

Annual Review of Linguistics

Coherence Establishment as a Source of Explanation in Linguistic Theory

Andrew Kehler

Department of Linguistics, University of California, San Diego, La Jolla, California, USA;
email: akehlr@ucsd.edu

Annu. Rev. Linguist. 2022. 8:123–42

The *Annual Review of Linguistics* is online at
linguistics.annualreviews.org

<https://doi.org/10.1146/annurev-linguistics-011619-030357>

Copyright © 2022 by Annual Reviews.
All rights reserved

**ANNUAL
REVIEWS CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Keywords

coherence relations, discourse structure, anaphora, ellipsis, Coordinate Structure Constraint, tense, accent placement, QUDs

Abstract

The primary goal of coherence theory is to provide an explanation for the coherence properties of discourse: what properties distinguish a discourse from a mere collection of utterances, and what drives comprehenders to draw inferences in service of establishing coherence. However, the importance of coherence theory goes well beyond that; it also plays a crucial role in theories of a variety of discourse-dependent linguistic phenomena. This article surveys some ways in which coherence theory has been leveraged in this way, appealing to both Relational analyses and Question-Under-Discussion models. Theories of coherence establishment should therefore have a place in the linguist's toolbox as a source of explanation in linguistic theory.

1. INTRODUCTION

Suppose that a linguist (call him Andy) receives a call from an old colleague he has not seen in years (call him Jerry). Jerry says that he is in town for the evening and wonders if Andy would like to grab dinner and catch up. In response, Andy utters:

- (1) I won't be able to go to dinner tonight. I have a paper to finish for the *Annual Review of Linguistics*.

All is well until Jerry, after eating alone and then stopping at a brewery afterward, sees Andy hanging out there with a group of friends, clearly having been there a while, and hence presumably never having intended to work on his article. Feeling that he was lied to, Jerry confronts Andy, but Andy objects, pointing out that his two utterances were both true: He indeed was not able to go to dinner that night, and he does have to work on his paper. He never said that he was going to work on the paper that night, after all.

Whereas Andy may be technically correct, his position is not likely to garner much sympathy. For Jerry did what language users typically do when interpreting a multi-utterance discourse: They do not interpret co-occurring sentences as expressing independent statements but instead assume that such statements are relevant to one another. Of the various possibilities, in this case Jerry assumed that the two clauses were related by what various authors (Hobbs 1990, Lascarides & Asher 1993, *inter alia*) have called Explanation, according to which the second utterance describes a cause of, or reason for, the eventuality described by the first. Establishing this relation thus requires that a sequence of inferential steps be carried out to identify a causal link between the contents of the two clauses. As I discuss in greater detail below, a natural inference for Jerry to make is that Andy intends to work on his paper that night, hence tying him up during the time at which he and Jerry could otherwise have dinner. Jerry further reasons, under his assumption that Andy is a cooperative language user, that Andy intended for him to draw those inferences, and hence to convey their content to him. Hence his lack of sympathy for Andy's protests: Without having drawn the inferences he did, Jerry would be at a loss regarding how Andy's two utterances are related, a matter quite separate from whether each utterance is true as stated.

The primary goal of coherence theory is to provide an explanation for the coherence properties of discourse: what properties distinguish a discourse from a mere collection of utterances, and what drives comprehenders to draw inferences in service of establishing coherence. Describing how one approach to coherence—the Relational view—accomplishes these goals is the topic of Section 2. An equal goal of the current contribution is to describe how coherence establishment can provide a source of explanation in linguistic theory. Some of my own work has focused on this question, which I briefly summarize in Section 3. A competing theory of coherence that has also shown explanatory power is the Question-Under-Discussion (QUD) model, which I briefly describe and compare with the Relational view in Section 4. I conclude in Section 5.

2. COHERENCE RELATIONS

The primary desiderata of a theory of coherence are twofold. First, theories should explain how discourses gain their coherence, such that they are perceived differently from a series of unrelated utterances. For instance, we want to capture the reasons for which example 1 comes across as coherent compared with example 2, which will typically be judged as incoherent outside of an unusual context:

- (2) I won't be able to go to dinner tonight. I went to high school in Pennsylvania.

Second, we want an explanation of how the inference of unstated propositions (i.e., pragmatic enrichment) results from the coherence establishment process. This would include Jerry's inference upon hearing example 1 that Andy intended to work on his paper that evening.

According to the Relational view of coherence (Hobbs 1979, Mann & Thompson 1987, Polanyi 1988, Hobbs 1990, Sanders et al. 1992, Kehler 2002, Asher & Lascarides 2003, Wolf & Gibson 2006, *inter alia*), there exists an inventory of relations (henceforth, coherence relations) that are each associated with constraints that must be met for the relation to be established. It is the identification of these relations within a discourse that gives rise to its structure. This identification process will typically require inferences to satisfy the constraints associated with the relations being established, which will in turn give rise to pragmatic enrichments.

As such, a Relational theory of coherence necessarily comes with an inventory of relations that a speaker can employ in structuring her discourse, and which the addressee is expected to be able to identify from her utterances and other contextual clues. A wide variety of relation sets have been on offer, starting with several early proposals (Hobbs 1979, Mann & Thompson 1987, Hobbs 1990, Sanders et al. 1992, *inter alia*), and others since. The obvious question, then, is how one goes about evaluating the adequacy of such inventories. Sanders et al. (1992) propose two theoretical virtues, which often come into conflict: descriptive adequacy and cognitive plausibility. First, obviously theories should be descriptively adequate—ultimately able to describe the coherence structure of any given passage that one might encounter in the wild. The Rhetorical Structure Theory (RST) of Mann & Thompson (1987) is an early example of an inventory of this sort, where an extensive analysis of a corpus of data yielded a proposed set of 23 relations. A more recent inventory of relations was developed for the Penn Discourse Treebank (PDTB) annotation effort (Prasad et al. 2008, Webber et al. 2019), which as of this writing utilizes 36 relations. The third version of this corpus contains over 53,000 relation tokens annotated in *Wall Street Journal* newspaper articles. Unlike RST, the PDTB relations are classified hierarchically in terms of four different relation types (Temporal, Contingency, Comparison, and Expansion) at the highest level.

Building a relation inventory purely bottom-up is not completely satisfactory, however, in that it leaves open the questions of why this particular set of relations is the right one as opposed to other possibilities, what theoretical grounding underlies the inventory (versus merely being a laundry list), and how the inventory can be acquired during language acquisition. Such questions lead us to Sanders et al.'s (1992) second desideratum, cognitive plausibility. One would expect that there are fundamental cognitive principles at work that will serve both to constrain the set of possible relations and to provide an explanation for why a particular set of relations is to be preferred to one that contains more, fewer, or different relations. Hobbs (1990, p. 101) offers the following take:

It is tempting to speculate that these coherence relations are instantiations in discourse comprehension of more general principles of coherence that we apply in attempting to make sense out of the world we find ourselves in, principles that rest ultimately on some notion of cognitive economy. . . . Recognizing coherence relations may thus be just one way of using certain very general principles for simplifying our view of the world.

To put it another way, it is undeniable that the human cognitive apparatus comes with ways of establishing the coherence of our surroundings, such that we make inferences (e.g., causal ones) that go beyond the information made available to our perception. If we use this apparatus to interpret the situations that we encounter in the world, it stands to reason that we might employ the same machinery when we interpret natural language descriptions of those situations and that a speaker might take advantage of this when conveying her message.

What are the types of relationships we use to connect our perceptions, and more generally, thoughts? Many takes on this question have been offered. Here I briefly describe the take that I have pursued in past work, as it lays the groundwork for the discussion of linguistic phenomena in Section 3. Specifically, the account described in Kehler 2002 draws on the philosophy of David Hume, who famously ponders this question in his *Inquiry Concerning Human Understanding*, and ultimately posits only three such relationships: Resemblance, Contiguity in Time or Place, and Cause–Effect [Hume 1955 (1748), p. 32]. Taking my cue from Hobbs (1990), I have argued (Kehler 2002) for an inventory of coherence relations that are classified in terms of these three more general principles. The Cause–Effect category contains four relations—Result, Explanation, Violated Expectation, and Denial of Preventer—instances of which are given in examples 3–6, respectively:

- (3) Donald Trump is by far the least experienced in government of all the presidential candidates. He will probably not win the election. [Result]
- (4) Donald Trump will probably not win the election. He is by far the least experienced in government of all the presidential candidates. [Explanation]
- (5) Donald Trump is by far the least experienced in government of all the presidential candidates, but he will probably win the election. [Violated Expectation]
- (6) Donald Trump will probably win the election, even though he is by far the least experienced in government of all the presidential candidates. [Denial of Preventer]

What all of these relations have in common is that their establishment engages the presupposition of a defeasible causal relationship—in this case, that not being an experienced politician might ordinarily lead to not being elected president. The effect of the Result and Explanation relations is to establish that this general causal connection holds in the particular case being described. The difference between them is clause order; in Result relations the cause precedes the effect, whereas in Explanation relations the opposite is the case. The role of the Violated Expectation and Denial of Preventer relations, in contrast, is to express that the highlighted causal relationship does not in fact hold in the situation being described. The difference between the two relations is again clause order. Therefore, while instantiated as four distinct relations here, they can ultimately be seen as variants of the same underlying relation, with clause order and polarity of the effect as parameters (see also Sanders et al. 1992). In this sense, the four relations can be seen not as atomic but instead as molecular, all instantiating a single and more general cognitive principle. This feature distinguishes the analysis from those that feature laundry list inventories of coherence relations that lack a deeper theoretical grounding.

In contrast to the Cause–Effect relations, establishing Resemblance relations is a fundamentally different process, one that requires that commonalities and contrasts among corresponding sets of parallel relations and entities be recognized. The canonical case of a Resemblance relation is Parallel, which is exemplified in passage 7:

- (7) Biden has been touting his years of experience from being a Senator and Vice President. Sanders has stressed the fact that he’s leading a revolution.

Example 7 is characterized by a pair of contrastive topics—Biden and Sanders—and a pair of parallel predications that together can be analyzed in terms of a more general common topic (e.g., “the things *X* has been doing to advertise *X*’s leadership skills”).

Other relations in the Resemblance category follow this general pattern. For instance, the Exemplification relation requires that the same type of parallelism be established, in this case between a general statement and an instantiation of it:

- (8) Presidential candidates stand much to gain by touting examples of their leadership skills. Giuliani has been reminding everyone about September 11th at every conceivable opportunity.

Note that the establishment of Exemplification here requires that the hearer know (or accommodate) that Giuliani was a presidential candidate and, likewise, that he displayed leadership skills during September 11. Other relations in this class include Contrast, Generalization, Exception, and Elaboration (for further details, see Kehler 2002).

Finally, the third class of relation is Contiguity, instantiated by the Occasion relation:

- (9) Trump traveled to Vietnam. He met with the leader of North Korea, Kim Jong-un.

Occasion allows a speaker to describe a situation centered around a system of entities by using intermediate states of affairs as points of connection between partial descriptions of that situation. As such, the inference process that underlies Occasion attempts to equate the initial state of the eventuality denoted by the second utterance with the final state of the first. Additional inferences that result from establishing Occasion in passage 9 come effortlessly—for instance, that Trump met with Kim in Vietnam and that this occurred soon after his travel. It is therefore worth noting that merely establishing temporal progression between events is not enough (pace Halliday & Hasan 1976, Longacre 1983, *inter alia*).

With an inventory of relations in hand, we now ask how they can be used to address the two goals laid out in the Introduction: the manner in which they can be used to explain the coherence properties of discourses, and how the coherence establishment process itself can give rise to pragmatic enrichments. The two best-known and well-developed analyses are those of Hobbs et al. (1993), who posit a mechanism based on a cost-based notion of logical abduction, and Asher & Lascarides (2003, *inter alia*), whose approach is based on nonmonotonic deduction. The reader is referred to these works for further detail. Both systems explain enrichments by coupling an inference system with logical statements that (*a*) express the constraints that have to be met to establish each coherence relation and (*b*) represent axiomatic knowledge about the world. For instance, when interpreting passage 1, it is plausible to believe that a typical hearer would have a variety of relevant defeasible axioms at his disposal: that if someone needs to do something, they may take on the intention to do it; that if someone has work to do, they might stay at work or home to do it; that people cannot be at two places at once; and so forth. Such axiomatic knowledge, when instantiated with the details specified by the utterances being related, can then be used to form a chain of inference between the two clauses in passage 1, thereby establishing the Explanation relation: Andy has a paper to finish, therefore Andy needs to work on it tonight, and since he cannot be doing two things in different places at once, Andy cannot go to dinner with Jerry. The pragmatic enrichments noted above will fall out from this procedure—for instance, that Andy intends to work on his paper this evening—despite not being asserted by the passage itself.

Of course, the process by which this chain of inference occurs—presumably stable enough such that a speaker can rely on her hearer to make the inferences that she intends him to—remains mysterious. There exists no workable theory to answer many questions: what the structure of the knowledge base should be, what axioms should be in it, how much various assumptions (i.e., enrichments) should “cost” to make, how one selects the correct proof of coherence out of the many that are possible, and so forth. As such, there is no predictive theory that can tell us what coherence relation hearers will establish for arbitrary examples. But such inabilities are independent of the larger question of whether the Relational approach captures how interlocutors establish and manage the coherence of discourses. As such, the failure to answer all of the questions associated with a theory should not cast doubt on the theory itself. And further, as discussed below, the lack of a fully predictive theory does not mean that coherence theory cannot be profitably used in accounts of linguistic phenomena.

Having articulated a set of relationships that can hold between sets of utterances, we are naturally led to the question of how they can be used to offer a coherence criterion on larger discourses. This question also remains an open one. The debate has historically centered around the data structures required to characterize discourse structure: that is, whether tree structures suffice or the expressive power of graphs is required. On the tree-based view (Polanyi 1988, Hobbs 1990, Asher & Lascarides 2003, *inter alia*), the coherence criterion is straightforward to state: Starting with the assumption that sentences are discourse segments, larger segments (all the way up to the entire discourse) result by relating smaller segments by a coherence relation. In contrast to sentence structures, the “syntactic” component for building discourse structures is therefore trivial; the real work lies in the process of establishing coherence relations. That said, relatively little work has been carried out on the difficult task of explaining how coherence relations apply to collections of larger segments of texts.

The tree-based conception has been argued by various researchers to be too simplistic, however. For instance, Wolf & Gibson (2006) argue that the full power of a graph structure is necessary. On their analysis, the coherence criterion is less restrictive, requiring only that discourses be characterizable as a connected graph that, unlike tree structures, allows for nodes with multiple parents and for crossing dependencies. Webber et al. (2003) offer a third view: that while multiple parents are permissible in discourse structures, the need for crossing dependencies is eliminated when one distinguishes between discourse relationships that are structural and those that are anaphoric—the latter of which are not represented in the discourse structure. A full exploration of the issues is beyond the scope of this review; suffice it to say that future work is necessary to resolve these questions and many others regarding theories of discourse coherence and structure.

3. COHERENCE THEORY AS AN EXPLANATORY DEVICE

A core thesis of the Relational view is that the establishment of coherence is an automatic, inevitable component of discourse processing. That is, just as addressees naturally attempt to parse a sequence of words with which they are presented into a syntactic structure, they just as automatically attempt to analyze a sequence of sentences with which they are presented into a discourse structure. The crucial difference is that whereas rules of grammar are responsible for building larger constituents out of smaller ones in syntax, it is the establishment of coherence relations between discourse segments that is responsible for building larger segments at the level of discourse.

If coherence establishment is a pervasive aspect of discourse interpretation, then it would be unsurprising if one were to find that it affects the manner in which linguistic phenomena are interpreted, particularly those that rely heavily on context for their interpretation. This observation entails that adequate theories of such phenomena will require not just an account of their linguistic properties but also an analysis of how those properties interact with coherence establishment to yield the patterns of behavior witnessed in the data. As an illustration, here I briefly survey examples in which coherence establishment has been brought to bear in my own theories, focusing on the acceptability and interpretation of a range of specific linguistic phenomena. These of course are not the only problems to which coherence theory has been applied (see, e.g., Asher & Lascarides 2003 and their papers since for discussions of the relationship between coherence establishment and topics including presupposition, implicature, and indirect speech acts).

3.1. Pronoun Interpretation

Pronoun interpretation may be the first linguistic phenomenon to which coherence theory has been applied. Hobbs (1979) famously developed an account whereby pronoun interpretation is not even an independent process but instead results as a by-product of more general reasoning

about the most likely interpretation of utterances in a discourse context, including the coherence relationships that they participate in. Pronouns are modeled as free variables in logical representations, which become bound during these inference processes; potential referents of pronouns are therefore those that result in valid proofs of coherence. The classic example shown in examples 10a and b provides an illustration:

- (10) The city council denied the demonstrators a permit because. . .
 - (10a) . . .they feared violence.
 - (10b) . . .they advocated violence.
- (Adapted from Winograd 1972, p. 33)

In Hobbs's (1979) account, the correct assignment for the pronoun in each case falls out as a side effect of the process of establishing the Explanation relation, which here is signaled explicitly by *because*. It seems plausible to posit that language users have general causal knowledge that, when instantiated with the event participants in the example, would support the following inference: If the city council fears violence, the demonstrators advocate violence, and a permit would enable the demonstrators to bring about violence, then this might cause the city council to deny the demonstrators a permit. With this knowledge, the Explanation relation can be straightforwardly established for passage 10a, but with a caveat: The pronoun *they* has to refer to the city council since our causal rule says that they have to be the ones doing the fearing. Analogously, Explanation can be likewise established for passage 10b, but in this case *they* has to refer to the demonstrators since our rule applies only if they are doing the advocating. So by instantiation of the unbound variable as specified by the rule, the correct referent for the pronoun is identified as a by-product of establishing coherence.

A more recent line of experimental work (Kertz et al. 2006, Rohde et al. 2007, Kehler et al. 2008, Kaiser 2011, Kehler & Rohde 2013, Kaiser & Cherqaoui 2016, Kehler & Rohde 2019, Zhan et al. 2020, inter alia) has further established the role that coherence establishment plays. For example, Rohde et al. (2007) elicited passage completions that utilized stimuli containing a transfer-of-possession context sentence followed by an ambiguous pronoun prompt, as in passage 11:

- (11) John handed a book to Bob. He _____

In such sentences, the subject fills the Source thematic role, and the object of the preposition fills the Goal role. Participants were asked to provide a natural completion to the pronoun prompt, in one of two conditions. In one version, the participants were asked to write a completion that answered the question *What happened next?* This condition was expected to elicit more continuations participating in an Occasion relation, which were known from previous studies to be referentially biased toward the Goal entity. In the other condition, participants were to answer the question *Why?*, which was expected to yield more Explanation relation continuations, which were known to be biased toward the Source. The outcome was as predicted. On the one hand, the referential biases associated with each coherence relation were consistent between the two conditions. On the other hand, the distribution of coherence relations varied, and as a result, the overall pattern of pronoun interpretations shifted dramatically: There were significantly more pronoun references to the Source in the *Why?* condition than in the *What happened next?* condition. Since the stimuli themselves were identical between conditions, the difference can only be attributed to the different distribution of coherence relations.

Subsequent work, however, has painted a more complicated picture than that suggested by Hobbs (1979), one in which coherence establishment affects pronoun interpretation only

indirectly. According to the Bayesian model of pronoun use (Kehler et al. 2008, Kehler & Rohde 2013), a comprehender interprets a pronoun by reverse-engineering the speaker's referential intentions by combining his estimates of the speaker's production biases (i.e., the likelihood that she will use a pronoun to refer to each potential referent) and his estimates of the probability that the speaker will mention a potential referent at that point in the discourse regardless of the form of referring expression that she chooses (called the “prior”). Research has suggested that semantic and pragmatic factors, including expectations about the ensuing coherence relation, primarily condition the prior, whereas production biases appear to be grammatical (e.g., grammatical role obliqueness) and/or information structural (e.g., topicality) in nature (Rohde 2008, Fukumura & van Gompel 2010, Rohde & Kehler 2014; but cf. Rosa & Arnold 2017). On the Bayesian model, therefore, pronouns do not function merely as unbound variables in logical forms but rather carry their own referential biases. Coherence establishment instead affects predictions about what entities will be mentioned next in the discourse, which, along with production biases, in turn affects how pronouns will be interpreted.

3.2. Ellipsis

When contextual conditions are right, natural languages permit speakers to omit linguistic material that expresses information that they nonetheless intend to communicate to their hearers, via a class of phenomena collectively known as ellipsis. One such form is verb phrase (VP) ellipsis, illustrated in example 12:

- (12) Joe looked into the problem, and Don did too.

The stranded auxiliary in the second clause (henceforth referred to as the ellipsis clause) marks a vestigial verb phrase, a meaning for which must be recovered by the hearer. In many cases, the recovery of the meaning is enabled by the occurrence of another linguistic expression (the antecedent), which in this case is located in the first clause.

Perhaps the most active area of debate in the VP-ellipsis literature questions the conditions under which a representation of an utterance may serve as a suitable antecedent for interpreting a subsequent ellipsis—particularly, whether VP-ellipsis interpretation directly engages the syntactic structure of the antecedent clause or whether it operates more as a referential expression that engages instead with a mental representation of an eventuality. The question is highly vexed. On the one hand, critics have pointed to cases like examples 13 and 14 in support of a syntactic theory:

- (13) # This problem was looked into by Joe, and Donald did too. [look into the problem]
 (14) # Donald defended himself and Mike did too. [defend Donald]

Specifically, example 13 sounds degraded, plausibly due to the voice mismatch between the antecedent and ellipsis clauses. Similarly, many informants have a strong bias against a strict reading of example 14—that is, to read it as necessarily meaning that Mike defended himself (and not Donald), as would be predicted on accounts that require strict syntactic identity in light of the reflexive pronoun.

Examples 15 and 16, on the other hand, provide evidence in the opposite direction. Specifically, example 15 is felicitous despite having the same voice mismatch as 13, and the strict reading of example 16 is readily accessible despite the same prospects for a Binding Theory violation as in example 14:

- (15) This problem was to have been looked into, but obviously nobody did. [look into the problem]
(Vincent Della Pietra, in conversation, cited in Kehler 1993, example 3)
- (16) Donald defended himself because Mike wouldn't. [defend Donald]
(Adapted from Dalrymple 2005, example 75a)

So we are faced with a mystery: What difference between these two sentence pairs could explain their varying behavior?

Kehler (2000, 2002) argues that the divergent judgments in these cases and a range of others are due to the nature of the coherence relations operative in each case: The two clauses in examples 13 and 14 are related by the Resemblance relation *Parallel*, whereas those in examples 15 and 16 are related by *Cause–Effect* relations (*Violated Expectation* and *Explanation*, respectively). According to this theory, VP-ellipsis is anaphoric, but it interacts with the coherence establishment process such that elided syntactic material is reconstructed when a Resemblance relation, but not a Cause–Effect relation, is established. Whereas I forgo further specifics, a property worth highlighting is that syntactic reconstruction is not performed to recover the meaning of the missing VP but instead is done as part of establishing a limited set coherence relations.

It is also argued by Kehler (2000, 2002) that the account naturally explains an otherwise mysterious fact about the gapping construction, which to my knowledge was first noticed by Levin & Prince (1986). Consider that example 17 has two salient coherence construals:

- (17) Trump became upset and Pence became downright angry.
- (18) Trump became upset and Pence, downright angry.

First, it has a *Parallel* reading, whereby Trump and Pence both display negative emotions, perhaps (but not necessarily) in response to the same external stimulus. Second, it also has a *Result* reading, whereby Pence's anger is a direct response to Trump's getting upset. Intriguingly, the gapped version in example 18 has only the first of these readings.

This pattern is predicted by the analysis simply because gapping is not anaphoric; unlike VP-ellipsis, it does not have the hallmarks of referential behavior (e.g., the ability to refer cataphorically, find antecedents more than one sentence back, take split antecedents, and so forth). As such, it has no way to obtain an interpretation without the syntactic reconstruction process that is triggered by the establishment of *Parallel* relations. Since the process of establishing the *Result* relation does not recover elided syntax, example 18 is unable to receive that reading.

If this theory is on the right track, it is not the properties of VP-ellipsis and gapping themselves that are responsible for the conflicting data in examples 13–16 and 17–18, but instead the interaction between these properties and coherence establishment processes. That having been said, there has since been a robust debate focused primarily on the syntactic mismatch facts in examples 13 and 15, with some authors arguing for syntactic repair mechanisms (e.g., Arregui et al. 2006, Grant et al. 2012) or instead that the key factor is actually information structure rather than coherence (Kertz 2013). In addition, several studies have examined the claims of the coherence analysis experimentally (Frazier & Clifton 2006; Kim & Runner 2009, 2011; SanPietro et al. 2012; Kertz 2013). For a more extensive discussion, readers are referred to Kehler (2019).

3.3. Coordinate Structure Constraint

Another phenomenon to which coherence theory has been applied involves so-called extraction from coordinate structures. Ross (1967) famously argues for the existence of a *Coordinate Structure Constraint* (CSC) that, in part, bans the movement of constituents within a conjunct

to a position outside of the larger conjoined phrase. Cases like examples 19–21 motivate the constraint:

- (19) Trump signed the tax bill and vetoed the national emergency resolution.
- (20) # This is the tax bill that Trump signed and vetoed the national emergency resolution.
- (21) # What emergency resolution did Trump sign the tax bill and veto?

Whereas sentence 19 is perfectly acceptable, the variants in which extraction has taken place out of one of the two conjuncts per examples 20 and 21 are infelicitous, as predicted by the constraint.

However, no sooner does Ross (1967) posit the constraint than he starts listing counterexamples to it. One class of examples involves across-the-board (ATB) extraction, in which extraction occurs from all conjuncts, as in example 22:

- (22) This is the emergency resolution that Congress passed and Trump vetoed.

Such cases lead Ross to posit an ATB exception to the CSC. However, various authors have offered other examples that appear to be felicitous even though extraction occurs out of only a subset of the conjuncts. Examples 23–25 are from Ross (1967) himself, Goldsmith (1985), and Lakoff (1986), respectively:

- (23) What did you go to the store and buy?
- (24) How much can you drink and still stay sober?
- (25) That's the stuff that the guys in the Caucasus drink and live to be a hundred.

The relevant syntactic properties of these examples are similar to those in examples 20 and 21, involving conjoined VPs in each case.

However, the examples are not similar with respect to the coherence relations that are operative. Building on insights from Lakoff (1986) and others, Kehler (2002) argues that the constraints responsible for licensing extraction from coordinate structures are not grammatical but instead result from an information-structural constraint on extraction—in particular, one akin to the Topic-hood Condition of Kuno (1976, 1987)—and the manner in which the condition interacts with different coherence relations.

Roughly speaking, Kuno's constraint states that an extracted constituent must be a candidate for the sentential topic in the nonextracted counterpart of the sentence. For instance, whereas sentences 26 and 27 are perfectly acceptable, the acceptability of question 28 and the pragmatic oddity of question 29 result from the fact that people typically read, but do not lose, books because of their content:

- (26) I read a book about Marilyn Monroe.
- (27) I lost a book about Marilyn Monroe.
- (28) Who did you read a book about?
- (29) # Who did you lose a book about?

The question, then, is what constituents can serve as a topic for a set of conjoined clauses so as to allow for extraction. The answer depends on what coherence relation is operative. Specifically, by definition, clauses that participate in a Parallel relation are bound by a common topic; therefore, any constituent that serves as a topic must necessarily be mentioned in all of the conjuncts. Thus, *the emergency resolution* can be the topic of the nonextracted variant of example 22 but not of examples 20 or 21.

The situation is different for Occasion relations, however, as narratives that center around a topic need not mention it in every clause. Examples include clauses that set the scene or describe preparatory conditions, as with the first conjunct of example 23, repeated below as example 30. What results is a pattern such that, whereas not all clauses need to have constituents extracted from them, the last one does, lest the discourse no longer be about the extracted topic, as shown in examples 30–33:

- (30) What did you go to the store and buy?
- (31) # What did you go to the store, buy, and come home?
- (32) What did you go to the store, buy, come home, and devour in thirty seconds?
- (33) # What did you go to the store, buy, come home, devour in thirty seconds, and take a bath?

Importantly, establishing Occasion for cases like example 30 requires an inference—specifically, that the item in question was bought at the store. Without this inference, the conjoined phrases can only be construed as participating in a Parallel relation (the addressee did two independent things: went to the store, and bought something somewhere). On this construal, the example suddenly becomes unacceptable.

Finally, examples 24 and 25, repeated below as examples 34 and 35, participate in the Cause–Effect relations Violated Expectation and Result, respectively:

- (34) How much can you drink and still stay sober?
- (35) That's the stuff that the guys in the Caucasus drink and live to be a hundred.

Unlike the case for Occasion relations, Cause–Effect relations allow for cases in which no extraction occurs from the final clause, insofar as an extracted constituent from a clause describing a cause can serve as a topic for a larger structure that describes both the cause and its effect.

Viewing the data with respect to coherence, there does not seem to be much evidence for the existence of a CSC, as extraction from a coordinate structure is allowed in certain scenarios with all three types of coherence relation. The data suggest instead that an interaction between syntax, information structure, and coherence establishment determines when a constituent can be felicitously extracted from a coordinate structure.

3.4. Tense Interpretation

Another phenomenon to which coherence theory has been applied is tense interpretation and its role in establishing the temporal relationships among eventualities in discourse. Numerous researchers (Partee 1984, Hinrichs 1986, Nerbonne 1986, Webber 1988, *inter alia*) have sought to explain the temporal relations induced by tense by treating it as anaphoric, drawing on Reichenbach's (1947) separation between event (*E*), speech (*S*), and reference (*R*) times. To illustrate, consider examples 36 and 37, taken from Lascarides & Asher (1993):

- (36) Max slipped. He spilt a bucket of water.
- (37) Max slipped. He had spilt a bucket of water.

On Reichenbach's account, the difference between the simple past and past perfect lies in how the reference time is ordered with respect to the event time: The simple past equates them ($E =_t R <_t S$), whereas the past perfect orders the reference time before the event time ($E <_t R <_t S$). This approach works well for cases like example 37: If one equates the reference time

associated with the spilling and the event time associated with the slipping, the fact that the spilling occurred before the slipping is derived. It does not work for cases like example 36, however, as the same assumptions predict that the two events occurred simultaneously. Whereas modifications to Reichenbach's system have been offered to derive the forward progression of time (e.g., Hinrichs 1986), multiple orderings of events are actually possible: It could be that the spilling occurred after the slipping, that it occurred before the slipping (e.g., if the spilling caused the slipping), or that no ordering may be being conveyed, for instance, if example 36 was offered in answer to the question *What happened to Max today?*

To my knowledge, Lascarides & Asher (1993) were the first to investigate the relevance of coherence establishment to the inference of temporal relations. In their analysis, in fact, no tenses are treated as anaphoric; instead, temporal relations between events are recovered purely as a by-product of establishing coherence relations. The simple past and past perfect tenses differ not in their referential properties but instead in the coherence relations with which they are compatible: Whereas the simple past is compatible with any relation, the past perfect is compatible only with a limited set.

Kehler (2002) argues for a different analysis, one in which the referential properties of tense and coherence establishment interact. In this account, tense places certain grammatically encoded constraints on temporal relations, but then coherence establishment may further enrich them, so long as the constraints imposed by tense are not overridden. This account agrees with Reichenbach's in attributing a reference time, and hence anaphoric behavior, to the past perfect. However, the simple past has no reference time: It simply orders the event time before the speech time ($E <_t S$). (For a broader analysis that shares these properties, see Comrie 1985.) As such, the simple past tense plays no role in ordering the event times in example 36. As with Lascarides & Asher (1993), coherence establishment is instead responsible: Occasion and Result enforce the forward movement of time, Explanation enforces backward movement, and Parallel enforces no ordering. However, treating the past perfect as anaphoric captures the fact that only coherence relations compatible with the backward movement of time can be inferred, avoiding the need to stipulate this fact as part of axiomatic knowledge as in Lascarides and Asher's account.

Whereas I forgo further details, the key takeaway is that attempts to ascribe the full range of temporal relations that can be inferred to hold among eventualities in discourse to the semantics of tense are doomed to fail. In contrast, a fairly straightforward account of the semantics of tense is possible when the interaction between its linguistic properties and the establishment of coherence relations is taken into account.

3.5. Accent Placement

Finally, there is evidence to suggest that coherence establishment is a crucial factor in determining the distribution of focus and accent within utterances. This is notable because, on some influential theories, the rules that govern focus marking and accent placement are sensitive only to the properties of constituents and their relationships to the preceding context. An example is the analysis of Schwarzschild (1999), in which a wide range of accent patterns fall out from an optimization-driven interaction among a small set of rules and violable constraints. Schwarzschild's elegant account demonstrates how a single mechanism can simultaneously capture a variety of ways in which new information has been defined in the literature. The coverage of the system is impressive in light of the fairly limited set of theoretical tools it employs.

Kehler (2005), however, argues that any empirically adequate account will need to be sensitive to the inferred coherence relation between the target utterance and the context. Compare the following two examples from Kehler 2005 (p. 99):

- (38a) {John cited Mary, but}
 (38b) he DISSED_F SUE_F.
 (38c) (# he [dissed_F SUE_F]_F. (Schwarzschild's ex. 60)
 (39a) {Fred read the menu and then}
 (39b) # he ORDERED_F [a HAMBURGER]_F.
 (39c) he [ordered_F [a HAMBURGER]_F]_F.

(The F subscripts denote focus, and all caps indicates accent.) These examples are equivalent with respect to the information status of the three constituents in the second clause: The subject is given in the discourse, and the verb and object are new. Yet they receive different accent patterns: Both the verb and direct object of the second clause of sentence 38 receive accent, whereas in sentence 39 only the object does. The difference between them lies in their coherence construals: Example 38 participates in a Parallel relation, whereas example 39 participates in an Occasion relation.

The role of coherence becomes especially clear when comparing different coherence construals of the same passage. Consider example 40:

- (40) Powell defied Cheney, and Bush punished him.
 (Kehler 2002)

Example 40 admits of both Parallel and Result interpretations, and the accent pattern assigned depends on which one is operative. On a Parallel interpretation, the pronoun *him* does not receive accent if it is coreferent with Cheney, and it does if it is coreferent with Powell, as shown in example 41:

- (41a) BUSH_F PUNISHED_F him. (*him* = Cheney)
 (41b) BUSH_F PUNISHED_F HIM_F. (*HIM* = Powell)

In contrast, on a Result interpretation—whereby the punishment is a result of the defiance—*him* receives no accent when it is coreferent with Powell. Indeed, placing accent on *him* in this case yields an interpretation in which it refers to Cheney, signaling a contrast with an expectation about who would be punished in such a situation. These accent patterns and corresponding F-markings are shown in example 42:

- (42a) BUSH_F [PUNISHED_F him]_F. (*him* = Powell)
 (42b) BUSH_F [punished HIM_F]_F. (*HIM* = Cheney)

It is worth noting that there is nothing special about the function of accented pronouns in such examples. Indeed, the placement of accent is independent of the choice to pronominalize; the same accent patterns result if repeated proper names are used instead.

The takeaway message is that coherence construals, by virtue of the differing ways in which they structure the discourse informationally, contribute an implicit factor that influences the manner in which utterances receive accent (for further details, see Kehler 2005). Theories therefore need to go beyond working with only the information status (given, new, entailed) of the constituents in the utterance.

4. RELATIONSHIP TO OTHER MODELS OF COHERENCE

The previous section has surveyed several linguistic applications of a Relational theory of coherence. The Relational view is not the only theory of coherence on offer, however. For instance,

Grosz & Sidner (1986, *inter alia*), following work in speech act theory and plan recognition (e.g., Cohen & Perrault 1979, Allen & Perrault 1980), argue that the role of the speaker's utterance in her larger plan underlying the discourse is the determining factor of coherence. This view has come to be known as the Intentional approach. In this view, a hearer considers utterances as actions and infers the plan-based speaker intentions underlying them to establish coherence.

Historically, whereas Relational theories have been applied primarily to monologues, the Intentional approach has been applied predominantly to dialogues. The following interchange from Cohen et al. (1990, p. 1) illustrates the intuitions behind the approach:

- (43) Customer: Where are the chuck steaks you advertised for 88 cents per pound?
Butcher: How many do you want?

Typically, the most coherent way to respond to a question is to provide an answer (e.g., “behind the counter”), but here the butcher responds instead with a question of his own. But rather than being annoyed, the customer is likely to find it perfectly cooperative. Why?

Because the customer, like any typical hearer, will draw an extended chain of inferences to establish the relevance of the butcher's question. Specifically, she will presumably infer that the butcher (*a*) understands that knowing the location of the steaks is a precondition in a plan for coming to obtain steaks, (*b*) has inferred from the customer's inquiry about the location of the steaks that she is executing a plan to obtain steaks, (*c*) that this plan in turn serves a larger plan to satisfy the higher-level goal of purchasing the steaks, (*d*) realizes that the customer's plan for obtaining the steaks will not work because she cannot get to the steaks herself, (*e*) comes up with a different plan, in which the customer obtains the steaks by having him give them to her, (*f*) realizes that an unsatisfied precondition to carrying out his plan is knowing the number of steaks she wants, and (*g*) has asked her the question he did as a recipe for satisfying this precondition. And after realizing all this, our customer will presumably respond with an answer, having inferred all of the above in establishing the coherence of the dialogue without so much as a blink. The Relational approach to coherence leaves us wanting when we consider passages like example 43, since no coherence relation posited to apply to the two questions could capture this rich, hierarchical structure of plan-based inferences.

That having been said, it is far from clear how the Intentional approach can account for the coherence of even simple monologic passages like example 1, repeated below as passage 44, let alone provide an explanatory analysis of the interaction with linguistic phenomena outlined in Section 3:

- (44) I won't be able to go to dinner tonight. I have a paper to finish for the *Annual Review of Linguistics*.

For instance, it is not clear what constraints exist on plausible goals such that hearers typically are not happy with merely assuming that the speaker intends to inform them of some facts, with no further inferences being necessary. Alternatively, we could say that the hearer recognizes the speaker's intention to use the second sentence to perform the action of explaining (justifying, motivating) the eventuality denoted by the first, but it is only through the coherence-driven reasoning offered by the Relational view that the hearer is able to figure out what the relevant intention-level action was. Whereas Hobbs (1997) concedes that “the Intentional Perspective is the correct one when we look at things from the broadest possible point of view” (p. 139), he argues that, particularly in written text, “there is a level of detail that is eventually reached at which the Intentional Perspective tells us little” (p. 140).

A theory that can in some ways be seen to sit between these two approaches is the QUD analysis, proposed by Carlson (1983) and elaborated by van Kuppevelt (1995), Ginzburg (1996),

and most influentially, Roberts (2012). In QUD analyses, discourses are structured by question–answer relationships rather than coherence relations. Roughly speaking, an utterance is coherent insofar as it provides an answer to a (generally implicit) question that is relevant to the preceding discourse. So for an example like passage 44, the task is not to identify which of a set of coherence relations relates the two clauses but instead to identify a question that the second utterance could be understood to answer in relation to the first. Here, the natural implicit question to infer is *Why?*

The question that immediately arises is what type of process is responsible for determining what questions are suitably relevant, as well as for establishing coherence at higher levels of the hierarchy, such that coherent discourses can be distinguished from incoherent ones. According to Roberts’s (2012) analysis, discourses are organized around the conversational goals that interlocutors have and the “strategies of inquiry” they employ to satisfy them. Specifically, Roberts suggests, following Stalnaker (1979), that discourse is to be viewed as an attempt by conversational participants to discover and share “the ways things are,” or to put it in QUD terms, to answer the question *What is the way things are?*. Thus, by engaging in a conversation, the interlocutors agree to jointly adopt goals that center around finding the answers to this question. This in turn will generally necessitate the adoption and satisfaction of subgoals centered on answering subquestions, giving rise to a hierarchical discourse structure. Understanding a discourse therefore requires that hearers not only understand the particular utterances in the discourse, but infer the questions they are intended to answer, and further situate these questions within the underlying strategy of inquiry.

The fact that Roberts’s (2012) QUD theory is rooted in the conversational goals of the interlocutors means it shares properties of the Intentional approach of Grosz & Sidner (1986) and related works, although intentions in Roberts’s analysis appear to be more narrowly restricted to providing full or partial answers to QUDs. But as Roberts briefly notes, the Relational and QUD analyses share properties as well, in that coherence relations can often be characterized in terms of implicit questions that intervene among utterances. We have seen in our discussion of example 44, for instance, that cases that the Relational view would characterize as Explanation relations can typically be characterized in a QUD analysis as having the implicit QUD *Why?* intervening between the effect and its cause or reason. Note that the constraints on recognizing the coherence of such discourses are similar on the two analyses: The processes of recovering the implicit Explanation relation or the implicit QUD *Why?* both require that the hearer use world knowledge to infer that the second sentence in example 44 describes the reason for the eventuality described in the first. Other corresponding relationships between coherence relations and questions readily come to mind: *How come?* and *What for?* also correspond to Explanation, *What happened next?* corresponds to Occasion or Result, and *Where?*, *When?*, and *How?* correspond to Elaboration.

The Parallel relation also receives a natural treatment within the QUD analysis, and it is worth considering separately because of its configurational properties. In Section 2, I characterized utterances that participate in a Parallel relation as being related by a common topic that each utterance instantiates; in the case of example 7, repeated below as examples 45b and c, this topic was the things *X* has been doing to advertise *X*’s leadership skills. These utterances are equally naturally characterized as providing partial answers to an implicit QUD (example 45a). In this case the QUD sits above both utterances rather than intervening between them as in our treatment of example 44.

- (45a) [Who has done what to advertise their leadership skills?]
- (45b) Biden has been touting his years of experience from being a Senator and Vice President.
- (45c) Sanders has stressed the fact that he’s leading a revolution.

This analysis captures the oft-cited intuition that the Parallel relation is “multi-nuclear” in the terminology of RST, as well as why, unlike other coherence relations, it can relate more than two clauses. Finally, Roberts (2012) and Büring (2003) demonstrate that the intonational properties of such passages follow predicted patterns of question–answer congruence under this model. This having been said, it is less clear how other relations commonly found in Relational theories (e.g., Violated Expectation, Contrast, Generalization) naturally map onto QUDs, and thus a fuller understanding of the relationship between the two accounts remains an important and active question for research.

Like coherence relations, QUDs have been used as a source of explanation in linguistic theories. In fact, the phenomena addressed to date include two discussed in the previous section, particularly pronouns and ellipsis. In the case of pronouns, the experiments conducted by Rohde et al. (2007) discussed in Section 3.1 can also be seen to examine the effects of QUDs on interpretation, since the method used for influencing the distribution of coherence relations was to ask participants to complete passages in a way that answered either the question *Why?* or *What happened next?*. Kehler & Rohde (2016) subsequently showed that the same QUD manipulation affects pronoun interpretation during self-paced reading.

The linguistic phenomenon to which QUDs have perhaps been applied most extensively, however, is ellipsis. For instance, Kehler & Büring (2008) appeal to them in their treatment of Dahl-style missing readings in VP-ellipsis (Dahl 1974)—that is, to explain why the sentence *John thinks he loves his wife and Bill does too* lacks the reading in which Bill thinks John loves Bill’s wife. Keshet (2013) likewise appeals to QUDs to account for examples with sloppy readings when the pronoun in the antecedent fails to stand in a c-command relation with its binder, and Elliott et al. (2014) appeal to them in handling what they call “sticky” interpretations. Kehler (2015) uses QUDs to explain a mysterious interpretation of an example first discussed by Hardt (1992, 1999)—*Every boy in John’s class hoped she would pass him. In John’s case, I think she will.*—which appears to have a sloppy reading (*I think she will pass John*) even though it lacks the type of syntactic and semantic parallelism between the antecedent and ellipsis clauses that virtually all theories require. QUDs have also played an instrumental role in theories aimed at accounting for the conditions under which syntactically mismatched antecedents are acceptable (Kertz 2008, 2013; Grant et al. 2012), exophorically resolved cases (Miller & Pullum 2014), and cases with nominalized antecedents (Miller & Hemforth 2014). Weir (2014), interestingly, posits a QUD-based theory of fragment ellipsis, whereby certain contrasting behaviors of fragment ellipsis and VP-ellipsis (Jacobson 2016) are explained via the claim that VP-ellipsis, unlike fragment ellipsis, is not similarly constrained by the operative QUD.

QUDs have also been used in analyses of sluicing, such as *Someone murdered Smith, but I don’t know who* (Ginzburg & Sag 2000; AnderBois 2010, 2014; Ginzburg 2012; Barros 2014; Poppels & Kehler 2019). According to AnderBois (2014), for instance, such “inquisitive” linguistic phenomena—indefinites, disjunction—evoke QUDs that are capable of licensing a subsequent sluice. The acceptability of example 46a and unacceptability of example 46b are explained by the claim that the indefinite *someone* is necessary to evoke the QUD *Who was Smith murdered by?*:

- (46a) Smith was murdered by someone, but I don’t know who.
- (46b) # Smith was murdered, but I don’t know who.

Poppels & Kehler (2019) similarly appeal to QUDs and their gradient accessibility to explain the relative acceptability of cases that go beyond those captured by AnderBois’s proposal (*Regarding Trump’s impeachment, the only question is when*).

5. CONCLUSION

The primary goal of coherence theory is to provide an explanation for the coherence properties of discourse: what properties distinguish a discourse from a mere collection of utterances, and what drives comprehenders to draw inferences in service of establishing coherence. As I hope to have demonstrated, however, the importance of coherence theory goes well beyond that. Indeed, it stands to reason, in light of the inherent discourse dependence of a variety of linguistic phenomena, that their (in many cases, puzzling) behavior may result not merely from their linguistic properties but also from the interaction of those properties with the constraints and interpretive processes associated with coherence establishment. Theories of coherence establishment should therefore take their place in the linguist's toolbox as a source of explanation in linguistic theory.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

LITERATURE CITED

- Allen JF, Perrault CR. 1980. Analyzing intention in utterances. *Artif. Intell.* 15(3):143–78
- AnderBois S. 2010. Sluicing as anaphora to issues. *Proc. Semant. Linguist. Theory* 20:451–70
- AnderBois S. 2014. The semantics of sluicing: beyond truth-conditions. *Language* 90(4):887–926
- Arregui A, Clifton C, Frazier L, Moulton K. 2006. Processing elided verb phrases with flawed antecedents: the recycling hypothesis. *J. Mem. Lang.* 55(2):232–46
- Asher N, Lascarides A. 2003. *Logics of Conversation*. Cambridge, UK: Cambridge Univ. Press
- Barros M. 2014. *Sluicing and identity in ellipsis*. PhD Thesis, Rutgers Univ., New Brunswick, NJ
- Büring D. 2003. On D-trees, beans, and B-accent. *Linguist. Philos.* 26:511–45
- Carlson L. 1983. *Dialogue Games: An Approach to Discourse Analysis*. Dordrecht, Neth.: Reidel
- Cohen PR, Morgan J, Pollack ME. 1990. Introduction. In *Intentions in Communication*, ed. PR Cohen, J Morgan, ME Pollack, pp. 1–13. Cambridge, MA: Bradford/MIT Press
- Cohen PR, Perrault CR. 1979. Elements of a plan-based theory of speech acts. *Cogn. Sci.* 3:177–213
- Comrie B. 1985. *Tense*. Cambridge, UK: Cambridge Univ. Press
- Dahl Ö. 1974. *How to open a sentence: abstraction in natural language*. Log. Gramm. Rep. 12, Univ. Göteborg, Göteborg, Swed.
- Dalrymple M. 2005. Against reconstruction in ellipsis. In *Studies in Linguistics and Philosophy*, Vol. 81: *Ellipsis and Nonsentential Speech*, ed. R Elugardo, RJ Stainton, pp. 31–55. Dordrecht: Springer Netherlands
- Elliott PD, Nicolae A, Sudo Y. 2014. The sticky reading: VP ellipsis without parallel binding. In *Proceedings of the 24th Conference on Semantics and Linguistic Theory (SALT 24)*, pp. 640–55. Washington, DC: Linguist. Soc. Am.
- Frazier L, Clifton C. 2006. Ellipsis and discourse coherence. *Linguist. Philos.* 29(3):315–46
- Fukumura K, van Gompel RPG. 2010. Choosing anaphoric expressions: Do people take into account likelihood of reference? *J. Mem. Lang.* 62:52–66
- Ginzburg J. 1996. Dynamics and the semantics of dialog. In *Language, Logic and Computation*, Vol. 1, ed. J Seligman, pp. 221–37. Stanford, CA: CSLI Publ.
- Ginzburg J. 2012. *The Interactive Stance*. Oxford, UK: Oxford Univ. Press
- Ginzburg J, Sag IA. 2000. *Interrogative Investigations: The Form, Meaning, and Use of English Interrogatives*. Stanford, CA: CSLI Publ.
- Goldsmith J. 1985. A principled exception to the coordinate structure constraint. In *Papers from the Twenty-First Annual Regional Meeting of the Chicago Linguistic Society*, pp. 133–43. Chicago: Chicago Linguist. Soc.
- Grant M, Clifton C, Frazier L. 2012. The role of non-actuality implicatures in processing elided constituents. *J. Mem. Lang.* 66(1):326–43

Published version of
technical report
circulated in 1991.

- Grosz BJ, Sidner CL. 1986. Attention, intentions, and the structure of discourse. *Comput. Linguist.* 12(3):175–204
- Halliday M, Hasan R. 1976. *Cohesion in English*. London: Longman
- Hardt D. 1992. Verb phrase ellipsis and semantic identity. In *Proceedings of the 2nd Conference on Semantics and Linguistic Theory (SALT 2)*, pp. 145–61. Washington, DC: Linguist. Soc. Am.
- Hardt D. 1999. Dynamic interpretation of verb phrase ellipsis. *Linguist. Philos.* 22(2):187–221
- Hinrichs E. 1986. Temporal anaphora in discourses of English. *Linguist. Philos.* 9:63–82
- Hobbs JR. 1979. Coherence and coreference. *Cogn. Sci.* 3:67–90
- Hobbs JR. 1990. *Literature and Cognition*. Stanford, CA: CSLI Publ.
- Hobbs JR. 1997. On the relation between the informational and intentional perspectives on discourse. In *Computational and Conversational Discourse*, ed. EH Hovy, DR Scott, pp. 139–57. Berlin: Springer-Verlag
- Hobbs JR, Stickel ME, Appelt DE, Martin P. 1993. Interpretation as abduction. *Artif. Intell.* 63:69–142
- Hume D. 1955 (1748). *An Inquiry Concerning Human Understanding*. New York: Lib. Arts Press
- Jacobson P. 2016. The short answer: implications for direct compositionality (and vice versa). *Language* 92(2):331–75
- Kaiser E. 2011. On the relation between coherence relations and anaphoric demonstratives in German. In *Proceedings of Sinn und Bedeutung 15*, pp. 337–51. Saarbrücken, Ger.: Saarland Univ. Press
- Kaiser E, Cherqaoui B. 2016. Effects of coherence on anaphor resolution, and vice versa: evidence from French personal pronouns and anaphoric demonstratives. In *Empirical Perspectives on Anaphora Resolution*, ed. A Holler, K Suckow, pp. 51–78. Berlin/Boston: De Gruyter Mouton
- Kehler A. 1993. The effect of establishing coherence in ellipsis and anaphora resolution. In *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, pp. 62–69. Stroudsburg, PA: Assoc. Comput. Linguist.
- Kehler A. 2000. Coherence and the resolution of ellipsis. *Linguist. Philos.* 23(6):533–75
- Kehler A. 2002. *Coherence, Reference, and the Theory of Grammar*. Stanford, CA: CSLI Publ.
- Kehler A. 2005. Coherence-driven constraints on the placement of accent. In *Proceedings of the 15th Conference on Semantics and Linguistic Theory (SALT 15)*, pp. 98–115. Washington, DC: Linguist. Soc. Am.
- Kehler A. 2015. On QUD-based licensing of strict and sloppy ambiguities. In *Proceedings of the 25th Conference on Semantics and Linguistic Theory (SALT 25)*, pp. 512–32. Washington, DC: Linguist. Soc. Am.
- Kehler A. 2019. Ellipsis and discourse. In *The Oxford Handbook of Ellipsis*, ed. J van Craenenbroeck, T Temmerman, pp. 314–41. Oxford, UK: Oxford Univ. Press
- Kehler A, Büring D. 2008. Be bound or be disjoint! In *Proceedings of the 38th Annual Meeting of the North East Linguistic Society (NELS 38)*, Vol. 1, ed. A Schardl, M Walkow, M Abdurrahman, pp. 487–501. Amherst, MA: Grad. Linguist. Stud. Assoc.
- Kehler A, Kertz L, Rohde H, Elman JL. 2008. Coherence and coreference revisited. *J. Semant.* 25(1):1–44
- Kehler A, Rohde H. 2013. A probabilistic reconciliation of coherence-driven and centering-driven theories of pronoun interpretation. *Theor. Linguist.* 39(1–2):1–37
- Kehler A, Rohde H. 2016. Evaluating an expectation-driven QUD model of discourse interpretation. *Discourse Process.* 54:219–38
- Kehler A, Rohde H. 2019. Prominence and coherence in a Bayesian theory of pronoun interpretation. *J. Pragmat.* 154:63–78
- Kertz L. 2008. Focus structure and acceptability in verb phrase ellipsis. In *Proceedings of the 27th West Coast Conference on Formal Linguistics*, pp. 283–91. Los Angeles: Univ. Calif., Los Angeles
- Kertz L. 2013. Verb phrase ellipsis: the view from information structure. *Language* 89(3):390–428
- Kertz L, Kehler A, Elman JL. 2006. Grammatical and coherence-based factors in pronoun interpretation. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, pp. 1636–41. Austin, Tex.: Cogn. Sci. Soc.
- Keshet E. 2013. Sloppy identity unbound. In *Proceedings of the 23rd Conference on Semantics and Linguistic Theory (SALT 23)*, ed. T Snider, pp. 412–31. Washington, DC: Linguist. Soc. Am.
- Kim C, Runner J. 2009. Strict identity, coherence, and parallelism in VP ellipsis. In *Proceedings of the 19th Conference on Semantics and Linguistic Theory (SALT 19)*, pp. 275–87. Washington, DC: Linguist. Soc. Am.

- Kim C, Runner J. 2011. Discourse structure and syntactic parallelism in VP ellipsis. In *UMass Occasional Papers in Linguistics: Processing Linguistic Structure*, ed. J Harris, M Grant, pp. 75–102. Amherst, MA: Grad. Linguist. Stud. Assoc.
- Kuno S. 1976. Subject, theme, and the speaker's empathy—a reexamination of the relativization phenomena. In *Subject and Topic*, ed. CN Li, pp. 419–44. New York: Academic
- Kuno S. 1987. *Functional Syntax—Anaphora, Discourse and Empathy*. Chicago/London: Univ. Chicago Press
- Lakoff G. 1986. Frame semantic control of the coordinate structure constraint. In *CLS-22, Part 2: Papers from the Parasession on Pragmatics and Grammatical Theory at the Twenty-Second Annual Regional Meeting*, pp. 152–67. Chicago: Chicago Linguist. Soc.
- Lascarides A, Asher N. 1993. Temporal interpretation, discourse relations, and common sense entailment. *Linguist. Philos.* 16(5):437–93
- Levin N, Prince EF. 1986. Gapping and causal implicature. *Papers Linguist.* 19:351–64
- Longacre RE. 1983. *The Grammar of Discourse*. New York: Plenum
- Mann WC, Thompson SA. 1987. *Rhetorical Structure Theory: a theory of text organization*. Tech. Rep. RS-87-190, Inf. Sci. Inst., Univ. South. Calif., Marina del Rey
- Miller P, Hemforth B. 2014. *VP ellipsis beyond syntactic identity: the case of nominal antecedents*. Work. Pap., Paris Diderot Univ., Paris, Fr.
- Miller P, Pullum GK. 2014. Exophoric VP ellipsis. In *The Core and the Periphery: Data-Driven Perspectives on Syntax Inspired by Ivan A. Sag*, ed. P Hofmeister, E Norcliffe, pp. 5–32. Stanford, CA: CSLI Publ.
- Nerbonne J. 1986. Reference time and time in narration. *Linguist. Philos.* 9:83–95
- Partee B. 1984. Nominal and temporal anaphora. *Linguist. Philos.* 7:243–86
- Polanyi L. 1988. A formal model of the structure of discourse. *J. Pragmat.* 12:601–38
- Poppels T, Kehler A. 2019. *Ellipsis and the QUD: evidence from sluicing with nominal antecedents*. Paper presented at the 8th Experimental Pragmatics Conference (XPRAG 2019), Edinburgh, UK, June 19–21
- Prasad R, Dinesh N, Lee A, Miltsakaki E, Robaldo L, et al. 2008. The Penn Discourse Treebank 2.0. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC)*, pp. 2961–68. Paris: Eur. Lang. Res. Assoc.
- Reichenbach H. 1947. *Elements of Symbolic Logic*. New York: Macmillan
- Roberts C. 2012. Information structure in discourse: towards an integrated formal theory of pragmatics. *Semant. Pragmat.* 5(6):1–69**
- Rohde H. 2008. *Coherence-driven effects in sentence and discourse processing*. PhD Thesis, Univ. Calif., San Diego, La Jolla
- Rohde H, Kehler A. 2014. Grammatical and information-structural influences on pronoun production. *Lang. Cogn. Neurosci.* 29(8):912–27
- Rohde H, Kehler A, Elman JL. 2007. Pronoun interpretation as a side effect of discourse coherence. In *Proceedings of the 29th Annual Conference of the Cognitive Science Society (CogSci 2007)*, pp. 617–22. Austin, TX: Cogn. Sci. Soc.
- Rosa EC, Arnold JE. 2017. Predictability affects production: thematic roles can affect reference form selection. *J. Mem. Lang.* 94:43–60
- Ross JR. 1967. *Constraints on variables in syntax*. PhD Thesis, MIT, Cambridge, MA
- Sanders TJM, Spooren WPM, Noordman LGM. 1992. Toward a taxonomy of coherence relations. *Discourse Process.* 15:1–35
- SanPietro SA, Xiang M, Merchant J. 2012. Accounting for voice mismatch in ellipsis. In *Proceedings of the 30th West Coast Conference on Formal Linguistics*, pp. 32–42. Somerville, MA: Cascadia Proc. Proj.
- Schwarzschild R. 1999. Givenness, AvoidF, and other constraints on the placement of accent. *Nat. Lang. Semant.* 7:141–77
- Stalnaker R. 1979. Assertion. In *Syntax and Semantics 9: Pragmatics*, ed. P Cole, pp. 315–32. New York: Academic
- van Kuppevelt J. 1995. Discourse structure, topicality, and questioning. *J. Linguist.* 31:109–47
- Webber BL. 1988. Tense as discourse anaphor. *Comput. Linguist.* 14(2):61–73
- Webber BL, Prasad R, Lee A, Joshi A. 2019. *The Penn Discourse Treebank 3.0 Annotation Manual*. <https://catalog.ldc.upenn.edu/docs/LDC2019T05/PDTB3-Annotation-Manual.pdf>
- Webber BL, Stone M, Joshi A, Knott A. 2003. Anaphora and discourse structure. *Comput. Linguist.* 29(4):545–

Published version of
draft circulated in 1996
and amended in 1998.

- Weir A. 2014. Fragment answers and the question under discussion. In *Proceedings of the 38th Annual Meeting of the North East Linguistics Society (NELS 38)*, ed. J Iyer, L Kusmer, pp. 255–66. Amherst, MA: Grad. Linguist. Stud. Assoc.
- Winograd T. 1972. *Understanding Natural Language*. New York: Academic
- Wolf F, Gibson E. 2006. *Coherence in Natural Language: Data Structures and Applications*. Cambridge, MA: MIT Press
- Zhan M, Levy R, Kehler A. 2020. Pronoun interpretation in Mandarin Chinese follows principles of Bayesian inference. *PLOS ONE* 15(8):e0237012