

Coherence and coherence establishment: Lessons from eliciture *

Jonathan Cohen[†] and Andrew Kehler[‡]

Abstract

The observation that the same coherence relations serving to constrain interpretation at the discourse/intersentential level also operate on constituents within sentences (we label enrichments resulting from such intrasentential coherence relations ‘eliciture’) has surprising and far-reaching consequences for our understanding of coherence and coherence establishment.

Current accounts construe discourse coherence establishment — and, therefore, discourse level coherence-based enrichment — as resulting from a bottom-up search for ways in which elements of contents expressed might be coherently related to one another, and mandatorily triggered by a requirement to bind components into wholes. But such accounts leave us without an obvious explanation for intrasentential coherence. One problem is that intrasentential coherence establishment (unlike intersentential coherence establishment) is not required for felicity: hence, one cannot see coherence-based enrichments uniformly as the mandatory downstream consequence of a trigger. A second problem is that, because elicitures can arise from complex interactions between any combination of a sentence’s constituents, an account rooted in a search for possible coherence relations between expressed contents quickly runs into trouble: such a search would have to compare the contents expressed by every pair of constituents, then every triple, and so on. This is clearly not computationally tractable.

These and related considerations suggest a quite different picture of the inferences arising from coherence establishment — one on which such inferences are not results of triggered searches, but the

*This work is fully collaborative; the authors are listed in alphabetical order.

[†]Department of Philosophy, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0119, cohen@ucsd.edu

[‡]Department of Linguistics, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0108, akehler@ucsd.edu

inevitable upshots from entertaining combinations of linguistically expressed contents, analogous to our recognition of causal and other relations obtaining between components of the non-linguistically presented world. We'll develop this picture by starting with eliciture, then show how it can be extended to intersentential coherence establishment, and finally draw out consequences resulting from this reconceptualization. Among other benefits, we'll contend that our account provides explanations for features (such as the preference for causal interpretations) that have required special principles in the more traditional accounts developed with only intersentential coherence in mind.

1 Introduction: Two flavors of coherence-based enrichment

A substantial line of work during the last several decades has investigated the nature of coherence and the inferential processes by which it is established by language users (inter alia, [Halliday and Hasan 1976](#); [Hobbs 1979](#); [Longacre 1983](#); [Mann and Thompson 1987](#); [Hobbs 1990](#); [Martin 1992](#); [Knott and Dale 1994](#); [Sanders 1997](#); [Kehler 2002](#); [Asher and Lascarides 2003](#); [Pagin 2014](#)). The large majority of this work has focused on the level of discourse, whereby coherence is established via the recognition of coherence relations that hold among sentences. Importantly, however, it appears that closely analogous coherence relations obtain between elements within sentences, and that they give rise to closely analogous interpretive effects (cf. [Kronfeld 1990](#); [Hobbs 2010](#); [Rohde et al. 2011](#); [Kehler and Rohde 2019](#); [Cohen and Kehler 2021](#)). [Cohen and Kehler \(2021\)](#) offer the label ELICITURE for enrichments resulting from such intrasentential coherence relationships.

As a first demonstration of the intersentential/intrasentential parallel, consider the pair (1a–b):

- (1) a. The company fired the manager. He embezzled money.
- b. The company fired the manager who embezzled money.

Here both the intersentential example (1a) and the intrasentential example (1b) defeasibly invite the inference that the embezzling was the reason for the firing. In both cases, these enrichments plausibly arise as an effect of the establishment of a coherence relation between two contents (*the employee embezzled money, the boss fired the employee*), i.e., by treating the latter as an explanation for the former. It is natural to think of both these enrichments

as falling under the heading [Hobbs \(1990\)](#) calls EXPLANATION. Or consider the pair (2a–b):

- (2) a. The employee went to the store. He bought a bottle of scotch for the office party.
- b. The employee who went to the store bought a bottle of scotch for the office party.

Both (2a) and (2b) defeasibly invite a coherence-driven expansion to the effect that the scotch was bought at the store. It appears that, in both cases, the enrichment is driven by an OCCASION relation that treats the event of the scotch buying and the event of the employee's going to the store as segments of a single spatiotemporally connected series of events. Likewise, utterances of both the intersentential example (3a) and the intrasentential example (3b) invite the defeasible coherence-driven conclusion that the employee's donating vacation time to Sue is an instance of the generalization that she does favors for all her co-workers (this is an instance of what [Hobbs \(1990\)](#) calls EXEMPLIFICATION).

- (3) a. The employee does favors for all her co-workers. She donated some of her vacation time to Sue.
- b. The employee who does favors for all her co-workers donated some of her vacation time to Sue.

The parallel behavior observed in inter- and intra-sentential examples is striking, and, we think, theoretically significant. In particular, we see in these data an impressive alignment in both the coherence relations themselves and the interpretive enrichments they give rise to in both sorts of cases. To be fair, there are interesting differences between the intersentential and intrasentential cases that deserve comment, and to which we'll return below. However, given the striking parallelism, the default view is clearly that there is one system at work here, i.e., that an adequate account must extend to both inter- and intra-sentential cases.

2 Theories of coherence establishment

Frameworks for thinking about coherence relations in the literature (e.g. [Hobbs 1979, 1990](#); [Sanders 1997](#); [Asher and Lascarides 2003](#)) were developed around intersentential cases. Many of the early analyses of coherence focused on providing inventories of relations (inter alia, [Halliday and Hasan 1976](#); [Longacre 1983](#); [Mann and Thompson 1987](#); [Martin 1992](#))

as a means of characterizing coherent discourse, but without much in the way of computational models that could put such relations to use. Since then, accounts containing fuller computational models have emerged, most prominently including those of [Hobbs et al. \(1993\)](#) and [Asher and Lascarides \(2003\)](#). Though these two accounts differ in a variety of respects, they share an important characteristic that will be central in what follows: they treat the establishment of coherence as resulting from a bottom-up search for ways in which elements of contents expressed might be coherently related to one another, triggered by a mandate to bind components into wholes.¹ It will be useful for the ensuing discussion to consider the desiderata for a computational account of coherence establishment, which we take to include the following:

- It should provide a means for determining the ‘best explanation’ of the coherence of the passage among alternative coherence proofs.
- It should provide a means for distinguishing coherent from incoherent discourses.
- It should capture the logical reasoning process that hearers are taken through when establishing the coherence of a passage.
- It should explain the source of pragmatic enrichments, that is, additional inferences to unexpressed content that hearers make in service of establishing coherence.
- It should draw inferences that are defeasible, and hence capable of being retracted as further information becomes available.

In order to illustrate how a computational theory of coherence establishment might meet these explanatory goals, consider the following minimal theory of coherence establishment modeled on ideas of [Asher and Lascarides \(2003\)](#) and (especially) [Hobbs et al. \(1993\)](#), which, for ease of exposition, we henceforth refer to as BRACE — the Basic Relational Account of Coherence Establishment. The starting assumption that underlies relational theories of coherence is that hearers are mandated to identify *some* coherence relation between clauses. It is this mandate that gets the inference process going. [Hobbs et al. \(1993\)](#), for instance, adopt two axioms to capture

¹On our usage, the distinguishing feature of what we’re calling “bottom-up processes” is that they unfold as a response to a detected trigger of some kind (and not, say, that they avoid reliance on long term memory). Thanks to David Danks for pressing us on this issue.

this process. The first, shown in statement (4), states that a sentence is a discourse segment.²

$$(4) (\forall i, j, e) s(i, j, e) \supset \text{Segment}(i, j, e)$$

The second axiom, shown in statement (5), allows for two adjacent segments to be composed into a larger one if a coherence relation can be established between the two.

$$(5) (\forall i, j, k, e_1, e_2, e) \text{Segment}(i, j, e_1) \wedge \text{Segment}(j, k, e_2) \\ \wedge \text{CoherenceRel}(e_1, e_2, e) \supset \text{Segment}(i, k, e)$$

Then, to interpret a coherent text comprised of words 0 through N, it must be proven that it is a segment, as expressed by statement (6).

$$(6) (\exists e) \text{Segment}(0, N, e)$$

The task of discourse parsing on this model is therefore analogous to that of sentence parsing: coherent discourses require their clauses to be structurally related in much the same way that grammatical sentences require their syntactic constituents to be, and hence discourses admit of hierarchical structures in the way that sentences do. The fundamental difference is that, whereas the constraints on when adjacent constituents can combine to form larger ones in building a sentence structure are dictated by grammatical rules, in the realm of discourse ‘grammar’ it is the ability to infer coherence relations that determines when two clauses (or collections thereof at higher levels of structure) can be combined.

Evidence for a mandate to establish coherence among co-occurring clauses can be seen by consideration of the following two examples.

(7) The company fired the manager. He embezzled money. (=1a)

(8) The company fired the manager. He has red hair.

When interpreting either passage, a hearer will not be satisfied to merely update their beliefs with the two propositions that the speaker conveyed. Nor are they content to view the discourses as coherent merely by virtue of centering around the same topic — in this case, the manager. They want more. According to relational theories of coherence, they are required to

²Here, i and j are the start and end indices on words in a word sequence, and e is the sentence’s assertion or topic.

identify which instance out of an inventory of coherence relations serves as the glue that establishes their mutual relevance.

In most contexts, the two clauses in (7) are most naturally construed as participating in an EXPLANATION relation, whereby the second clause expresses a cause or reason for the eventuality described in the first. This process engages with causal generalizations that the speaker presupposes (i.e., believes to be in the common ground). Using the formalism from [Hobbs et al. \(1993\)](#), a relevant rule might look something like this:

$$(9) (\forall e_1, e_2, e_3, e_4, x, y, z) \text{embezzle}(e_1, x, y)^{c1} \wedge \text{money}(y)^{c2} \\ \wedge \text{owned_by}(e_2, y, z)^{c3} \wedge \text{employer}(z, x)^{c4} \wedge \text{discover}(e_3, z, e_1)^{c5} \\ \wedge \text{fire}(e_4, z, x)^{c6} \wedge \text{etc}(e_1, e_2, e_3, e_4, x, y, z)^{c7} \\ \supset \text{cause}(e_1, e_4)$$

Setting the details aside temporarily, this axiom captures a causal generalization according to which if an employee is caught embezzling money by his company and he is fired, then it is likely that the embezzling caused the firing. Equipped with this knowledge, our hearer will use it to establish the Explanation relation, by virtue having inferred a causal relation per the consequent of (9) and the following rule:

$$(10) (\forall e_1, e_2) \text{cause}(e_2, e_1) \supset \text{Explanation}(e_1, e_2)$$

And since Explanation is a coherence relation:

$$(11) (\forall e_1, e_2) \text{Explanation}(e_1, e_2) \supset \text{CoherenceRel}(e_1, e_2, e_1)$$

the mandate imposed by (5) has been satisfied.

On the analysis we're considering, pursuing the causal relationship necessary to establish Explanation will utilize backchaining on axiom (9): the meaning of the first clause of (7) instantiates the sixth conjunct in the antecedent of the axiom, and the second clause instantiates the first and second conjuncts. The other terms in the antecedent will need to be accommodated, and hence this inference process will bring with it corresponding pragmatic enrichments: that the money mentioned in the second clause belonged to the company mentioned in the first clause (the third predicate in the antecedent of (9)), that the manager in fact works for the company (the fourth predicate) and that the company did in fact discover the manager's crime (the fifth predicate). Whereas (7) neither asserts nor entails any of these, they are assumptions that have to be made

for the hearer to use the causal generalization encoded by (9) as the basis for establishing the coherence of the passage, and the hearer will presumably accommodate them in the absence of evidence to the contrary. These accommodations necessarily come at a cost, however. For example, in the formalism of Hobbs et al. (1993), the superscript on each predicate in (9) represents the price of having to assume its truth rather than being able to deduce it. The winning proof of coherence is the one associated with the lowest total cost.³

The establishment of coherence in passage (8), on the other hand, comes less readily. Here a typical hearer will ponder the question of why having red hair would cause someone to be fired. Presumably, outside of a special context, no causal generalization analogous to (9) exists that can be similarly applied to this case. Does the company have a hair color policy? Did the CEO's spouse run off with a redhead? Any passage can be rendered coherent if enough assumptions are accommodated, but the lack of a salient, low-cost explanation renders the passage incoherent. The important point here is that, as noted by Hobbs (1979), the very fact that hearers are led to entertain such explanations is itself proof that a mandate to establish coherence is in effect. Otherwise, our hearer would be satisfied to have simply learned two new things about the manager and happily move on.

With this sketch of BRACE in hand, we can now step back and ask how this framework satisfies the desiderata listed at the top of this section. First, it provides, in principle, a way of determining the preferred explanation of the coherence of a passage among alternative coherence proofs. All proofs of coherence come at a cost, which will be higher or lower dependent on the number of assumptions that need to be made and their relative costs. By performing a search over alternative coherence analyses and calculating the cost of each, the one with the lowest cost can be selected as the winner. Note, of course, that this feature of the system is only conceptual: There exists no established way of encoding facts and axioms in a knowledge base (countless degrees of freedom exist), no coherent and consistent basis for assigning costs to terms in such axioms, nor any suitable mechanism for inference capable of arriving at proofs of coherence efficiently. These

³The *etc* predicate in (9) represents a term that cannot be proven and hence must be accommodated, again at a cost. Since (9) represents a mere causal generalization to which there may be exceptions (it's not strictly guaranteed that if an employee is caught embezzling and he gets fired, that the embezzling is the cause of the firing), this predicate can be seen as representing the cost of assuming that all other contingent facts are in order such that it is operating as a causal law for the particular example to which it is being applied.

difficulties notwithstanding, BRACE nonetheless gives us a framework for thinking about how coherence establishment might be carried out.

This cost-based search-and-select procedure also provides a handle on the second desideratum, the ability to distinguish coherent from incoherent discourses. We've seen how this is accomplished in our comparison of examples (7) and (8). Since all proofs of coherence will come at a cost, we can hypothesize the existence of a threshold at which the cost of the assumptions necessary to establish coherence under a particular analysis is too high. In light of a sea of competing, excessively costly proofs (the company has a hair color policy; the CEO's spouse ran off with a redhead), incoherence results. Note that if a suitable number of otherwise high-cost assumptions are provided by the context — e.g., it is in fact already known that the CEO's spouse ran off with a redhead — then the overall cost may drop below threshold, and the passage achieves coherence.

By way of our example, we've already briefly explained the manner in which the third and fourth desiderata are met — capturing the logical reasoning process that hearers are taken through, and explaining the source of pragmatic enrichments. According to BRACE, the processing of discourse triggers a search for coherence, which entails traversing paths of inference capable of connecting the meanings of two utterances in a way that satisfies the constraints imposed by a coherence relation. As we saw in our discussion of (7), often these constraints are satisfied neither by the hearer's prior knowledge nor by entailments from the utterances themselves and hence need to be accommodated, albeit at a cost. When the winning coherence construal is adopted, the accommodated components of meaning serve as pragmatic enrichments.

And finally, the nature of the inference processes that serve coherence establishment, like all forms of pragmatic enrichment, generate enrichments that are defeasible. Consider again the discourse shown in (12a–b), but in this case with the follow-on shown in (12c).

- (12) a. The company fired the manager.
- b. He embezzled money.
- c. But the reason he was fired is that he repeatedly comes in late.

As before, upon reading (12a–b) a hearer is likely to infer an Explanation relation, whereby the embezzlement was a reason for the firing. The follow-on shown in (12c), however, will then lead the hearer to retract that inference, forcing a reanalysis of the coherence structure of the passage that engages different causal generalizations. Note that this switch will likely

lead to the retraction of other enrichments that came along for the ride when applying the causal generalization in (9). For instance, we no longer know that the company discovered the embezzlement, since this constraint plays no role in a causal generalization that relates coming in late and getting fired.

This analysis of coherence establishment thus allows for such inferences to be contingent and retractable, and yet still demand that *some* coherence relation be established between co-occurring clauses. For instance, if we consider only the discourse in (12a–b), a possible alternative construal is as a Result relation: if we allow for the assumption that the company didn't immediately show the manager the door upon firing him, the manager's anger about being fired may have driven him to embezzle money before leaving. In this case, the proof of coherence doesn't require that the company be aware of the embezzlement. Likewise, if we change the assumptions about the context and stipulate that the sentences in (12a–b) provide partial answers to the question *What are some events that the manager was involved with today?*, the enrichments that the embezzlement was the reason for the firing and that it was discovered by the company become far less inevitable. This is again predicted, as the trigger for those enrichments — axiom (9) — would not come into play during the process of coherence establishment. Thus, as we have pointed out, many proofs of coherence will generally be possible for a particular passage — and what may initially be a dispreferred construal might become the preferred one when more information becomes known.

3 Problems with eliciture

The BRACE framework thus satisfies the five desiderata for a computational theory of coherence establishment, and undoubtedly holds considerable intuitive appeal as an account of intersentential coherence-driven pragmatic expansions, such as those on display in (1a), (2a), and (3a). However, despite its attractions, it is less clear that this style of analysis can be extended to cases in which the very same inferences are drawn among intrasententially introduced contents, as in the instances of eliciture in (1b), (2b), and (3b). Moreover, we contend that the difficulties in extending the analysis to capture eliciture stem from fundamental commitments that BRACE makes, such that no mere tweak of the framework is likely to overcome these hurdles.

At least two significant features of eliciture pose obstacles for extending existing accounts of intersentential coherence-driven expansions to intrasentential cases. The first is that (as noted in §2) while coherence establishment is required for the felicity of intersentential examples like (1a), (2a), and (3a), the expansions in the counterpart intrasentential examples (1b), (2b), and (3b) are optional. But if coherence establishment is optional in intrasentential examples, it follows immediately that one cannot see eliciture — *a fortiori*, that one cannot see all coherence-based enrichment — as an unfolding of a bottom-up interpretive process initiated by the detection of a trigger, per extant accounts of the sort reviewed in §2.

Moreover, there is a second threat to the prospect of extending extant bottom-up search-based accounts to intrasentential cases: in many intrasentential cases, it is unclear just what could serve as a trigger, or where a bottom-up procedure should be looking for one. To see why, note that elicitures are not the results of any single constituent, but of interactions among (interpretations of) constituents. We've already seen examples of elicitures that arise from an interaction between a matrix verb and either a relative clause modifying the subject NP (as in (2b), (3b)) or a relative clause modifying the object NP (as in (1b)). But it's easy to find elicitures that don't involve the matrix at all, such as (13).

(13) The drunk kid who got into a car accident is home now.

The eliciture arising in (13) to the effect that the drinking led to the accident results from a relation between the content expressed by two modifiers of the subject NP: a proposition derived from an adjectival and the nominal that it modifies (*the kid was drunk*) and a proposition derived from an RC and the NP to which it attaches (*the kid got into a car accident*). The proposition denoted by the matrix (*the kid is home now*) doesn't come into play. These considerations suggest that elicitures can arise from interactions between constituents standing in a wide variety of syntactic configurations. However, it won't suffice to consider interactions only between pairs of constituents. To see why, consider the (as usual, optional) intrasentential coherence-driven expansion, invited by (14a), that the pilot was arrested because he was flying (or perhaps preparing to fly) while inebriated.

- (14) a. The drunk pilot was arrested.
b. The 53 year-old pilot was arrested.
c. The drunk person was arrested.
d. The drunk pilot was walking down the street.

Crucially, the variants in (14b–d) invite no analogous inferences, despite the fact that (14b) contains the same head noun and verb phrase as (14a), (14c) contains the same adjective and verb phrase as (14a), and (14d) contains the same head noun and adjective as (14a). The reason for this is intuitively clear: the eliciture depends on a causal-explanatory generalization that requires all three constituent meanings provided by (14a) to be instantiated, that is, that (i) a pilot, (ii) when flying or preparing to do so, (iii) can be arrested for being inebriated.

We take these considerations to show that elicitures can arise from complex interactions between any combination of a sentence’s constituents. This makes it hard to see how any account rooted in a bottom-up, triggered, search for possible coherence relations between expressed contents could succeed. Such a search for a trigger would have to compare the contents expressed by every pair of constituents, then every triple, and so on, in a way that would quickly overrun the limits of computational tractability.

4 Towards a unified account

The observations of §3 suggest that extant accounts of intrasentential coherence in terms of a triggered bottom-up search, such as those surveyed in §2 (or descendants thereof), are unlikely to extend to the treatment of eliciture. One possible reaction to this lesson would be to give up our initial aspiration for a unification, and, instead, pursue separate theoretical accounts of inter- and intra-sentential coherence-driven expansion. We believe this would be a mistake.

Indeed, we contend that a unified account will not only have wider scope, but will also make available an attractive, systematic explanation of the differences between inter- and intra-sentential cases. This explanation is provided by the interaction of a single underlying mechanism with the distinct felicity requirements of discourses on the one hand and sentences on the other. That coherence establishment is mandatory in intersentential cases derives from the observation, noted in §2, that discourses themselves are not simply concatenations of independent sentences, but hierarchically structured wholes built from clauses (or larger units comprised from clauses at higher levels of structure) standing in relations of coherence (Hobbs 1990; Kehler 2002; Asher and Lascarides 2003). The felicity of a discourse requires that such structural relations obtain between its constituents, just as the grammaticality of a sentence requires certain grammatical relations to obtain between its constituents. This means that, barring some other source

of structure, the failure of coherence establishment in an intersentential example like (1a) breaks the unity of the discourse whole, leaving us with an infelicitous mere concatenation of clauses. In contrast, though sentences are also hierarchically structured wholes, their unity does not ordinarily require coherence establishment. Rather, sentences count as unified when their constituents stand in the right syntactic and semantic relations, as specified by grammatical rules. For this reason, though treating the matrix and restrictive relative clauses of intrasentential examples like (1b) as standing in a coherence relation is an available interpretive option, it is not mandated. One can also simply read the restrictive relative clause in this example as merely having its typical function of restricting the range of candidate referents for the head noun *employee*, without attempting to establish a coherence relation between the matrix and relative clause at all. This is to say that the coherence reading of intrasentential examples like (1b) is merely invited. It is also to say that, in intrasentential cases where no such coherence relation presents itself, interpretation can proceed with a bare identificational reading, and that this will leave the felicity of the sentence undamaged. If this much is correct, we should expect to see examples of multiple predication whose intersentential expressions are infelicitous although their intrasentential versions are unimpeachable. And this is exactly what we find, as in examples (15–16):

- (15) a. # The employee broke his leg. He likes plums.
b. The employee who likes plums broke his leg.

- (16) a. # The company fired the manager (even though) he drives a blue Lexus.
b. The company fired the manager who drives a blue Lexus.

To our minds, the *prima facie* plausibility of this explanation of the contrast in optionality between coherence-based enrichments in inter- and intra-sentential cases adds to the motivations for unification already noted in §1. To see why, it may be helpful to consider what we posit to be a parallel theoretical situation with respect to the referential phenomenon of bridging inferences (Clark 1975; Prince 1981; Hawkins 1984), on display in (17), from Kehler (2015):

- (17) a. I went to a wedding on Saturday and befriended *the maid of honor*.
b. I went to a wedding on Saturday and befriended *a bridesmaid*.

There are two key features of interest in these examples. First, both the formulation involving a definite in (17a) and the formulation involving an indefinite in (17b) invite a so-called bridging inference to the effect that the wedding participant mentioned in the second clause is part of the wedding mentioned in the first. Second, despite this commonality, there is an interesting difference between the two cases: the inference with the indefinite is cancelable (18b), whereas the inference with the definite is not (18a).

- (18) a. I went to a wedding on Saturday and befriended *the maid of honor*.
It turned out that she was in a different wedding.
b. I went to a wedding on Saturday and befriended *a bridesmaid*. It turned out that she was in a different wedding.

So, once again, there are similarities but also differences between the interpretive effects arising from the two configurations.

But, faced with these facts, it would surely be inadvisable to posit separate mechanisms for bridging inferences involving indefinites and those involving definites. Among other things, it is hard to see why two such distinct mechanisms would manifest such similar behavior. And, in any case, it is possible to derive the difference in cancelability from the (independently motivated) constraints that definite and indefinite NPs place on their referents. Specifically, we can understand the non-cancelability of a bridging inference in the presence of a definite NP as arising from the unique identifiability constraint associated with definites (Gundel et al. 1993): if the maid of honor was not part of the aforementioned wedding then the *the*-NP would not have been licensed. On the other hand, there is no such constraint imposed by *a*-NPs; consequently, while the bridging inference will typically be drawn in a configuration like (17b), it is not required for the referring expression to be licensed, and hence the inference can be felicitously canceled.

We see the situation with respect to the similarities and differences in optionality between intersentential and intrasentential instances of coherence establishment in analogous terms. A unified account is warranted here, just as a unified account of the two forms of bridging inferences is warranted, because a unified account makes available an elegant and parsimonious explanation of both what is shared and what is not shared between the two forms, while a disunified account does not. With respect to coherence establishment, positing separate mechanisms for intersentential and intrasentential coherence establishment might capture the differences

between such instances, but would leave their extensive similarities wholly unexplained. On the other hand, a unified account allows us to explain the similarities between the two forms in terms of a shared mechanism, while also explaining the observed differences in their behavior in terms of the interaction of that mechanism with the fact that coherence is ordinarily necessary for establishing discourse felicity, but not sentence felicity.

Given these powerful motivations for unification, the observation that BRACE cannot be extended to intra-sentential cases (§3) suggests that that theoretical framework will not, by itself, explain everything we should hope to understand about coherence-driven expansion. Despite this, we see BRACE as continuing to play an important explanatory role within a fuller theory. To be sure, we take the considerations adduced above to show that a satisfactory explanation of coherence-driven expansion will ultimately require an account of the non-obligatory expansions that arise intrasententially (i.e., in the absence of a mandate), and that, despite its other virtues, BRACE cannot be that account. However, whatever the envisaged account of non-obligatory expansions looks like (cf. §5 for our positive suggestions), it is reasonable to think that, once in hand, it could, in turn, feed a BRACE-like explanation of the coherence-driven inferences/expansions that occur in intersentential cases (i.e., where coherence-driven expansion is mandated). On this picture, then, there is reason for retaining BRACE as a component of the unified total theory. (Having said this, we shall argue in §§6–7 that thinking about BRACE in the context of eliciture suggests modifications that improve the view and solve longstanding puzzles it has faced.)

5 Ampliative inference

If we are right, prevailing accounts directed at intersentential cases are incapable of serving as the unified theory for inter- and intra-sentential coherence-driven expansions we are seeking. A new theoretical account is needed.

The positive suggestion we want to advance is that the expansions in question can be understood as special cases of a more general form of inference for which there is plenty of independent evidence, and whose existence we take to be uncontroversial. In particular, there is reason for viewing the coherence-based inferences arising in the inter- and intra-sentential examples in (1)–(3) as instances of the same type of ampliative

inferences to causal and other conclusions that we make when we encounter the world non-linguistically.

To see what we have in mind, consider an occasion in which an employee encounters strong evidence that her manager has embezzled money. Then, a few days later, she witnesses the manager being fired. Our employee will likely infer (provisionally) that the firing was due to the embezzling. Note here that the inference is not the result of a triggered, bottom up, search: neither the cognitive procedure of making the inference, nor the conclusion that results from such processing, is mandated by the evidence itself, or by the realization that that evidence (or anything deductively following therefrom) violates any rule/norm/constraint. Moreover, it is ampliative (its conclusion goes beyond the evidence and its logical entailments), top-down (it is strongly guided by world-knowledge), and highly defeasible. Now consider the same context — work again, and suppose that our agent subsequently sees a customer asking the manager where the automotive department is. In this case, our employee is unlikely to infer any relationship between the customer's question and the employee's embezzlement. Why? Because our world knowledge does not support a causal connection between the events, and so does not underwrite the kind of ampliative conclusion reached in the first case. The lesson appears to be that when an inference of this sort suggests itself as in the firing scenario, a thinker will likely draw it, at least provisionally. However, the world remains perfectly coherent when no such inference between eventualities presents itself, as in the customer question scenario.

Our suggestion is that the very same ampliative strategies for making sense of the non-linguistically presented world should show up in the course of interpreting discourses that describe such situations linguistically. Thus, when a hearer interprets (19),

(19) The company fired the manager who embezzled money.

he will reasonably associate the firing with the embezzlement in the same way in which they are associated in the (non-linguistic) case considered above, in which the events were perceived. Likewise, our addressee will presumably not draw a causal inference for the variant in (20):

(20) A customer asked the manager who embezzled money where the automotive department is.

Here, all that is necessary for the object NP to be felicitous is that it allows the addressee to identify the referent.

Perhaps needless to say, the question of exactly how such ampliative inference works as a general matter remains a (perhaps *the*) deeply unresolved mystery about how the mind works. We do not pretend to have offered a solution to that question. What we are suggesting is, rather, that the inter- and intra-sentential coherence-based inferences we make in interpreting (1)–(3) do not amount to a separate, further, question. If this is right, then there is no need to provide a separate account of the special mechanisms underpinning these inferences as part of our theory of extrasemantic interpretation. Rather, the inferences in question can be viewed as the results of applying to linguistic materials the general mechanisms for ampliative inference that are uncontroversially part of our mental endowments.

Indeed, reflection on how the search procedure might unfold provides independent support for the idea that we should treat the extrasemantic inferences under study as instances of general strategies for ampliative inference rather than results of a special form of bottom up, triggered, searches. For instance, when we interpret (1b), it is surely not that our understanding of the semantic contents of the matrix verb and the relative clause first provides a trigger to the effect that those contents stand in some kind of relevance relation, and which a search is required to spell out (*hmm, firing and embezzling money; those two things are somehow related to one another, now I need to go figure out how*). On what basis could one identify any such relation if one didn't *already* understand how those contents were relevantly related? Surely the correct picture must be that one is already computing whether the target contents stand in the right kind of relation or not. The recognition/absence of recognition of an instance of the so-called trigger is a result of this computation, not a prompt to initiate it (cf. [Sperber and Wilson 1986](#); [Cohen and Kehler 2021](#)).

These considerations suggest to us that, even before taking eliciture into account, the cognitive mechanisms underlying extrasemantic enrichment and its interpretation are often better understood as more top-down, and less triggered/bottom-up, than has been appreciated. And, indeed, for reasons we have discussed in §3, this conclusion seems to us yet more pressing once we take seriously the properties of intrasentential coherence-based enrichments. Though it is perhaps surprising that a high level characterization of a pragmatic phenomenon like eliciture would turn out to have implications for what might have seemed to be more implementation-level questions of cognitive processing, we don't see how this conclusion can be avoided.

6 Reconsidering costs in the theory of intersentential coherence

Let us take stock. In §1 we began by reviewing evidence of parallel coherence-based enrichments in inter- and intra-sentential examples, and in §2 we outlined BRACE, our characterization of the process through which information can be accommodated in service of the mandate to establish coherence between sentences intersententially. Such accommodations come at a cost, which then provides a method for calculating the overall cost of a proof of coherence: it is simply the sum of the costs of the accommodations.

As we argued in §3–§5, however, certain of the inferences that are leveraged for establishing coherence appear to have a ‘top-down’ character; hearers are inclined to draw them, at least provisionally, even when no interpretive mandate to do so exists. Otherwise, for example, we would never draw them when confronted with non-linguistic situations in the world; after all, there is no point to making inferences that come at a cost when one doesn’t have to. Similarly, we would never draw elicitures where there is no mandate to associate (causally or otherwise) contents from different parts of a clause that bear no direct syntactic relationship.

Taking these two ideas together, we are left with the conclusion that two types of enrichment must be distinguished in a theory of coherence establishment: those that are served up by our cognitive apparatus by virtue of the joint evocation of contents that can be related, and those that are specifically triggered by the mandate to establish coherence. Whereas both come at a cost in current conceptions of how coherence establishment works, we maintain that only those of the second type should.

Let us first revisit our core examples:

- (21) a. The company fired the manager. He embezzled money.
b. The company fired the employee who embezzled money.

As we discussed in §2, on the standard model, various enrichments are typically made when establishing the coherence of (21a) that correspond to terms in the antecedent of axiom (9): that the money mentioned in the second clause belonged to the company mentioned in the first clause, that the manager in fact works for the company, and that the company did in fact discover the manager’s crime. These all came at a cost. But the fact that the same causal inferences are drawn in (21b) suggests that these enrichments are not triggered by the mandate to establish coherence — as far as the mandate is concerned, they come for free. In light of how (21b) is

typically interpreted, one would conclude that the same inferences would come about for (21a) even if there was no mandate to establish coherence. As such, the overall cost of establishing the coherence of (21a) should not be burdened by the costs of these enrichments. The cost of establishing coherence should include only those of inferences that are drawn in service of the mandate.

To see a case in which the cost of accommodations required to establish a causal relation should burden the cost of coherence establishment, consider examples (22a-b).

- (22) a. The student took a train to San Francisco. He likes donuts.
- b. The student who likes donuts took a train to San Francisco.

Assuming that the referring expression *the student who likes donuts* in (22b) is a felicitous expression for the speaker to have chosen in the context (e.g., the referent's name is unknown to the interlocutors, and the common ground is such that it provides the best way to successfully refer), (22b) is perfectly acceptable without any need to draw a causal eliciture that links the donut liking to the travel. That is, without some tip-off that a causal eliciture was intended by the speaker, a hearer will not pursue the fact that a causal generalization that would support a causal eliciture — e.g., that San Francisco is world-renowned for its donuts — is not part of her beliefs, and hence not in the common ground. On the other hand, this is precisely how our hearer would likely react when hearing (22a): *Wait a minute, is San Francisco known for their donuts or something?* The signaling of presupposition failure here is a direct reaction to the mandate to establish coherence; the candidate inference was *not* already served up. This accommodation should, therefore, figure in the cost assigned to the proof of coherence.

Finally, let's consider one more example pair, which are variants of an example from [Knott and Dale \(1994\)](#):

- (23) a. The student broke his leg. He likes plums.
- b. The student who likes plums broke his leg.

Here there can be little doubt that, exceptional contexts aside, (23b) will not result in a causal eliciture linking the plum liking to the leg breaking. After all, there is no reason to expect that the hearer's cognitive apparatus will serve a candidate inference up. The lack of such an eliciture, however, does nothing to threaten the utterance's felicity. This fact stands in stark

contrast to the situation for (23a), which sends the hearer searching for a way to establish coherence between the clauses. Possible explanations may come to mind via offline analysis — e.g., the story is about the student’s unsuccessful climb up a plum tree — but without our cognitive apparatus volunteering the necessary accommodations, they come at too high of a cost for the hearer to run with.

Consideration of the behavior of the foregoing examples as an ensemble leads us to conclude that the existence of eliciture forces a reconceptualization of how theories of coherence establishment should calculate the cost of a coherence explanation. In particular, the cost assigned to a proof of coherence should include only those accommodations that are drawn in service of the mandate that coherence establishment imposes.

7 On the priority of causal interpretations

The BRACE framework for coherence establishment that we outlined in §2 and variations on it have been associated in the literature with a longstanding puzzle: Why do hearers ever enrich passages to have a causal coherence construal like Explanation or Result, when such interpretations would seem to carry a higher cost than interpretations that are not so enriched?⁴ In this section, we suggest that the revelation concerning costs associated with proofs of coherence presented in the last section offers a new solution of this puzzle.

Let’s step through the issue by reconsidering our example:

(24) The company fired the manager. He embezzled money.

As we discussed in §2, several coherence construals are possible for (24). One is the analysis as an Explanation relation that we have already described in detail, according to which the embezzling was the reason for the firing. Another is as a Parallel relation, in which the two clauses are understood as each providing a partial answer to a common question, e.g. *What are some events that the manager was involved with today?*. As we have seen, the Explanation relation analysis comes at a cost, as the hearer needs to make the accommodations we previously discussed: that the manager worked for the company, that the money belonged to the company, that the crime was discovered, and that there is nothing else true that would block

⁴Assuming, of course, that they are devoid of causal connectives — e.g., *because*, and *as a result* — which explicitly instruct the hearer to infer a causal relation.

the applicability of axiom (9) to the situation at hand. On the other hand, establishing the Parallel relation might seem cheap by comparison: a hearer merely needs to accommodate that the speaker wishes to convey two events that involve the manager, with no further enrichments necessary. So why would a hearer strengthen the interpretation to a costlier causal one, which is no doubt what typically occurs for this example?

Several researchers have engaged with this question (Levinson 2000; Asher and Lascarides 2003; Sanders 2005; Pagin 2014). Whereas their respective accounts differ in their details, they all posit a principle that explicitly favors causal coherence relations over others when possible (cf. Levinson’s corollary to his enrichment rule imploring recipients to “Assume the richest temporal, casual and referential connection between described situation or events, consistent with what is taken for granted” (Levinson 2000, p. 114), Lascarides and Asher’s partial order on relations — e.g., Explanation > Background— as part of their Maximize Discourse Coherence principle (Asher and Lascarides 2003, p. 231), Sanders’s “causality by default” hypothesis, according to which hearers will initially assume a causal connection exists, and will fall back on a different relation only when establishing causation isn’t possible (Sanders 2005, p. 113), and Pagin’s Principle of Pragmatic Enrichment, which calls for coherence to be strengthened to the highest available degree, with causality being the highest (Pagin 2014, p. 76)). With respect to such accounts, it is important to acknowledge that principles that say hearers should “infer causality when they can” are only helpful when there’s an underlining theory of when they can, and here the aforementioned proposals give at best limited guidance.

To see the point, let us revisit example (23a), repeated here as (25).

(25) The student broke his leg. He likes plums.

As we have discussed, there are possible scenarios that would support an Explanation relation; for instance, it could be a story about a student’s unsuccessful climb up a plum tree. Here only a small number of accommodative inferences are necessary to connect the two eventualities expressed: The student likes plums (second sentence), tried to climb a plum tree to get some (accommodation), fell out (accommodation), and broke his leg (first sentence). The question for the theories just surveyed is why hearers don’t in fact draw this chain of inference in service of the mandate to establish causality. It can’t be the fact that accommodative inferences are required, including ones involving entities and/or events that have not been evoked by the discourse; it is the norm for such inferences to

be made to establish coherence, as we saw in our discussion of example (7). Intuitively, the problem is that the inferences required for (25) are simply *too* costly; too much needs to be assumed without further contextual support to connect the sentences in this way. Without an account capable of determining when inferring causality costs too much, principles that posit that hearers should “infer causality when they can” are of little predictive value.

The analysis we have developed provides a different solution to the puzzle presented by causal priority. Recall that in the BRACE framework we described in §2, there are no prejudgments about what relationships are preferred over others, i.e., there are no overlaid preference rankings among them. The model contains only a cost-based system in which the lowest-cost interpretation(s) win, and hence, according to the prevailing wisdom, we’re left without an explanation for causal priority. However, as we argued in §6, typically many of the enrichments drawn in support of causal interpretations come about in absence of a mandate to establish coherence and hence, as far as the cost of satisfying the mandate goes, they come for free, since they are already available when the time comes to establish causal coherence. Since only those inferences necessary to establish coherence come with a cost, we have an explanation for how the establishment of a causal coherence relation can ultimately come at a low cost, and hence win the competition among competing coherence analyses. For instance, it seems perfectly plausible that, in a case such as (24), the costs of establishing Parallel — which still requires that the contextual relevance of a question such as *What are some events that the manager was involved with today?* be inferred or accommodated — would come at a greater cost than the establishment of Explanation, where the most central of the necessary accommodations come at a low or even no additional cost. On the other hand, if the aforementioned question had been made explicit in the discourse, with (24) offered as an answer, then Parallel now comes at a low cost since the question that lends it coherence no longer needs to be accommodated.

We believe that the account we have offered provides a more natural solution to the puzzle of causal priority, in deriving from more primitive and independently-motivated principles (the non-mandated causal inferences that an agent’s cognitive apparatus serves up in the non-linguistic world, and the evidence provided by eliciture that similiar non-mandated inferences carry over to linguistic descriptions of the world) and therefore not requiring special rules. It is worth being clear that our proposed explanation of the priority of causal interpretations helps itself to, and does

not amount to an explanation of, the apparent fact that our psychologies are disposed to attribute causal relations to events. Moreover, we are (like other theorists) without a predictive theory of just what conditions underwrite such causal inferences. But we don't see these points as undercutting the significance of our proposal. After all, the psychological phenomena we are appealing to are uncontroversial, and whatever the right account of these phenomena is, we are claiming that it, alongside the evidence for its application to natural language interpretation provided by eliciture, gives us what is needed to explain the priority of causal interpretation within a coherence establishment framework. Once these elements are in place, there's no need to stipulate specific principles to resolve an additional special version of the same puzzle concerning language interpretation, as per other accounts.

Beyond this, there is an additional fact that favors our analysis over competitors: our view, but not the others we have discussed, predicts the potential for passages to be ambiguous between causal and non-causal coherence construals. Consider (26):

(26) Pence became really angry, and Trump threw a tantrum.

There are two salient coherence construals for (26). The first is a (relatively unenriched) Parallel construal: Trump and his Vice-President were each overcome with negative emotions, possibly (but not necessarily) due to the same external stimulus. The second is an (enriched) Result construal: Trump, being unsympathetic to Pence's show of emotion, threw a tantrum as a result. The important observation is that the existence of a salient causal interpretation does not impinge on the ability to construe the passage under a Parallel interpretation. Both, outside of a larger context, appear to be salient coherence construals. The question for the analyses described above is how a hearer would ever infer a Parallel construal: since the mandate to infer a causal connection succeeds, there is no reason to back off to a non-causal coherence relation.

Whereas we can of course offer no analysis capable of predicting the existence of ambiguity for this particular example, conceptually the fact that examples like (26) can have more than one salient coherence construal provides no conundrum for our account, since it contains no overlaid ranking of coherence relations. The close similarity in the events described naturally evokes a question that each sentence could provide a partial answer to — e.g., *How did the President and Vice-President react?* — which renders Parallel a low-cost relation to infer. The Result relation is likewise

low-cost, in light of a salient causal generalization that connects one person getting angry with another throwing a tantrum. So like ambiguities that reside at the lexical, syntactic, and semantic levels, (26) is simply a case in which two salient interpretations compete at the level of discourse interpretation. There is no need for a preference ranking of coherence relations to be invoked.⁵

8 Conclusion

The phenomenon of eliciture, we contend, sheds new light on the nature of coherence and capacities for coherence establishment central to language understanding, and ultimately motivates a reorientation of the BRACE framework for coherence establishment. Our argument for this reorientation came in several steps.

We began by noting that elicitures result from coherence-driven pragmatic enrichments made without any sort of linguistic mandate, and took this as evidence of a top-down component in such enrichments: the evocation of ideas suggestive of causal and other types of relationships leads hearers to infer that those relationships hold. Moreover, we have taken pains to note that such non-mandated, ampliative inferences parallel the inferences agents draw to establish the coherence of events they encounter non-linguistically in the world. This parallelism motivates a parsimonious account on which the kinds of coherence-driven inferences agents draw regarding non-linguistically-presented events are applicable, similarly, to the interpretation of natural language descriptions of such events.

The parsimony of this style of account runs yet deeper, in that it offers a unified treatment of elicitures, which operate intrasententially, and apparently parallel inferences arising from coherence establishment processes that apply intersententially. And, indeed, it allows us to understand the most important difference between intrasentential and intersentential coherence-based enrichments (*viz.*, that such enrichments are optional in intrasentential cases but mandatory in intersentential cases) in terms of the interaction between a single inferential mechanism and the distinct felicity requirements of sentences and discourses.

⁵Pagin (2014) discusses an example similar to (26) in a response to a reviewer's comment, and says "the greater availability of the [Parallel] reading (no enrichment is needed) is balanced by the higher degree of coherence of the second [Result]. Ambiguity is then predicted" (p. 81). Unfortunately, we don't see what independent grounds could justify Pagin's stipulation that the Parallel reading has "greater availability" than the Result reading.

We've also argued that the relationship between eliciture and intersentential coherence establishment has ramifications for current conceptions of coherence establishment. Whereas the BRACE framework described in §2 assigns costs to all ampliative inferences made in the service of establishing coherence, consideration of eliciture makes clear that two types of inferences need to be distinguished: those that are served up by our cognitive apparatus top-down despite the absence of a mandate to draw them, and those that occur bottom up, being initiated specifically by the need to establish coherence. We have argued that only the latter sort of inference should figure in the assessment of the cost of a candidate coherence construal. That is, not all of the inferences leveraged in a proof of coherence result specifically from the need to establish coherence; otherwise we would be left without an explanation of the enrichments associated with eliciture.

Finally, the recognition of this distinction led to an explanation of a puzzle concerning the priority of causal interpretations that has long hounded coherence theory: why do hearers establish causal coherence construals of passages when, *prima facie*, such construals come at a greater cost than less enriched ones? Previous authors have answered this puzzle by specifying additional principles and/or rankings among relations sitting alongside the coherence processor. We've argued that such stipulations are unnecessary. For, given that causal inferences often fall outside the class of inferences mandated by the need for coherence establishment, they will often make no contribution to the cost of causal proofs of coherence. This is just to say that, given the psychological operations at work in such cases, inferences to causal coherence relations will often come at low cost for reasons that are not specific to language understanding. While, admittedly, neither we nor anyone else has a quantitative mechanism capable of testing this prediction, the explanation we are offering is well-motivated, and draws only on uncontroversial underlying psychological phenomena, and so avoids the need to hard code into the theory of language interpretation a preference for causal interpretations, or even treat causal relations as special in any way.

Ultimately, we find that eliciture — which may at first blush appear to be an innocuous side phenomenon associated with well-known processes for establishing the coherence of discourses — in fact has compelling and far-reaching ramifications for how those processes are characterized. We have made similar points elsewhere in other domains, arguing that elicitures do not square with current conceptions of pragmatic enrichment (Cohen and Kehler 2021), yield an alternate explanation for certain cases of apparent

pragmatic intrusion into truth conditions (Kehler and Cohen 2016), and are problematic for accounts that reduce pragmatic enrichment to convention and disambiguation (Kehler and Cohen 2018). Though one might initially be inclined to retrofit eliciture into current views of coherence establishment, we urge resisting that temptation. If, as we believe, facts about eliciture sit uneasily with current theoretical views, this gives us reason for revising the theory, not the facts.

References

- Asher, Nicholas and Alex Lascarides (2003), *Logics of Conversation*, Cambridge University Press, Cambridge.
- Clark, Herb H. (1975), "Bridging," in *Theoretical issues in natural language processing*, ed. by R. C. Schank and B. L. Nash-Webber, New York: Association for Computing Machinery, pp. 175-196.
- Cohen, Jonathan and Andrew Kehler (2021), "Conversational Eliciture," *Philosophers' Imprint*, 21, 12, pp. 1-26.
- Gundel, Jeanette K., Nancy Hedberg, and Ron Zacharski (1993), "Cognitive Status and the Form of Referring Expressions in Discourse," *Language*, 69, 2, pp. 274-307.
- Halliday, M.A.K. and Ruqaiya Hasan (1976), *Cohesion in English*, English Language Series, Title No. 9, Longman's, London.
- Hawkins, John A. (1984), "A Note on Referent Identifiability and Co-Presence," *Journal of Pragmatics*, 8, pp. 649-659.
- Hobbs, Jerry R. (1979), "Coherence and Coreference," *Cognitive Science*, 3, 1, pp. 67-90.
- (1990), *Literature and Cognition*, CSLI, Stanford.
- (2010), "Clause-Internal Coherence," in *Constraints in Discourse 2*, ed. by Peter Kühnlein, Anton Benz, and Candace L. Sidner, John Benjamins, Amsterdam, pp. 15-34.
- Hobbs, Jerry R., Mark E. Stickel, Douglas E. Appelt, and Paul Martin (1993), "Interpretation as Abduction," *Artificial Intelligence*, 63, pp. 69-142.
- Kehler, Andrew (2002), *Coherence, Reference, and the Theory of Grammar*, CSLI, Stanford.
- (2015), "Reference in Discourse," in *The Handbook of Contemporary Semantic Theory*, ed. by Shalom Lappin and Chris Fox, 2nd ed., Wiley-Blackwell, pp. 625-654.

- Kehler, Andrew and Jonathan Cohen (2016), "On the Presuppositional Behavior of Coherence-Driven Pragmatic Enrichments," *Semantics and Linguistic Theory*, 26, pp. 961-979.
- (2018), "On convention and coherence," in *Beyond Semantics and Pragmatics*, ed. by Gerhard Preyer, Oxford University Press, Oxford, pp. 261-283.
- Kehler, Andrew and Hannah Rohde (2019), "Prominence and Coherence in a Bayesian Theory of Pronoun Interpretation," *Journal of Pragmatics*, 154, pp. 63-78.
- Knott, Alistair and Robert Dale (1994), "Using Linguistic Phenomena to Motivate a Set of Coherence Relations," *Discourse Processes*, 18, 1, pp. 35-62.
- Kronfeld, Amichai (1990), *Reference and Computation: An Essay in Applied Philosophy of Language*, Cambridge University Press, Cambridge.
- Levinson, Stephen C. (2000), *Presumptive Meanings: The Theory of Generalized Conversational Implicature*, MIT Press, Cambridge, MA.
- Longacre, Robert E. (1983), *The Grammar of Discourse*, Plenum Press, New York.
- Mann, William C. and Sandra A. Thompson (1987), *Rhetorical Structure Theory: A Theory of Text Organization*, tech. rep. RS-87-190, Information Sciences Institute, University of Southern California, Marina del Rey, California.
- Martin, James R. (1992), *English Text: Systems and Structure*, Amsterdam: Benjamins.
- Pagin, Peter (2014), "Pragmatic enrichment as coherence raising," *Philosophical Studies*, 168, pp. 59-100.
- Prince, Ellen F. (1981), "Toward a Taxonomy of Given-New Information," in *Radical Pragmatics*, ed. by Peter Cole, Academic Press, New York, pp. 223-255.
- Rohde, Hannah, Roger Levy, and Andrew Kehler (2011), "Anticipating Explanations in Relative Clause Processing," *Cognition*, 118, pp. 339-358.
- Sanders, Ted J. M. (1997), "Semantic and Pragmatic Sources of Coherence: On the Categorization of Coherence Relations in Context," *Discourse Processes*, 24, pp. 119-147.
- (2005), "Coherence, causality and cognitive complexity in discourse," in *Proceedings of the First International Symposium on the Exploration and Modeling of Meaning*, ed. by M. Aurnague and M. Bras, University of Toulouse-le-Mirail, pp. 31-46.
- Sperber, Dan and Deirdre Wilson (1986), *Relevance: Communication and Cognition*, Harvard University Press, Cambridge, MA.