlets for him.

Here the relation p is that of *doing something in support of*, a relation which subsumes both *organize rallies for* and *distribute pamphlets for*. So the relation of *doing something in support of* holds both between Gephardt and Gore and between Daschle and Gore. In

addition, Gephardt and Daschle have the shared property of being high-ranking democratic politicians.⁷

Kehler's cause-effect relations are *Result*, *Explanation*, *Violated Expectation*, and *Denial of Preventer*. An example of *Explanation* (Kehler 2002, 21) is

(13) George is dishonest. He is a politician.

This satisfies the *Explanation* relation insofar as it is *presupposed* that being a politician *implies* being dishonest. The *Result* relation has the opposite order of mention of properties, which is exemplified in

(14) George is a politician, and therefore he's dishonest.

The remaining two relations are negated versions of the two above, exemplified by

(15) George is a politician, but he is honest.

(16) George is honest, even though he is a politician.

Kehler has only one contiguity relation, in two versions, taken over from Hobbs:

(Occasion (i)) Infer a state of a system of entities from S_1 , inferring the final state for this system from S_2 .

(Occasion (ii)) Infer a state for a system of entities from S_2 , inferring the initial state for this system from S_1 .

Examples provided are

(17) George picked up the speech. He began to read.

(18) Larry went to into a restaurant. The baked salmon sounded good and he ordered it.

These examples illustrate the basic idea that the first mentioned event sets the stage for the second mentioned event. That implies that the first event resulted in a state that still obtained when the second event took place. As is illustrated in (18), the second event need not follow immediately on the first. This completes this brief summary of the Hobbs-Kehler account.

⁷Here a_1 and b_1 correspond to Gephardt and Daschle, respectively, a_2 and b_2 , to Gore and the referent of *him*. q_1 is understood to be the property of being a high-ranking democratic politician. Clearly, the general idea is that there is some clear analogy between objects mentioned in the respective conjuncts, and between the properties, respectively.

A well-known problem for discourse relations theory is the lack of clear and plausible criteria for postulating such relations and evaluating the proposals. The existence of many alternative theories that overlap and differ “in myriad of ways” (Kehler 2002, 23) aggravates the problem. Kehler’s own solution was to base the theory on the threefold Humean distinction. My own project is to use relations of the kinds proposed in discourse relations theory for an account of pragmatic enrichment. If the relations that will be proposed fulfil this purpose, no further criterion of adequacy will be required.

4. A coherence hierarchy

To repeat, the main idea is that (free) pragmatic enrichment occurs when it serves to raise the degree of coherence. This idea presupposes a *coherence scale*, i.e. a sequence of *degrees* of coherence ordered according to *strength*, and below I shall propose such a scale. I shall propose five degrees of *content* coherence. The reason for the attribute ‘content’ is that I shall focus only on relations between *what is represented*, not on relations that depend, over and above on content, on the linguistic means of representation. For instance, we intuitively perceive (19b) to be more coherent than (19a), since pronouns are used to indicate reference to persons that are already introduced in the context.

- (19) a. John₁ bought a new car. John₁ crashed it the same day.
b. John₁ bought a new car. He₁ crashed it the same day.

Issues of this and related kinds, turning on the choice between expressions that are representationally equivalent in context, do not belong to the present topic.

In distinguishing the types of coherence, I shall largely follow the Hobbs-Kehler scheme, although they do not offer a scale. Ordered from weakest to strongest, the proposed types are as follows:

Scale of coherence strength

- 0) Vacuity
- 1) Contiguity type relations
- 2) Resemblance type relations
- 3) Possibility type relations
- 4) Necessity type relations

The first of the five, *Vacuity*, is the zero degree of coherence, when the content of the new utterance has nothing (available) to do with the given salient background, linguistic or non-linguistic. A beautiful example is given by Alistair Knott and Robert Dale (1994, p. 39):

(20) John broke his leg. I like plums.⁸

Of course, with sufficient ingenuity, any example of Vacuity can be enriched into a non-vacuous one. For instance, maybe John slipped on a plum, which causes me to like plums.⁹ This enrichment would not count as available, however (see next section). Given that coherence relations are binary relations between *contents*, not between contents and contexts, any factor that renders a relation non-vacuous must be added to the background content.

Next, *Contiguity* type relations are relations in which the new content is in some salient respect *close to* the background content. Typically, this is exemplified with spatio-temporal contiguity, as in:

(21) The table is covered with books. A cat is lying on the sofa.

(22) John turned left. He sneezed.

In (21), the assumption of discourse coherence generates the interpretation that the sofa mentioned in the second sentence is in the same room as the table mentioned in the first, and preferably close to that table (at least closer than any other sofa). This exemplifies what is known as *bridging inferences*, where the inference would be to the conclusion *that there is a sofa close to the table* (cf. Levinson 2000, pp. 37-38, Wilson and Matsui 1998). According to Levinson, bridging inferences are examples of I-implicatures.

In (22), the coherence assumption generates the interpretation in which the pronoun 'he' in the second sentence is anaphoric on 'John', and the event reported in the second sentence, the sneezing, occurs immediately after the event reported in the first. The general idea is that, as far as possible, the new sentence should concern the same *topic* as the background content: the same or closely related objects, the same or a closely related spatio-temporal region.

In this case, achieving Contiguity is a matter of saturation: fixing values of indexicals and temporal parameters. However, it is also part of the present account that Contiguity level coherence is *generally* effected by saturation, not by enrichment. Saturation does not strictly speaking *raise* the degree of coherence, since before saturation (in the cases where it is needed), there is no proposition. Rather, in the choice between different candidate saturations, a preference is given to those that deliver a higher degree of coherence.

⁸This discourse satisfies the relation *Inform-accident-and-mention-fruit*, proposed by Knott and Dale to illustrate the arbitrariness of unmotivated discourse relations.

⁹Thanks here to Kathrin Glüer.

In Levinson's scheme, fixing the pronoun in (22) as coreferring with the name is again an I-implicature, according to the principle of *preferred local coreference* (Levinson 2000, p. 177).

The *Resemblance* type is taken over from Hobbs and Kehler, comprising the same subtypes. For an example that exhibits both contiguity and resemblance, (23) is a sentence reporting the heights of two bridges in Stockholm, both built in the mid-1930s.

(23) Tranebergsbron has a height of 25.2 meters, Västerbron a height of 26 meters.

The objects are close in space (as far as bridges go), and the two conjuncts instantiate the *Parallel* relation, each concerning the height of a bridge (of about the same age).

States or events related by Contiguity are merely described as close; there is no implication of any *influence*, causal or otherwise, of one on the other. This does not exclude dependence on a joint cause, or conspiring to a joint effect, but these are no more than two ways in which the states or events may be close. Parallel constructions often indicate a shared cause or contributions to a joint effect, but need not do so. With contiguity and resemblance, there must be some connection between objects, and some connection between properties, of the old and the new content, but there need not be any connection between the *combinations* of objects and properties, i.e. not between the *propositions*. This changes with the third and fourth degrees.

The third degree of coherence, *Possibility*, comprises relations where the new content represents a fact, state or event that has been *prepared* or *made possible* by what is represented in the background. This comprises what has been called "occasion" by Hobbs and Kehler. An example from Kehler (2002, p. 14):

(24) A flashy-looking campaign bus arrived in Iowa. Soon afterwards George W. Bush gave his first speech of the primary season.

Here the assumption of coherence naturally ties the two events together, but a contiguity assumption would deliver little more than the conclusion that Bush's first speech was delivered in Iowa. In this case, we tend to read more into the connection: that the arrival of the campaign bus was part of the preparations that enabled the speech mentioned in the second sentence.

Occasionally, the order of reports is reversed, so that the enabling event is mentioned second:

(25) John managed to look in through the window. He had climbed up a wooden ladder.

The coherence assumption delivers a reading where the climbing enables the looking.

Finally, *Necessity* comprises relations where the new content represents a state or event that in some respect is a *consequence* of, or *necessitated by*, what is represented in the background content, or vice versa. The consequence may be causal, as in:

(26) John turned left. He was facing the library building.

Here it is overwhelmingly natural to read the discourse as reporting implicitly that John faced the library building *as a consequence of* just having turned left.¹⁰ So the cause-effect relations belong here.

Teleological explanations also belong to this type:

(27) The man took out a knife. He was going to cut the rope.

Here the coherence assumption ascribes cutting the rope as the *purpose* of taking out the knife. The second sentence thereby provides a reasons explanation of the event reported in the first.

A third kind of *Necessity* relation is *justification*, for instance in reasoning. This is exemplified in

(28) The butler did it. It can only have been the butler or the gardener, and the gardener was on the phone.

It is implicit here that the gardener could not have done “it” (e.g. murdered the baron) because of having been on the phone, and therefore one can infer by a disjunctive syllogism that it was the butler.

This completes the brief sketch of the categories. I am not in a position to define them more explicitly. There is, however, a further question about the strength ordering. In what sense are some of these relations *stronger* than others? Ultimately, an ordering of this kind will be justified by success in predictions of (saturation and) enrichments. And ultimately, if it *does* yield successful predictions, the ordering also tracks cognitive features of the human mind, its tendency to look for connections and patterns in order to systematize and make sense of its environment.

But perhaps something can be said to motivate, in an abstract fashion, the proposed ordering over alternatives. We can set up a sequence parallel to the scale of coherence strength, as follows:

¹⁰As pointed out by a referee, we could also have the opposite relation: that John was facing the library *explains* why he turned left.

Scale of connection types

- 0) No connection
- 1) Object connections
- 2) Property connections
- 3) Weak propositional connections
- 4) Strong propositional connections

The coherence strength scale is related to the order of connection types in the following simple way: when coherence is satisfied to a particular degree by a linguistic unit (a sentence, or discourse, or utterance in context), it gives us information of a type at the corresponding degree of connection types.

Thus, with coherence at the contiguity degree, we do get new information about the same object or objects that were known from the background content, or about objects closely related in space and time. In (21), we get new information about two objects, the sofa and the cat, that are spatio-temporally close to the table mentioned in the first sentence. In (22), the new information concerns the same object, John, at an immediately following moment in time.

The assumption of contiguity raises the expectation that new information will concern closely related objects, but raises no expectations about what properties or relations will be attributed to these objects. In a sense, then, the assumption of contiguity raises expectations concerning what we will next be informed *about*, the topic of the information, but not what we will be informed of concerning that topic, the comment of the information.

With coherence at the resemblance degree, we do get information about related objects, but also information about related properties. That is, the property (or relation) attributed in the new linguistic unit is to be closely related to the property (relation) that is already salient in the background content. The property might be the same, or similar in some respect, e.g. as two sub-species, as in (12) (where two political activities are subsumed under a more general category). Compare

- (29)
- a. Twenty years ago my brother bought a car, and last week my boss bought a car, too.
 - b. Yesterday, my brother bought a car, and my sister had eggplant for dinner.
 - c. Yesterday, my brother bought a car, and my sister went to Greece with her family.

There is an intuitive lack of coherence in (29a), since the distance in time makes the sameness of activity seem insignificant. This is a lack of contiguity. There is also an intu-

itive lack of coherence in (29b), since the two activities, of buying a car and having dinner, seem both unlike (not least in expected importance) and not plausibly related. There is contiguity, but a lack of resemblance. Both (29a) and (29b) appear eminently pointless. (29c), by contrast, exemplifies both contiguity (concerning brother and sister, same time) and resemblance (both activities are examples of expensive private consumption), and seems completely normal.

The expectation of resemblance raises the expectation that the new information will concern related objects but also in related respects. But it does not raise the expectation that the *propositions* will be related, i.e. that the (alleged) *fact* that related objects have related properties is somehow connected to the (alleged) background facts.

When we move up to the third and fourth degrees of coherence, we add propositional connections. These can plausibly be characterized in terms of conditional probabilities, or in other words, the rate in which a situation represented by the added content [background content] occurs in connection with situations represented by the background content [added content]. Here there are two parameters that are crucial: both the conditional probability in itself, and the difference between the conditional probability and the prior probability of the added content [background content].

To illustrate. In cause-effect relations from background content to new content, the *causing* state more or less necessitates the state or event represented by the new content. The conditional probability of the truth of the new content, given the truth of the proposition that adequately represents the causing state is very high, close to 1. For example, in example (26), John is assumed to be in a certain location, turned in a certain direction. If under these circumstances John turns left, he will be facing the library building; it is almost bound to happen (unless in the library building in the meantime has been demolished or a huge screen has been placed in front of it, etc.).

Of course, John will not very often face a library building in general when turning left. In this case, only a very small part of the relevant background content is actually given in the immediately preceding part of the discourse, i.e. the sentence 'John turned left'. The discourse given may be a fragment of a larger discourse where causally relevant information is provided. It may also be that *not* enough causally relevant information is given. In such a case, as a matter of text interpretation, the reader who infers a cause-effect relation between the events/states reported in the two sentences of (26) *accommodates*, i.e. *infers* that the situation is such that John is bound to face the library building on turning left.¹¹

¹¹ For the original proposal of accommodation, concerning presupposition, see Lewis 1979.

It will often be the case that the background content that is explicitly given does not by itself provide a very high conditional probability. To take an example:¹²

- (30) Johnny has been getting too much homework in elementary school this year, and
[as a result] his parents are taking the school to court.

The information given in the first conjunct does not by itself give the truth of the second conjunct a very high conditional probability: most cases of excessive homework do not result in a lawsuit. Note, however, that we would expect a big difference between this conditional probability and the prior probability of the second conjunct; that is, we would expect the rate of lawsuits by parents of kids with excessive homework to be significantly higher than lawsuit by parents of school kids in general. This would be explained by the assumption of a causal connection. Tacit background conditions (choleric parents, arrogant headmaster, etc.) that raise the conditional probability to a value close to 1 can be added by the imaginative interpreter, but this is not needed to appreciate the difference itself.

We have a similar situation in cases of teleological explanation: given a certain goal or desire, together with certain obtaining circumstances and beliefs about those circumstances, the agent very probably decides on a particular action to reach the goal. In the example (27), we infer that the man desired to cut the rope and believed that an efficient step to reaching this goal was to take out his knife. Part of the reason why we infer this belief is that we assume a basic shared theory of the world, including the proposition that in general knives are good tools for cutting ropes. We should not expect teleological explanations to give in general as high conditional probabilities as causal explanations, for the simple reason that there often are alternative means of reaching a particular goal.

In justification relations we do reach conditional probability 1 when the justified content is a conclusion that follows logically from the premises. In non-logical cases it will be lower, but that it is high is part of the very idea of justification.

Turning to the *Possibility* type, conditional probabilities will in general be significantly lower, but the raising of prior probabilities still significantly high. When a state or event needs some preparation or enabling, then the prior probability that the event occurs is comparatively low, and when the enabling event has occurred, it is plausibly higher. In example (25), we infer a situation where the window is too high above the ground to look in through for a person standing on the ground. The prior probability of the event of looking in through a window that is too high is low, while the conditional probability of looking in given that an enabling event has taken place is considerably

¹²Thanks for the example to an anonymous referee.

higher. Often, cases of enabling will also be cases of teleological explanation, since what is made possible is often a goal, and what is enabling is often believed to be enabling. Nonetheless, the categories are conceptually distinct.

We cannot distinguish Possibility and Necessity probabilities on a case by case basis. The idea is only that we have *generally* higher conditional probabilities, and *generally* higher increases from prior probabilities, in the Necessity type than in the Possibility type, and this at least in part justifies treating Necessity type coherence as stronger than the Possibility type.¹³

5. Enrichment Theory

The main idea of the present proposal is that pragmatic enrichments occur as a means of *raising coherence*.¹⁴ The content of an utterance is enriched to the point of instantiating a *stronger* coherence relation with already given background content. The proposal can be summed up as follows:

(PPE) *Principle of Pragmatic Enrichment*

(Free) pragmatic enrichments strengthen the coherence between background content and new content to the highest available degree.

The (PPE) principle brings out the main idea of the proposal, but more precision is needed for stating the Enrichment Theory. We can spell it out a bit more as follows below. Here, a *potential enrichment* is conceptual unit that is a candidate for being added to a given prior (= prior to enrichment) content. If the enrichment in question is free, the *prior* content is a proposition, and the resulting enriched content is again a proposition. We say an enrichment *occurs* just in case an enrichment is in fact pragmatically added in interpretation.

Enrichment Theory:

- i) Free enrichment occurs iff an available potential enrichment gives the *enriched* new content a higher degree of coherence—at least 2—with the background content than the new content prior to enrichment.
- ii) If α and β are two potential enrichments, α and β are equally available, and

¹³In the scheme of coherence relations of Sanders, Spooren, and Noordman 1992, p. 6, the category of *causality/implication* corresponds to *strong* connections and the category of *addition/conjunction* corresponds to *weak* connections. This is said to be a matter of theoretical intuition.

¹⁴Leth (2010), chapter 4, argues for applying the theory of discourse relations (rhetorical relations) to pragmatic phenomena, but also for the view that the theory concerns intrinsic properties of texts rather than the psychology of communication. These two ideas are in some tension, it seems to me, and Leth does not try to develop any systematic application to pragmatics.

- α raises coherence more than β , or if α and β raises coherence equally much but α is (clearly) more available than β , then β does not occur.
- iii) If potential enrichments α and β are available, equally available, and raise coherence equally much, there is pragmatic ambiguity: both α and β can result from interpretation. Also, if α is (clearly) more available but β raises coherence more, then there is (or at least may be) pragmatic ambiguity.
 - iv) If α and β both are available non-free enrichments, α is preferred over β if α give rise to a higher degree of coherence than β —at least 1.
 - v) A potential enrichment is *available* only if it satisfies requirements of *plausibility* and *simplicity*.

The theory imposes two constraints on availability: *Plausibility* and *Simplicity*. Plausibility has two aspects. On the one hand, the Plausibility constraint requires that the enriched proposition is plausible, by the standards of interpreter, or the conversation. On the other hand, the proposition describing the coherence *connection* must also be plausible. For instance, with respect to the opening example (1), the assumed coherence connection, e.g. that having the key enabled her to open the door, must be plausible as well. In general, the enriched proposition is plausible if it is not in conflict with background theory, while the *connection proposition* is plausible only if it is a *consequence* of, or made highly probable by, the background theory and the background content together. These two aspects of plausibility must be combined.

That the enriched added content, or the connection proposition, is true, is neither necessary nor sufficient for their pragmatic role. The effect on interpretation depends only on whether they are *taken* to be true. What is plausible of course depends on the background empirical theory of the interpreter, and may vary between interpreters, especially if they belong to very different communities. The Plausibility requirement therefore imposes a context dependence on enrichments that is more radical than more standard versions in semantics.

The Simplicity constraint is difficult to spell out in a way that is both reasonable and precise. It concerns the cognitive accessibility of potential enrichments. That a potential enrichment must be cognitively accessible to occur, is obviously correct, and also pretty much analytic. It does not add a constraint that can be used for prediction as long as our knowledge of what is accessible coincides with our knowledge of which enrichments have in fact occurred.

Simplicity is a natural parameter of accessibility. What is very complex will not be accessible in rapid on-line interpretation (interpretation that is ongoing as the utterance unfolds). But what is a natural measure of complexity of contents? Contents themselves

do not have a size.¹⁵ What we can measure easily is the size of natural (without ad-hoc abbreviations) English expressions. A content is intuitively more complex, at least by and large, if it requires an expression (surface structure) of greater length. Virtually all enrichment examples in the literature can be expressed by phrases of less than ten words, often between one and five. Accordingly, expressibility by phrases of a size of that magnitude is a natural restriction. Other things equal, lower complexity makes for higher availability.

We can, however, say something more directly about content. The potential enrichment should not involve reference to any object, event or fact that is not already represented in the background content or in the added prior content, unless linked by some contiguity or resemblance relation. We can label this the *No New Entity* restriction. The following exemplifies a violation of it:

- (31) a. Uncle Otto prefers Claret.
b. Uncle Otto prefers Claret [*to rotten fish*].

The (non-free) enrichment in (31b) is clearly small enough and yields a plausible proposition (given what we know about normal human taste). But unless rotten fish has been raised to salience in the context of utterance (in which case it should be added to the background content), the enrichment is not easily available. If the speaker *meant* rotten fish, the hearer would need a great deal of ingenuity or luck to get the right enrichment. Without a No New Entity restriction, the interpretation possibilities are virtually endless. Predictability would go down.

ET is much in the spirit of Asher and Lascarides, who propose the following principle (p. 230):

Maximise Discourse Coherence (Asher and Lascarides)

The logical form for a discourse is always a logical form that's maximal in the partial order of the possible interpretation; i.e. those which are consistent with compositional semantics and monotonic constraints on anaphora resolution.

Asher and Lascarides define a technical notion of Maximal Coherence which includes an idea of a partial order on discourse relations (rhetorical relations) of producing a “closer connection” or a “better coherence” (p. 231). This idea is not elaborated, however, and the purpose of Asher and Lascarides is not, as here, to employ the ordering for explaining pragmatic enrichment.

¹⁵For propositions we consider a proposition *p* more complex than a proposition *q* if the set of *p*-worlds is a proper subset of the set of *q*-worlds, i.e. if *q* follows from *p*. This is not general enough, however. Between two propositions neither of which entails the other, we would want the one associated with the *smaller* set to be more complex, but the two sets are likely to have the same infinite cardinality.

6. Applications of Enrichment Theory

In the examples below the a. part is the sentence or discourse uttered, usually including expressions both of background content and prior new content, while the b. part has a phrase added in square brackets that expresses the intuitive enrichment. I shall also use italicized sentences in running text to refer to the propositions they express (given an imagined context). I shall take several examples from Carston 2002.

Enrichments are operations on content. They can operate on proper parts of complex contents, and that is the rule rather than an exception. For instance, as in (1) below, what is enriched is the VP meaning *opened the door* rather than the meaning of the sentence ‘she opened the door’. A formal representation would employ a schema of structured meanings in order to provide the proper arguments of enrichment operations.¹⁶ To keep things simpler here and focus on the main ideas, I shall simply make use of the syntactic surface structure of the example sentences to indicate where the enrichment goes.

Typically, the enriched proposition, such as *she opened the door with the key that he had handed her* entails the unenriched one, in this case *she opened the door*. But since enrichment can operate on a part of the proposition, this relation can be reversed. This happens when the enriched part is dominated by *negation*, i.e. is the semantic content of an expression that occurs syntactically in the scope of negation. If the enriched content q necessitates the unenriched content p , then $\neg p$ necessitates $\neg q$. This is relevant in some of the examples below.

We start with some examples of coherence of the third degree, i.e. of Possibility level coherence. The general idea here is that after the enrichment, the action described by one sentence or conjunct is made possible by a state or event described by the other, and that this is not the case before the enrichment. The first example is the opening example from Carston 2002, p. 71:

- (1) a. He handed her the key and she opened the door.
- b. He handed her the key and she opened the door [*with the key that he had handed her*].

In (1), as in other cases, we assume that a Contiguity level coherence is achieved already at the saturation stage. The pronoun ‘she’ is therefore read as coreferring with ‘her’ in the first conjunct, and the event described in the second conjunct is understood as taking place *immediately after* the event described in the first conjunct. We can achieve a Possibility level Coherence with the enrichment expressed in (1b): having the key (as a

¹⁶See Pagin and Pelletier 2007 for such an account.

result of Contiguity: the state of having the key is assumed to still obtain) *enables* the referent of ‘her’ to open the door. In this case there is no plausible Necessity level enrichment. Being given the key is not plausibly a sufficient *cause* of her opening the door (she might not have wanted to), and so the truth of the first conjunct is not a premise from which the truth of the second conjunct can be plausibly inferred, even after the enrichment.

Carston herself describes the relation between the contents of the conjuncts as an *instrument relation*. This is of course correct, and typically instrument relations *are* enabling relations. Another example of the same pattern:

- (32) a. He found a dollar on the floor and bought a chocolate bar from the vending machine.
 b. He found a dollar on the floor and bought a chocolate bar from the vending machine [*with the dollar that he had found*].

The next example is explanatory and teleological:

- (33) a. The man took out a knife. He was going to cut the rope.
 b. The man took out a knife. He was going to cut the rope [*with the knife he took out*].

After the enrichment, the truth of the second sentence provides a teleological explanation of the truth of the first. Given are background theory of knives and ropes, both the enrichment proposition and the connection are highly plausible, and we also get coherence of the fourth degree.

Could it be read otherwise? Could we take the truth of the first proposition to *enable* or *cause* the truth of the second? The progressive is in general used to describe a state or activity that takes place over an interval of time. The past progressive of the second sentence of (33a) describes a state which will end in a rope-cutting. That it will so end can be determined causally as a *tendency* that will make a rope-cutting occur in all possible developments (branches of time). It can also be determined in hindsight by being described *after* the cutting event, when only rope-cutting branches of time are left to be considered.

To get a left-to-right enabling or causal interpretation would require a so-called *inceptive* reading (Dowty 1986) of the second sentence: the ending-in-rope-cutting state *starts* immediately after the event of taking out the knife. It is, however, known that such readings of sentences in the progressive tense are almost impossible, whatever the explanation (cf. Dowty 1986, p. 38). Without explicit indication to the contrary (such as

‘Suddenly’, or ‘From that time on,’), the state will be taken to have obtained from some time *before* the reference time.

Moreover, the availability of such interpretations is very low by present criteria. For it is implausible that taking out a knife would either *cause* such a state to obtain or *make it possible*. Taking out a knife enables the cutting, it does not plausibly enable an ending-in-a-rope-cutting *state*. From the retrospective perspective, that state started much earlier. And in the *tendency* reading, the determination that rope-cutting will be in the future would plausibly depend on the *intention* of the agent to cut the rope, and it is very implausible that taking out the knife would cause this intention. All in all, a left-to-right coherence reading is ruled out.

The next example is from Carston 2002, p. 138:

- (34) a. He ran to the edge of the cliff and jumped.
b. He ran to the edge of the cliff and jumped [*over the edge of the cliff*].

Here Contiguity of person is already assured since the subject of the second conjunct is unpronounced, and therefore must be anaphoric on ‘he’ in the first conjunct. Again, saturation assures Contiguity of time: the jumping occurs immediately after reaching the edge. The intuitive enrichment is plausible and makes running to the edge an enabling condition of the jumping. Simply *jumping up and down* is something one does not need the location at the edge for doing. Similarly, *jumping away from the cliff* does not require being at the edge either, since only the direction matters.¹⁷

If we add a specification of where the jumping is done, the enabling condition is met:

- (34) c. He ran to the edge of the cliff and jumped [*up and down at the edge of the cliff*].

However, on the present scheme of interpretation, spatio-temporal determination takes place at the saturation stage, and does not enter into enrichments. That rules out (34c). We therefore reach a Possibility level coherence with (34b), which is predicted.

However, it is also easily inferred that jumping over the edge was the *purpose* of the running. Understood that way, we have Necessity level coherence; the running is teleologically explained by the jumping over the edge. The same enrichment serves both ways of strengthening.

Is this a case of concealed ambiguity? That is a reasonable result, in this case. It is reasonable because we must weigh the higher degree of coherence that we get from the teleological connection against the greater plausibility of the weaker, enabling connec-

¹⁷Thanks to Jeff King for the suggestion.

tion. The plausibility of the enabling connection is greater, since it is a consequence of our common background theory that a person can jump over an edge he is at, while it is difficult or impossible from a distance to the edge, whereas it is probable but less certain that a person runs to the edge with the purpose of jumping over the edge.

As far as I can tell, there is also no simple enrichment that either ensures or rules out the teleological interpretation. For instance

- (34) d. He ran to the edge of the cliff and jumped [*over the edge of the cliff, as he had intended*].

is dispreferred because of the greater complexity (which is borderline too high). Moreover, it violates the general pattern that enrichments characterize event or facts that are represented in the contents they enrich. In (34d), part of the enrichment is only concerned with what is represented in the *first* conjunct.

Two more examples of ambiguity. Wilson and Matsui (1998, p. 199) offer

- (35) a. I ran to the classroom from the playground. *The children* were making too much noise.
 b. [...]. The children [*of the playground*] were making too much noise [*for me to stand*].
 c. [...]. The children [*of the classroom*] were making too much noise [*for me not to interfere*].

In this case we can get an explanation of the running whether ‘the children’ refers to the children in the classroom or to the children at the playground. A non-free enrichment is required to fix the reference of the definite description. At the same time, a further, again non-free, enrichment (too much for what?), is required to get both a proposition from the second sentence as well as the causal connection between the screaming and the running. In this case, we have two alternatives which seem equally available, and which both deliver coherence of the fourth degree. Ambiguity is predicted.

The second ambiguity example is due an anonymous referee:

- (36) a. Mary embarrassed John, and Betty made fun of him.
 b. Mary embarrassed John, and Betty made fun of him [*because of that*].

As pointed out by the referee, there are two coherence construals of (36). The conjuncts satisfy a Parallel relation: two bad things happened to John, and they are described in the conjuncts, respectively. We therefore have a Resemblance level coherence without any enrichment. But we can also read into the second conjunct that it describes an *effect* of

the event described in the first. We then get a Necessity level coherence. This requires an enrichment that makes clear that the first event is the topic of ridicule in the second (otherwise there is no plausible connection). In this case, it seems, the greater availability of the first reading (no enrichment is needed) is balanced by the higher degree of coherence of the second. Ambiguity is then predicted.

The approach can of course be applied to dialogue. This adds the complication that the background content may be given in a question. The following example is from Carston 2002, p. 40:

- (37) a. A: Do you want to go to the cinema?
 B: I am tired.
 b. B: I am tired [*enough for not wanting to go to the cinema*].

In this case, B's answer implicates by relevance that B does not want to go to the cinema. We can see the enrichment either as resulting in a *justification* of the preference not to go to the cinema, or as representing a *cause* of the lack of a desire to go. In either case, the enrichment produces a Necessity level coherence between the enriched content and the implicature.

This example seems not to fit the schema that the enriched new content stands in a coherence relation to the background content, since the implicature *that B does not want to go* seems not to be part of the background. However, taking into account that the background content is given by a question, we must first determine what that content is. It is a standard ingredient in the semantics of questions (cf. Groenendijk and Stokhof 1994), that the content is *the set of possible answers* {*yes, no*}, i.e. the set of propositions

{*that B wants to go, that B does not want to go*}.

The background content then is a set of alternative background propositions. The enrichment establishes a Necessity level coherence with a member of that set. The resulting implicature is simply that member of the set that is selected by instantiating the coherence relation. And the enrichment is the enrichment needed for attaining the coherence relation. That some member of the background set must be selected is determined by the asking of the question, which generates an expectation that the answer by B will serve that function.

The background content can also be a proposition that is salient to both speaker and hearer, e.g. because of facts known by observation. The following is a classic example from Bach 1994, p. 278. A mother is talking to her child, who has had a cut in a finger:

- (38) a. You are not going to die.
 b. You are not going to: die [*from that cut*].

In this case, the salient background content is the content of the mutually shared knowledge that the child has a cut in the finger. The enrichment operates on a *part* of that content: it enriches *die* to *die from that cut* or *die because of that cut*. The unenriched utterance content has a Contiguity level coherence with the background content in so far as it concerns the same child. Perhaps it may also be understood as referring to the *near future* of the child even without the enrichment (in which case the utterance is true rather than false). The enrichment raises the coherence between *part of* the utterance content and the background content to Necessity level coherence, due to the causal connection.

This example has the complication that the unit of content that satisfies the coherence relation, i.e. the proposition that the child *is* going to die (soon) from the from the cut, is highly implausible. The entire proposition asserted, on this construal, (38b), is highly plausible, but does not itself instantiate a causal relation to the background content. That would have required an enrichment operating on the entire unenriched utterance content:

- (38) c. [*Because of that cut,*] you are not going to die.

This enrichment would also lead to Necessity level coherence, but is ruled out because it is wildly implausible.¹⁸

The theory also generates predictions about the interpretation of single sentences because of the interaction between content parts. Consider

- (39) a. A tall man picked up the book.
 b. A kind man picked up the book.

There is no obvious connection between being tall and picking up a book. By contrast, there is an obvious connection between being *kind* and picking up a book: picking it up out of kindness. In the parallel discourse (40a), we would predict the enrichment in (40b).

- (40) a. There was a kind man. He picked up the book.
 b. There was a kind man. He picked up the book [*in an act of kindness*].

¹⁸The dissociation here between what satisfies coherence (the negated proposition) and what satisfies plausibility (the negative proposition) instantiates a pattern that merits further investigation, but it falls outside the scope of the present paper. Note that negative relations do figure in the Hobbs-Kehler scheme.

The reason is that kindness, which is or involves a disposition to do kind things, is part of the causal background of the picking-up, provided the picking-up was an act of kindness. The enrichment in (40b) therefore raises coherence to Necessity level, since after the enrichment the first sentence describes an ingredient in the causal explanation of the event described in the second.

We get the same effect, it seems to me, in (39a). That is, we have

- (39) c. A kind man picked up the book [*in an act of kindness*].

The information that the man in question is kind is given as an attribute in the noun phrase of (39b), and the next property predicated of the man is that he picked up the book. There is a chance of making these two predications cohere causally by enriching the second. There is no corresponding chance with (39a).

The pattern is pervasive. Another example.

- (41) a. The temperature has risen to a dangerous level.
 b. The temperature has risen [*from a non-dangerous*] to a dangerous [*ly high*] level.

Strictly semantically, (41a) is true if the temperature has risen and after the rising is at a dangerous level, even if it is dangerous after the rising because it is *low* and the level was even more dangerous *before* the rising. To make more sense of it as a report, we read into it that the danger depends on the fact that the temperature is *high*, that it was not high enough to be dangerous before the rising, and that the danger therefore is *caused* by the rising.

In this case, what we enrich are parts of the meaning of the verb phrase by *adding* the meaning of a new prepositional phrase and also enriching the original prepositional phrase. We thereby raise the coherence between the main verb content (rising) and the content of the prepositional phrase to Necessity level.

We again find the same pattern in Rumelhart's (1979, p. 78) classic example:

- (42) a. The policeman raised his hand and stopped the car.
 b. The policeman raised his hand [*as a signal*] and stopped the car [*by the signal to the driver of the car*].

There seems to be one unique plausible enrichment, given our knowledge of traffic and police, which raises coherence to Necessity level. Firstly, we would like it to be relevant that the person is a policeman, and this we can achieve, since it is known that being a policeman (in uniform) *enables* a person to give orders to drivers by means of signalling.

Secondly, we would like to get a causal connection between the raising of the arm and the stopping of the car, and the unique way this is plausibly done by a policeman is exactly by raising his/her hand as a signal. We therefore get a Necessity level coherence between the conjuncts of the verb phrase, and a possibility level coherence between the noun phrase and the enriched verb phrase. This is all perfectly consistent with what Rumelhart himself says about the example, although he does not emphasize the search for coherence, only the appeal to background theory (“general knowledge”).

This concludes the examples of application. We shall next look at some related but arguably different phenomena.

7. Contrasting phenomena

For contrasting pragmatic phenomena, consider again proposed cases of *loosening* and *semantic shift*.

- (9) The ATM swallowed my credit card.

(page 9 above). The proposed loosening of the concept of *swallowing* achieves plausibility (a machine can swallow in the loose but not in the narrow sense). But there is no salient background content with which coherence is increased by this operation. Nor is it easy to see what possible enrichment would give the intuitive result. This comes out as a different phenomenon both by intuition and theory.

That case is different from

- (43) a. France is hexagonal.
b. France is [*roughly?*] hexagonal.

(Carston 2002, p. 41, going back to Austin.). There is a possible enrichment of (43a) that does deliver the intuitive content, as in (43b). This is of course not predicted by the present account, since no increase of coherence takes place. What does take place is an increase of plausibility.

Intuitively, we have the same phenomenon in (9) and (43), a loosening of a concept by relaxing its application conditions. The main difference between the two cases is that *hexagonal*, being a geometrical concept, has only one dimension of loosening, by allowing geometrical approximation. The concept of *swallowing* can be loosened in different respects. That is why it is unclear what is meant locutions such as “roughly swallowed” or “approximately swallowed”. I suggest, therefore, that the possibility of reaching the intuitive reading of (43a) by means of the enrichment in (43b) is really an artifact of the geometrical character of the example. What is taking place also in (43a) is precisely a

loosening of a concept, or the relaxing of a standard of precision (Lewis 1979), not an enrichment. These are distinct phenomena.

The same holds, I think, for semantic shift examples like

(10) The ham sandwich is sitting at table 20.

(page 9). It is *possible* in this case to reach an intuitive interpretation by means of an enrichment as in

(10') The ham sandwich *orderer* is sitting at table 20.

But this is only an increase of plausibility, not of coherence. And in this case too, I think we have to do with a different phenomenon than enrichment, a pragmatic mapping of content on content that in some cases can be *achieved* by means of enrichment, but in others not. Consider the following (real life) example, where A is a waiter, B a guest:

(44) A: Has he (*pointing to an empty chair next to B*) left?

B: No, he is sitting there (*pointing to (a person in) a chair across the table*). But *he* (*same pointing*) has left.

This is a double example of deferred ostension, a sub-species of semantic shift. The waiter A first points to an empty chair, but refers by 'he' to the person C that had been sitting there. The guest B replies, first using 'he' coreferentially, informing the waiter that C has moved to the chair across the table. B then points to C, or perhaps rather to the place where C is sitting, then refers to a third guest, D, who had been sitting where C now is located, informing A that D has left. There is no reasonable way in which these shifts could have been achieved by means of enrichment.

The upshot is that both loosening and semantic shift are different phenomena, as they also have been treated in the literature. Loosening and semantic shift mainly occur because of the drive to plausibility. This cannot be true of enrichment, because enrichment normally does *not* raise the chances of truth. Typically, a proposition *p* is mapped on a proposition *q* that *entails p*;¹⁹ in these cases there *cannot* be a shift from a false to a true proposition, and hence some other force must be operating.

Cases of so-called *narrowing* are more tricky, since it *prima facie* looks like they can always be achieved by enrichment. Both Recanati (example (11)) and Carston, in (45) (2004b, p. 639),

¹⁹As noted above, this is reversed under the scope of negation.

- (45) a. I have had a shower.
 b. I have had a shower [*today?*].

take implicit temporal reference to the day of utterance to be an enrichment. Recanati's example (having had breakfast *today*, and therefore not wanting anything to eat) can be accounted for by appeal to coherence, assuming that it is an enrichment, but Carston's example cannot, given that it works without any background whatsoever; that it concerns the day of utterance is simply the default interpretation (Assert the stronger?).

In contrast to Recanati and Carston, I take spatio-temporal determination to belong to the *saturation* stage, and so I don't want to treat the reference to the day of utterance as an enrichment at all. Rather, I prefer to see it as a value to a domain restriction on temporal quantification.

It is a common feature of several approaches to the semantics of tense that tense is associated with quantification over times, either time instances or time intervals (cf. Ogihara 2011), and such quantification often comes with a *restriction* on the domain of times. In particular, the present perfect, in the tradition starting with Reichenbach (1947), has been treated as involving existential quantification over times (cf. Portner 2011). On the domain restriction approach, there is a restriction *parameter* that needs a contextually given value. Providing that value would belong to saturation rather than to modulation, but it is still part of the pragmatic process. The analysis would be along the lines of

- (46) $\exists t \subseteq \text{TODAY} (t < \text{NOW} \ \& \ \text{I shower at } t)$

Here the subset relation says that the time/interval of shower-taking is included in TODAY. TODAY is the value of the restriction parameter.

A semantic tense apparatus is needed irrespective of potential enrichments, and with a theory that employs quantification plus restriction, no extra enrichments are required, since they can be handled by a feature that is built into the semantics itself. This speaks in favour of the semantic alternative.

It is also possible that coordination phenomena can favour one option over the other. Consider the following example. Before leaving on Wednesday for a business trip, Betty makes John promise to call the insurance company about a strange noise in the car. On Friday John calls the insurance company and then has a shower. On Saturday morning John has a shower and is then eating breakfast when Betty calls him.

- (47) a. Betty: Are you ready to go? Did you call the insurance company?
 b. John: I have had a shower and called the insurance company.

- c. John: I have had a shower [*today ?*] and called the insurance company [*during the recent couple of days ?*].

Given the context, (47b) seems an odd thing to say, since the relevant shower takes place on the day of utterance but the calling before the shower on the previous day. The likely reaction is: “Did you call them today? Today is Saturday”. That it is odd is predicted by the semantic theory. For if the restricted temporal domain includes the calling, then it also includes the first shower, and so the assertion would have been true even without a shower on the day of utterance. It is true but does not give the information that a shower has been taken on that same day. If the domain is restricted to the day of utterance, on the other hand, the assertion is false, since no calling takes place that day.

On the enrichment account, however, there is no obvious reason why the content should not be enrichment as in (47c), which would make (47b) perfectly fine. If (47b) is perfectly fine, this speaks for the enrichment account, and if it isn’t, the enrichment account is in trouble.

Another bone of contention with Carston concerns the narrowing of concepts. The following is an example from Carston 2004b, p. 639:

- (48) a. It’s going to take time for these wounds to heal.
b. It’s going to take [*considerable*] time for these wounds to heal.

Again, without any background content to cohere with, (48b) does seem to give the appropriate content of an assertion of (48a), and so we seem to have free enrichment without coherence. It is free, since without the enrichment the statement (48a) is true, and trivially so.

The impression is misleading, however, for what counts as *considerable time* varies between contexts, especially depending on the topic of conversation (the drying of paint, the healing of wounds of various kinds, the building of a tunnel). Because of this, (48b) is, strictly speaking, not a free enrichment at all, for it does not map a proposition on a proposition.

It might *look* like an enrichment, since in (48b) we add the *word* ‘considerable’, and that word is supposed to have a relevant contextual interpretation. However, the point of specifying enrichments as *contents* is that the content added is fully specified in the theoretical context itself, and this condition is not met. (48b) is more of a schema indicating where an enrichment will go in particular contexts.

It is different if a contextual background is added to the example. For instance, a patient asking a doctor:

- (49) a. Can I leave before Christmas?
 b. It's going to take [*much*] time [*in relation to time until Christmas*] for these wounds to heal.

In this case, we do have a real free enrichment, one that does generate a new specific content, but it is also one that is generated by coherence raising (causal connection: time for wound-healing preventing leave).

Finally, consider the simplest kind of what Levinson calls “conjunction buttressing”, that is, reading more into conjunctions than the logical content.

- (50) a. He sat down and took off his shoes.
 b. He sat down and [*directly afterwards?*] took off his shoes.

The question is whether we have an example of enrichment. On the present account we don't, since spatio-temporal determination belongs to saturation. In this case we give the time parameter of the second event a value that is just after the time of the first event. This is standard, and in some theories built into the semantics itself (e.g. in Kamp and Reyle 1993). On the present account, it is a pragmatic choice.²⁰

Now that the Enrichment Theory is in place, I shall make brief comparisons with Relevance Theory and with Levinson's theory of I-implicatures.

8. Relevance Theory

Relevance Theory (with Sperber and Wilson 1986/1995 as the basic work) is a collection of principles concerning communication and cognition. These principles are meant to describe the mechanisms of communication, among other things, the generation of enrichments. The principles have been presented in several places, with slight variations in wording. The following summary is taken from Wilson and Sperber 2004.

²⁰What is the general explanation? According to Carston, the explanation comes from the need of efficiency in processing: “It is just a fact about our minds that we find it easiest to process information when it is sequentially ordered. Imagine that the instructions for assembling a model aeroplane or for knitting a jumper were presented to you in random order or that the frames of a comic strip were not sequentially ordered” (Carston 1995, p. 233).

However, the main contrast is not, as in Carston's examples, between chronological order and random order, but e.g. between chronological order and reversed chronological order (then in general read the second conjunct as describing the earlier event). Still, I think it is right that processing efficiency is the key. We read from left to right, and since it is inefficient to *backtrack* in interpretation, we try to make pragmatic operations on the second conjunct, not on the first after having seen the second.

Also, in general, the presentation of the *later* event depends on the features of the *earlier* event. So there are in general pragmatic operations on the presentation of the later event that depend on features of the earlier event. That it is presented later then saves us from backtracking. At least, that would be my attempt at an explanation outline.

There are two cornerstones of the theory. One is a general principle of cognition, the other of communication (2004, pp. 610-12):

General Principle of Relevance

Human cognition tends to be geared to the maximization of relevance.

Communicative Principle of Relevance

Every ostensive stimulus conveys a presumption of its own optimal relevance.

The idea of relevance is technical, and explained as follows (2004, p. 609):

Relevance of an input to an individual

- a. Other things being equal, the greater the positive cognitive effects achieved by processing an input, the greater the relevance of the input to the individual at that time.
- b. Other things being equal, the greater the processing effort expended, the lower the relevance of the input to the individual at that time.

The presumption of relevance that governs interpretation is then spelled out as follows (2004, p. 612):

Presumption of optimal relevance

- a. The ostensive stimulus is relevant enough to be worth the audience's processing effort.
- b. It is the most relevant one compatible with communicator's abilities and preferences.

The interpretation process that is to result from these principles is described as follows (2004, p. 613):

Relevance-theoretic comprehension procedure

- a. Follow a path of least effort in computing cognitive effects: Test interpretive hypotheses (disambiguations, reference resolutions, implicatures, etc.) in order of accessibility.
- b. Stop when your expectations of relevance are satisfied (or abandoned).

A list of tasks to be performed during interpretation is provided:

Subtasks in the overall comprehension process

- a. Constructing an appropriate hypothesis about explicit content (explicatures) via decoding, disambiguation, reference resolution, and other pragmatic en-

richment processes.

- b. Constructing an appropriate hypothesis about the intended contextual assumptions (implicated premises).
- c. Constructing an appropriate hypothesis about the intended contextual implications (implicated conclusions).

These subtasks are of course not specific to Relevance Theory (except for the use of the word ‘explicature’).

Now, on the face of it, Relevance Theory and the present Enrichment Theory look rather similar. Both have a *driving* ingredient: trying to achieve relevance, and trying to achieve coherence, respectively. And both have a *restraining* ingredient: minimizing processing effort, and restricting to high, or maximal,²¹ availability, respectively. The question then is whether these two theories are compatible. I think the answer is in one sense yes, and in one sense no.

The answer is yes in the sense that ET is compatible with the wordings of the principles spelled out above. There is nothing in these principle that blocks equating *greater cognitive effects* with *higher degree of coherence*. Assuming that the interpreting mind simply, according to its nature, *does* strive to maximize coherence, then these just are the cognitive effects it seeks. And assuming we have a cost of more processing effort in case a resulting enrichment, or an assumed connection, is less plausible, as well as if the enrichment is more complex, then lower availability *does* mean higher processing efforts.²²

That Relevance Theory is nominally compatible with ET is due to the lack of specificity of definitions of the central terms. The concept of *relevance* introduced depends on the concept of *cognitive effects*. Sperber and Wilson, in the second edition of *Relevance*, define it as follows (Sperber and Wilson 1995, p. 265:

Let us define a *cognitive effect* as a contextual effect occurring in a system (e.g. an individual), and a *positive cognitive effect* as a cognitive effect that contributes positively to the fulfilment of cognitive functions or goals.

²¹Since alternative enrichments are compared and ranked according to availability, it also makes sense to speak of striving for maximal availability.

²²Hobbs et al. (1993, pp. 80-87), provide an elaborate system for estimating the *costs* of assumptions in interpretation. They have an informativeness-correctness trade-off for assumptions. The specificity of information should match the linguistic form. For instance, an interpretation that fails to provide a referent for a definite noun phrase is expensive. Longer derivations of logical form are more expensive than short ones. Etc.

The total system of Hobbs et al. 1993, within AI, is designed to move from natural language sentences to first-order logical forms with the help of a knowledge base. The main idea is that the truth of the text is *explained* by the interpretation, typically by means of abduction. In the course of this process, “assumptions are made when necessary” (1993, p. 79). The system is not concerned with enrichments, since enrichments are typically not necessary to render the resulting interpretation true; rather it is typically true before enrichment if true after. Thanks for the reference to an anonymous referee.

This definition is admitted to be “rather vague” (1995, p. 266), the reason being that of wanting to include phenomena such as “the reorganization of existing knowledge” (1995, p. 266). The most typical cognitive effect is still *an increase in knowledge*.

It seems to me that the goal of increasing knowledge does not lead to relevance in anything close to the intuitive sense of the word. Consider a world with two individuals, A and B, and the following report.

(51) A walks in the park. He whistles.

There is an overwhelming tendency to read the pronoun in the second sentence as anaphoric on the name ‘A’ in the first. This reading is also the one which gives a higher degree of coherence and which is *intuitively* more relevant. However, we gain exactly as great an increase in knowledge from learning that A whistles as from learning that B whistles. The goal of increasing knowledge does not favour either reading and so does not deliver any prediction.²³

In their specific analyses, however, Sperber and Wilson have much more definite ideas about what does and does not satisfy relevance. They are also critical of coherence theory. Concerning the example

(52) Sara left Australia for England. She hates *the sandy beaches*.

(their (22b)), which is a bridging inference task,²⁴ Wilson and Matsui say the following:

Here, the comprehension procedure works smoothly. On the effort side, *Australia* is the most accessible antecedent: first, because it is picked out by the Expected Focus algorithm as the direct object of *left*; and second, because of the encyclopaedic association between *Australia* and *sandy beaches*. Moreover, the second part of the utterance, on this interpretation, could be expected to achieve at least adequate cognitive effects by answering a question implicitly raised by the interpretation of the first part: why did Sara leave Australia for England? To say that the first part of the utterance implicitly raises a question amounts to no more than saying that it provides easy access to a context in which an answer to this question would be relevant enough to be worth the hearer’s attention. There is no need to appeal to

²³ Maybe it requires more *effort* to switch to thinking about B. What would be the reason for believing this?

Several critics (including Bach (2010) and Davis (2010)) have complained, correctly in my mind, that these principles are not descriptively precise, and hence are weak in predictive power. Wilson and Sperber (2004, pp. 626-27) seem to implicitly acknowledge the point: “For example, the Cognitive Principle of Relevance suggests testable predictions only when combined with descriptions of particular cognitive mechanisms (e.g. for perception, categorization, memory, or inference)”. That is, predictive power is achieved when combined with *other* theories.

²⁴ Are *the beaches* beaches of England or beaches of Australia? According to the experiment reported by Wilson and Matsui, 100% answer Australia.

a coherence relation of Explanation. From the relevance-theoretic point of view, computation of coherence relations is a waste of effort unless it gives rise to otherwise inaccessible effects. The expectation of optimal relevance is enough on its own to account for the interpretation of this example.

It is not clear to me how RT delivers the desired reading. Wilson and Matsui mention three factors: expected focus (the direct object of ‘left’), encyclopedic knowledge (Australia is famous for sandy beaches), and the fact that the first sentence provides access to a context where an answer to the question why Sara left Australia for England would be worth the processing effort. Appeal to explanation over and above this, is a waste of effort.

However, we can change the example by means of the minimal replacement of ‘hates’ by ‘loves’:

(53) Sara left Australia for England. She loves *the sandy beaches*.

All the features that Sperber and Matsui appeal to for giving the desired reading of (52) are features also of (53), but (53) does not easily give rise to the corresponding reading. By common sense psychology, you don’t leave a place because of a feature of it that you love (unless you think you love it too much, but that is not at issue). Hence, the *plausibility condition* for the explanation relation isn’t met, and Enrichment Theory does not predict that in this example *the beaches* is enriched into *the beaches of Australia*, just because the Explanation relation *is* crucial.

Equally, by common sense psychology, you don’t move to a place just because it has a particular feature (sandy beaches), if the place you leave has that feature to an equal or higher degree. Hence, an Explanation relation isn’t plausible either for the enrichment into *the beaches of England*. The coherence of (53) does not reach Necessity level. It seems we *can* reach Possibility level coherence with an enrichment into *the beaches of England*: after the move, Sara has come to appreciate the sandy beaches of England, something that was made possible by the move itself. This requires the assumption that she did not appreciate the sandy beaches of England before the move. To the extent this is plausible, the enrichment “she [*now*] loves [*but didn’t before*]” is predicted.

We at any rate reach Contiguity level coherence by enriching into *the beaches of England*, since it is more coherent to refer to beaches close to Sara’s location at the time of utterance, i.e. after the move (it is also part of contiguity—given the present tense of the second sentence—that she she still is where she moved to). To the extent that it is part of the background theory that England *does* have sandy beaches, this is the enrichment of the description which ET predicts. Note that the enrichment for (53) is non-free, since

the definite description ‘the sandy beaches’ is incomplete.

Why Relevance Theory would not predict enrichment into *the beaches of Australia* again is unclear, given that the discourse (53) shares all the adduced features of (52).

All in all, the relation between Relevance Theory and Enrichment Theory is ambivalent: on an abstract level they seem compatible, while in actual applications provided they clash.²⁵

9. I-implicatures

There are several features common to Enrichment Theory and the neo-Gricean theory of I-implicatures as developed by Levinson, especially in Levinson 2000. As noted in section 2, I-implicatures are GCIs, generalized conversational implicatures, which according to Levinson intrude into *what is said*, and form a separate layer of meaning between the semantic content and particularized conversational implicatures.

The general idea of I-implicatures is derived from Grice’s second Maxim of Quantity (Q): *Don’t make your contribution more informative than is required* (Grice 1975, p. 26). This is adapted by Levinson as a general recipe for saving on explicit information: we can always assume that states and events are *normal*, as long as they are described in normal ways:

Heuristic 2

What is simply described is stereotypically exemplified (Levinson 2000, p. 32).

Implicatures of this kind are *informational* GCIs, I-implicatures. What Levinson calls the *I-Principle* is spelled out as follows (2000, pp. 114-15)

I-Principle

Speaker’s maxim: the maxim of Minimization. “Say as little as necessary”; that is, produce the minimal linguistic information sufficient to achieve your communicational ends (bearing Q in mind).

Recipient’s corollary: the Enrichment Rule. Amplify the informational content of the

²⁵Wilson and Matsui (1998, p. 199), criticize Coherence Theory for allowing that two alternative explanations can satisfy *the same coherence relation*, as in

- (35) a. I ran to the classroom from the playground. *The children* were making too much noise.
 b. [...]. The children [*of the playground*] were making too much noise [*for me to stand*].
 c. [...]. The children [*of the classroom*] were making too much noise [*for me not to interfere*].

In this case we can get an explanation of the running whether ‘the children’ refers to the children in the classroom or to the children at the playground. Wilson and Matsui point out that one interpretation or another may be preferred in different circumstances. This is perfectly true, but as long as none of the interpretations is preferred *by default*, not predicting one reading over the other seems to me a feature, not a bug. Enrichment Theory here predicts ambiguity.

speaker's utterance, by finding the most *specific* interpretation, up to what you judge to be the speaker's m-intended point, unless the speaker has broken the maxim of Minimization by using a marked or prolix expression. Specifically:

- a. Assume the richest temporal, causal and referential connections between described situations or events, consistent with what is taken for granted.
- b. Assume that stereotypical relations obtain between referents or events, unless this is inconsistent with (a).
- c. Avoid interpretations that multiply entities referred to (assume referential parsimony); specifically, prefer coreferential readings of reduced NPs (pronouns or zeros).
- d. Assume the existence or actuality of what a sentence is about if that is consistent with what is taken for granted.

The immediate contact point with the present proposal is sub-clause (a) of the Enrichment rule, which seems to say pretty much, or almost, that coherence should be maximized as far as is plausibly possible. This holds in particular as Levinson (2000, p. 116) goes on to characterize the principles as follows: "[...]: they include the search for maximal cohesion, of temporal, spatial, causal, and referential connectedness, and the presumption of stereotypical relations and the actuality of referents."²⁶

Levinson (2000, pp. 117-18) provides a list of kinds of I-implicatures. This is a somewhat variegated collection, where only some of the members are related to coherence. Those that do are the well-known *conjunction buttressing*, *bridging*, and *preferred locational coreference*. Some seem to me more semantic in nature, where a variable requires a contextual value (the interpretation of *noun-noun compounds* and of *possessives*, and the *Mirror maxim*—a preference for collective over distributive readings of some noun phrases). Some have to do with stereotypes and normality:

- (54) *Conditional perfection*: 'If you mow the lawn, I'll give you five dollars' leads to *If and only if you mow the lawn, I'll give you five dollars*.
- (55) *Inference to stereotype*: 'John said "Hello" to the secretary and then he smiled' leads to *John said "Hello" to the female secretary and then John smiled*.
- (56) *Specialization of spatial terms*: 'The nail is in the wood' leads to *The nail is buried in the wood*.

This list, together with Levinson's characterizations of I-implicatures, lead to three questions.

²⁶Here I take Levinson to use 'cohesion' as equivalent with 'coherence'. In fact, Sanders, Spooren, and Noordman 1992, p. 2, and others in the discourse literature, distinguish between the two, taking the *cohesion* of a text to depend on the linguistic *marking* of textual relations.

1. Are the enrichments predicted by coherence considerations driven by expectations of normality?
2. Do expectations of normality invariably lead to new contents, whether by modulation or by implicature?
3. Is it a drive to *minimizing* on the part of the speaker that explains why sentences that are *subject to enrichment* are used?

The answer to the first question is negative, I think, as also Carston (1995, pp. 222-23) has argued. It is hardly *normal* that one opens a door with a key one is given, as contrasted with using a key one is already holding in one's hand at the time. Similarly, it is hardly *normal* that the temperature of something (e.g. body temperature) rises to a *more* dangerous level while *abnormal* that it rises to a *less* dangerous level (the former is true in cases of high fever, the latter e.g. in cases of dangerously low body temperature).

Enrichment Theory does require plausibility of the enriched propositions and the connections appealed to, but this requirement is a *restriction* on enrichment, not what *generates* enrichment, as Levinson's appeal to normality is assumed to do. The drive to coherence and the drive to normality are two quite different phenomena.

The answer to the second question is again negative, I think. For one thing, the expectations of normality on the part of the hearer is a common source of misunderstanding, sometimes depending on prejudices or rash conclusions. In such cases, it is clearly not intended by the speaker by the drive to minimize the expression used.

For another thing, one would indeed expect the speaker to avoid being overly verbose, but one would also expect the choice of words to be guided by what is *relevant* in the context, even if the situation described is abnormal in some currently irrelevant respect. If the relevant contrast is whether some keys are lying on the table or in the car, then

(57) The keys are on the table.

is perfectly adequate, even if the keys are lying on the top of a high stack of books, which is an abnormal way for keys to be on a table. Similarly, if it is relevant that the person referred to is a police officer, but not that the officer is female, the use of 'the police officer' seems more adequate than 'the female police officer', even if female police officers are unusual and the hearer comes to expect the referent to be male.²⁷

²⁷The following is taken (and translated) from the Swedish daily newspaper *Dagens Nyheter* (DN), 14 July 2013, concerning an event that made the headlines during the work on this paper:

The video that DN published on the web last Friday, showing how a police officer hits a man with a baton, has caused strong reactions. Several members of parliament interviewed by DN think that the police officer uses more violence than required.

In some cases certain aspects of normality are indeed communicated, but it is hardly, it seems to me, a general recipe for the generation of content.

The third question is more difficult, because it can concern *personal* speech strategies as well as *sub-personal* speech strategies. On a personal level, the answer is pretty clearly negative. In many and probably most cases, the speaker is simply not *aware* that the strictly semantic content (plus saturation) of a sentence used allows an interpretation clearly deviant from what is intended. The speaker has a content in mind and is often not aware of ambiguities in the generated surface structure or of other non-salient ways for the sentence to be true (the case of the risen temperature in (41) is a likely case in point).

Similarly, the hearer often jumps to conclusions by automatic enrichments, and easily loses awareness of alternative interpretations. It may well be the case that much of the rate of communicative success is explained by a general harmony between our shortcomings as speakers and our shortcomings as interpreters. Whether on the *sub-personal* level, at the level of cognitive design, this should be described as a minimizing *strategy* is a question for some other occasion.

Levinson's talk of "the search for maximal cohesion, or temporal, spatial, causal, and referential connectedness" is clearly much in line with the present account, and Enrichment Theory can be viewed as a proposal for making that ingredient in Levinson's theory more precise. But I don't think it fits his general characterization of I-implicatures.

10. Systematic pragmatics

The main aim of Enrichment Theory is to *predict* pragmatic enrichments. There is no further claim that it *explains* enrichment phenomena, over and above the familiar idea that you can explain individual cases to the extent that your theories predicts them. I do not, however, claim to be able to explain the principles of ET themselves.

To what extent does the theory live up to its predictive ambitions? I'll list a number of limitations to its predictive power.

The police officer in question was female (obvious from the video), but that is not mentioned anywhere in the article. In some other articles it is revealed only by the use of the feminine possessive pronoun, needed for reporting that the officer's police dog helped in attacking the man. In yet others it is mentioned explicitly, but only in passing and further down in the article.

Is the paper deliberately playing down the gender aspect, or is it simply treated as irrelevant? A possible explanation for the former would be that the paper wanted to avoid suggesting either that being a woman is a contributing factor in *causing* the violence (hits a man with a baton [*as a female reaction*]), or that the transgression is more/less *serious* if done by a woman (more violence then required [*for a female police officer*]). If so, the paper used coherence raising considerations.

Limitations to the predictive power of ET

- i) ET relies on the theory of coherence relations for the Resemblance category, but the field of coherence theory has not settled on a definitive list of coherence relations. This leaves some uncertainty as to what predictions result. The problem is provisionally solved by following the Hobbs-Kehler theory of Resemblance.²⁸
- ii) The concepts of *plausibility* and of *simplicity* are underspecified. When is an assumption plausible enough? When is an enrichment simple enough? To some extent even comparisons with respect to plausibility and simplicity may suffer from underdetermination.
- iii) Plausibility relies on a General Background Theory of the world (assumptions about “world knowledge”). To the extent that it is indeterminate what to include in such a general theory, it also remains indeterminate what the enrichment predictions are.
- iv) In some cases the theory requires *weighing* the simplicity of one enrichment against the plausibility of another. Needless to say, we have no idea of how to perform such weighing except by appeal to intuition. At present, this might be the greatest drawback of the account.

These limitations are serious, and they cannot be removed altogether. Nonetheless, there are two mitigating considerations, from opposite sides, that might save the account from being altogether worthless.

On one side, the *computational* side, I think it is possible to make a completely precise theory within an artificially limited domain. We would represent the General Background Theory by some database. We would provide measures of plausibility of the propositions and the laws of the database, provide a measure of complexity of enrichments, and a joint scale for weighing plausibility against simplicity. This would be a kind of AI system to compute enrichments. As far as I can see, Enrichment Theory would lend itself to such an implementation.

On the other side, the *empirical* side, we have a distinction between *central* and *peripheral* phenomena. The central phenomena are those agreed on by virtually everyone, those that are very robust and widespread. We might compare here with the notion of *observation*. As Quine (1960) noted, some observation based judgments depend on rather specific background knowledge. One of Quine's examples ('there was copper in it', Quine 1960, p. 11) is that of a chemist who concludes from the colour of a flame that the ma-

²⁸We can in fact think of predictions that rely on this theory as an empirical test of this part of Coherence Theory.

terial that is burned contained copper. *Observation sentences*, by contrast, such as ‘that is blue’, are sentences for which only minimal background knowledge is needed for the truth value judgment in a particular context, so that every competent member of the speech community would agree. By analogy, we would want to be able to predict those pragmatic effects that virtually all speakers agree on. Predictive success to this extent suffices for descriptive adequacy. Hopefully, these central cases correspond to what is close to universally included in the General Background Theory and to clear cases of assessing plausibility and simplicity.

Because of limitations of predictive capacity, a theory will undergenerate predictions. When a high predictive capacity has been attained, on the other hand, there is a risk of *overgeneration*, i.e. of predicting enrichments that will not in fact take place. That the theory is precise enough to run this risk means that it is falsifiable, which is a good thing (being false is not). The only way to deal with the possibility of false predictions is to test the theory on many and varied examples.²⁹

However, even if a theory in this area is descriptively adequate, it does not follow that it also adequately models the cognitive processes involved in communication. It would of course be strange if a theory reached descriptive adequacy without getting *anything* right about processing. Nevertheless, even if ET does provide an accurate model of processing, as far as it goes, it is clearly *incomplete*. For it does not provide an *algorithm*, i.e. an efficient method for computing the predicted enrichment. The enrichment that satisfies (PPE), if any, does so by *constraint satisfaction*: it is the value that satisfies, or best satisfies, the conditions of coherence raising and availability. Adding the content *with the key that he had handed her* to the semantic content of (1) is what best satisfies these conditions. But how do we arrive at that conclusion?

If we do have a complete but finite database for the General Background Theory, together with plausibility and simplicity measures, as imagined above, then we can also

²⁹There are bound to be complications. Here is one. Consider

- (i) Adam saw a bird and Bill wanted to buy it.
- (ii) Adam saw the bird and Bill wanted to buy it.

The examples are identical except that (i) has an indefinite and (ii) a definite NP. ET predicts that a causal connection will be sought between Adam’s seeing and Bill’s desire. Our background theory tells us that *making Bill aware of the bird* fits the requirement. In this way Adam causes Bill to know about the bird, and thereby *enables* him to desire to buy it. The predicted enrichment would be something like: Bill wanted to buy it [*after Adam had made him aware of it*]. This fits the intuition for (i), but not completely for (ii).

Why? In (ii), the causal interpretation is partly blocked by a Parallel reading: there were two bird-related events; Adam’s seeing it and Bill’s wanting to buy it. ET predicts ambiguity between the Resemblance interpretation and the Possibility interpretation. Then why not ambiguity also in (i)? For some reason the Parallel interpretation appears blocked by the fact that the bird is introduced in the left conjunct by an indefinite NP. This creates an asymmetry that we don’t have in (ii), but why it has this effect on interpretation, I don’t know.

have an algorithm: just test all candidate enrichments up to the maximal degree of complexity and compare them pairwise for coherence, simplicity, and plausibility. Clearly, this is not the way it works in the human mind (for one thing, it would take far too much time). Somehow we manage to find the good candidates in some quick, intuitive and intelligent way. ET does not say how, for it does not model the formation of candidates. It is still possible that ET actually is implemented in the cognitive process, perhaps as a filter on candidates generated by *association* mechanisms of some kind.^{30 31}

Even if we don't have an adequate model of processing, descriptive adequacy is essential to the overall theory of communication. If semantics studies exactly what linguistic meaning contributes to explaining communicative success, then we need to be able to separate the semantic and the pragmatic contributions to communicated content. A descriptively adequate pragmatic theory then helps us getting a handle on linguistic meaning, even if it has shortcomings as a model of pragmatic processes.

³⁰For instance, some theories of *attention* postulate a cognitive filter: stimuli that pass through the filter are raised to attention (cf. Mole 2009).

Some examples in particular suggest that plausibility operates as a filter even at a conscious level. The following is (summer 2013) written as public notices on subway doors in Athens:

- (i) Do not lean against the door. It opens automatically.

(thanks to Kathrin Glüer for the observation). The intuitive enrichment gives the result that the door opens automatically *when one leans against it*, which is also predicted, since it makes the leaning the cause of the opening. Thereby we have Necessity level coherence. The reading is automatically generated, but consciously discarded: the reader thinks such a mechanism would be criminally insane.

³¹On the question of association vs. inference/computation, Recanati (2004, pp. 28-29) favors association, for completely different reasons. According to Recanati, association is *local*, and therefore an enrichment made at one point on *part* of a sentence, does not require computing the semantic value of the entire sentence, the minimal proposition (for a critical discussion of this point, see Pagin 2007, pp. 24-25). Inference, by contrast, would require computing the minimal proposition, according to Recanati.

Sperber and Wilson, by contrast, do regard the comprehension process as inferential, more precisely it is "a dedicated inferential mechanism, a "fast and frugal heuristic," which automatically computes a hypothesis about the speaker's meaning on the basis of the linguistic and other evidence provided" (Wilson and Sperber 2004, p. 625).

Indeed they provide the following argument (Wilson and Sperber 2012b): "We will argue that, within the specifically communicative domain, it is indeed rational for hearers to follow a path of least effort in constructing a hypothesis about the speaker's meaning, and that the pragmatic interpretation process is therefore genuinely inferential [...]."

I don't find this very convincing. If for a process to be *rational* only amounts to being *successful*, there is no reason why it should involve *inference*. And if rationality requires doing something *for a reason*, then this seems to incur an extra cost. If I do ϕ *for the reason* that doing ϕ would be to follow the path of least effort, then it seems I am wasting energy on an inference that gives the same result as just following the path of least effort *without* a reason. Hence, because of the extra reasoning effort, I am not following the path of least effort. What rationality amounts to here is unclear.

References

- Asher, Nicholas and Alex Lascarides (2003). *Logics of Conversation*. Studies in Natural Language Processing. Cambridge: Cambridge University Press.
- Bach, Kent (1994). 'Conversational implicature'. In: *Mind & Language* 9, pp. 124–62.
- Bach, Kent (2010). 'Implicature vs. Explicature: What's the Difference?' In: *Explicit Communication. Robyn Carston's Pragmatics*. Ed. by Belén Soria and Esther Romero. London: Palgrave MacMillan, pp. 126–37.
- Carston, Robyn (1995). 'Quantity Maxims and Generalized Implicature'. In: *Lingua* 96, pp. 213–44.
- (2002). *Thoughts and Utterances*. Oxford: Blackwell.
- (2004a). 'Relevance Theory and the Saying/Implicating Distinction'. In: *The Handbook of Pragmatics*. Ed. by Laurence R. Horn and Gregory Ward. Oxford: Blackwell Publishing, pp. 633–56.
- (2004b). 'Truth-Conditional Content and Conversational Implicature'. In: *The Semantics/Pragmatics Distinction*. Ed. by Claudia Bianchi. Stanford, Cal.: CSLI Publications, pp. 65–100.
- Davidson, Donald (1984). *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press.
- Davis, Wayne (2010). *Implicature*. URL: <http://plato.stanford.edu/archives/win2010/entries/implicature/>.
- Dowty, David (1986). 'The Effects of Aspectual Class on the Temporal Structure of Discourse: Semantics or Pragmatics?' In: *Linguistics and Philosophy* 9, pp. 37–61.
- Grice, Herbert Paul (1975). 'Logic and Conversation'. In: *Speech Acts*. Ed. by Peter Cole and Jerry L. Morgan. Vol. 3. Syntax and Semantics. Academic Press. Reprinted in Grice 1989, chapter 2, pp 22–40. Page references to the reprint.
- (1989). *Studies in the Ways of Words*. Cambridge, Mass.: Harvard University Press.
- Groenendijk, Jeroen and Martin Stokhof (1994). 'Questions'. In: *Handbook of Logic and Language*. Ed. by Alice ter Meulen and Johan van Benthem. Elsevier. Chap. 19, pp. 1055–1124.
- Hobbs, Jerry R (1985). *On the Coherence and Structure of Discourse*. Tech. rep. CSLI-85-37. Stanford University: CSLI.
- Hobbs, Jerry R et al. (1993). 'Interpretation as Abduction'. In: *Artificial Intelligence* 63, pp. 69–142.
- Horn, Laurence R. (2004). 'Implicature'. In: *The Handbook of Pragmatics*. Ed. by Laurence R. Horn and Gregory Ward. Oxford: Blackwell Publishing, pp. 3–28.
- Horn, Laurence R. and Gregory Ward, eds. (2004). *The Handbook of Pragmatics*. Oxford: Blackwell Publishing.

- Hume, David (1748). *An Inquiry Concerning Human Understanding*. New York: Liberal Arts Press. 1955 edition.
- Kamp, Hans and Uwe Reyle (1993). *From Discourse to Logic: Introduction to Model Theoretic Semantics of Natural Language, Formal Logic and Discourse representation theory*. Dordrecht: Kluwer.
- Kaplan, David (1989). 'Demonstratives'. In: *Themes from Kaplan*. Ed. by J Almog, J Perry, and H Wettstein. Oxford: Oxford University Press, pp. 481–563.
- Kehler, Andrew (2002). *Coherence, Reference, and the Theory of Grammar*. Stanford University: CSLI Publications.
- Knott, Alistair and Robert Dale (1994). 'Using Linguistic Phenomena to Motivate a Set of Coherence Relations'. In: *Discourse Processes* 18, pp. 35–62.
- Leth, Palle (2010). 'Paraphrase and Rhetorical Adjustment. An Essay on Contextualism and Cohesion'. PhD thesis. Gothenburg: University of Gothenburg and Université Paris-Est Créteil.
- Levinson, Stephen C (2000). *Presumptive Meanings. The Theory of Generalized Conversational Implicature*. Cambridge, Mass.: MIT Press.
- Lewis, David (1979). 'Scorekeeping in a Language Game'. In: *Journal of Philosophical Logic* 8. Reprinted in Lewis 1983, 233–49. Page references to the reprint, pp. 339–59.
- (1983). *Philosophical Papers. Volume I*. Oxford: Oxford University Press.
- Maienborn, Claudia, Klaus von Heusinger, and Paul Portner, eds. (2011). *Semantics. An International Handbook of Natural Language Meaning*. Berlin: Mouton de Gruyter.
- Mole, Christopher (2009). *Attention*. URL: <http://plato.stanford.edu/archives/fall2009/entries/attention/>.
- Nunberg, Geoffrey (1979). 'The Non-Uniqueness of Semantic Solutions: Polysemy'. In: *Linguistics and Philosophy* 3, pp. 143–184.
- (1993). 'Indexicality and Deixis'. In: *Linguistics and Philosophy* 16, pp. 1–43.
- (1995). 'Transfers of Meaning'. In: *Journal of Semantics* 12, pp. 109–32.
- Ogihara, Toshiyuki (2011). 'Tense'. In: *Semantics. An International Handbook of Natural Language Meaning*. Ed. by Claudia Maienborn, Klaus von Heusinger, and Paul Portner. Vol. 2. Berlin: Mouton de Gruyter. Chap. 57, pp. 1483–64.
- Pagin, Peter (2005). 'Compositionality and context'. In: *Contextualism in Philosophy*. Ed. by Gerhard Preyer and Georg Peter. Oxford: Oxford University Press, pp. 303–48.
- (2007). 'Pragmatic composition?' In: *Svet jazyka a svet za jazykom*. Ed. by Marian Zouhar and Tomasz Marvan. Vol. Supplement volume of *Organon F*. Organon F, pp. 11–26.

- Pagin, Peter and Francis Jeffry Francis Pelletier (2007). 'Content, context and communication'. In: *Context-sensitivity and Semantic Minimalism. New Essays on Semantics and Pragmatics*. Ed. by Gerhard Preyer and Georg Peter. Oxford University Press, pp. 25–62.
- Portner, Paul (2011). 'Perfect and Progressive'. In: *Semantics. An International Handbook of Natural Language Meaning*. Ed. by Claudia Maienborn, Klaus von Stechow, and Paul Portner. Vol. 2. Berlin: Mouton de Gruyter. Chap. 49, pp. 1217–61.
- Quine, Willard Van Orman (1960). *Word and Object*. Cambridge, Mass.: MIT Press.
- Recanati, François (2004). *Literal Meaning*. Cambridge: Cambridge University Press.
- (2010). *Truth-Conditional Pragmatics*. Oxford: Oxford University Press.
- Reichenbach, Hans (1947). *Elements of Symbolic Logic*. New York: Macmillan Publishing Co.
- Rumelhart, D (1979). 'Some Problems with the Notion of Literal Meanings'. In: *Metaphor and Thought*. Ed. by A Ortony. 2nd ed. Cambridge: Cambridge University Press, pp. 71–82.
- Sanders, Ted J. M., Wilbert P. M. Spooren, and Leo G. M. Noordman (1992). 'Toward a Taxonomy of Coherence Relations'. In: *Discourse Processes* 15, pp. 1–35.
- Sperber, Dan and Deirdre Wilson (1986). *Relevance. Communication & Cognition*. 1st ed. Oxford: Blackwell.
- (1995). *Relevance. Communication & Cognition*. 2nd ed. Oxford: Blackwell.
- Stanley, Jason (2000). 'Context and Logical Form'. In: *Linguistics & Philosophy* 23, pp. 391–434.
- Wilson, Deirdre and Tomoko Matsui (1998). 'Recent Approaches to Bridging: Truth, Coherence, Relevance'. In: *UCL Working Papers in Linguistics* 10, pp. 1–28. Reprinted in Wilson and Sperber 2012a, pp. 187–209. Page references to the reprint.
- Wilson, Deirdre and Dan Sperber (2004). 'Relevance Theory'. In: *The Handbook of Pragmatics*. Ed. by Laurence R. Horn and Gregory Ward. Oxford: Blackwell Publishing, pp. 607–32.
- (2012a). *Meaning and Relevance*. Cambridge: California University Press.
- (2012b). 'Pragmatics, Modularity, and Mindreading'. In: *Meaning and Relevance*. Cambridge: California University Press. Chap. 12, pp. 261–78.
- Zobel, Sarah (2010). 'Non-Standard Uses of German 1st Person Singular Pronouns'. In: *JSAI-isAI, LNAI 6284*. Ed. by K Nakakoji, Y Murakami, and McCreedy E. Berlin-Heidelberg: Springer, pp. 292–311.