

On Convention and Coherence*

Andrew Kehler[†] and Jonathan Cohen[‡]

1 The controversy

The idea that literal linguistic meaning underdetermines communicated content is central for contemporary theorizing about language. This broad idea is more or less universally accepted, and has proven immensely fruitful for research in semantics, pragmatics, and the interface between the two. However, the broad idea leaves open the hotly disputed question of just how to draw the partition that separates linguistic meaning from the rest of communicated content.

In one of the most influential articulations of the broad idea — and one that kicked off an industry of explanation in semantics and pragmatics, [Grice \(1975\)](#) argues for what, in hindsight, amounts to an extreme position. He proposes that meanings of utterances remain quite close to their explicit linguistic form, differing from the latter only in allowing for resolution of indexicals, tense, and ambiguity. He holds that everything else communicated by utterances falls outside of literal linguistic meaning — that it is supplied by a particular form of pragmatic enrichment (“implicature”), the explanation for which lies in general properties of our cognitive apparatus, such as the rationality of conversational participants and the presumption of cooperativity as they pursue their conversation as a joint collaborative activity. In broad strokes, then, the Gricean program holds that linguistic meaning hews closely to what is specified by overt linguistic form, and that the large residual gap between linguistic meaning and communicated content is filled in by implicature.

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

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

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

Subsequent researchers have offered a range of refinements to the program, often reeling in Grice by pointing out that not all forms of enrichment satisfy the hallmarks of implicature, and that the Gricean schema for working out implicatures does not readily apply to all cases (e.g. [Bach 1994](#); [Sperber and Wilson 1986](#)). Still, as far as we can tell, none of these critics have disavowed the central concept of implicature or its role in bridging the gap between linguistic meaning and communicated content.

Lepore and Stone (henceforth, L&S) propose to depart from this consensus in ways that ultimately amount to dismantling the Gricean program almost entirely. Though they nominally agree with Griceans that the meaning of an utterance ("what is said") lies close to what is overtly encoded in its linguistic expression — again, allowing only for the fixing of indexicals and other context-sensitive expressions, tense, and ambiguity resolution, they give a far larger role to ambiguity resolution than Griceans allow. Indeed, L&S argue that nearly everything Griceans characterize as going beyond the recovery of literal meaning is better understood as disambiguation on the way to fixation of literal meaning ("where alternative approaches have postulated pragmatic processes of enrichment, what's really going on is disambiguation: finding the right reading of the utterance, understood as a [conventionally,] grammatically specified pairing of form and meaning" p. 88). Since, on this view, nearly all conveyed content is the product of disambiguation processes, there is little or no role left for any additional processes of extrasemantic enrichment or expansion by Gricean or other tools. This leads L&S to deny a (perhaps *the*) central tenet of pragmatic theory — that speakers construct their utterances to take advantage of their addressees' extralinguistic cognitive apparatus to communicate more content than they can be said to linguistically express.

Seen from afar, Grice's and L&S's views amount to two opposite but equally extreme positions about the relationship between literal meaning and communicated content. On the one pole, Grice holds that meaning encodes much less than what is communicated, and that, consequently, implicature must play a large role in bridging the chasm between the two. On the other, L&S hold that suitably disambiguated, conventionally specified linguistic literal meaning encodes (nearly?) all of what is communicated, and that, consequently, there's very little for implicature or other forms of pragmatic enrichment to do.

In what follows we argue against both extremal positions, and instead contend that the truth lies somewhere in the middle. On the one side, we believe L&S have performed an important service to the field by showing how Grice's program assigns too large a role to implicature — that it

“underestimated the scope of linguistic conventions and overestimated the reaches of communicative intentions” (Szabó 2016, p. 169; cf. Horn 2016, p. 151). Indeed, we’ll give additional arguments against this sort of Gricean overreach in §2. However, on the other side, in §§3–6, we’ll also deny that convention and disambiguation can do all the work that L&S assign to them, and therefore will contend that their ambiguity resolution view fails as well. In §7 we’ll argue for a middle ground account that recognizes robust contributions from both convention and pragmatic enrichment (and their interaction), using the interpretation of tense and event structure as a testing ground for theorizing. In §8, we’ll conclude.

2 On Gricean overreach: Coherence-driven pragmatic enrichments

L&S object to Gricean explanations by appealing primarily to three phenomena: indirect speech acts, English intonation and its conventional marking of information structure, and enrichments that result from the establishment of discourse coherence. Although we are not prepared to endorse L&S’s claim that there is no role for implicature in a theory of language meaning, we agree that these are cases in which Gricean analysts have overplayed their hand.¹ To keep our contribution focused, however, we will concentrate our discussion on the last of these phenomena, which we term COHERENCE-DRIVEN ENRICHMENTS. In this section we’ll argue, alongside L&S, that such enrichments are not the result of implicature.

We begin with the oft-made observation that, when confronted with a set of co-occurring utterances in a discourse, comprehenders are not content to merely update their world models with the meanings of each utterance (Asher and Lascarides 2003; Hobbs 1990; Kehler 2002, inter alia). Instead, they seek to establish some sort of semantic relationship — a COHERENCE RELATION — between them. For example, suppose a faculty

¹There are other cases as well. Levinson (2000, pp. 117–118), for instance, offers a diverse list of phenomena that he argues involve implicature, including noun-noun compound interpretation (*The oil compressor gauge*), possessive interpretation (*Wendy’s children*), bridging (*The picnic ... the beer*), preferred local coreference (*John came in and he sat down.*), inference to stereotypes (*secretary* \Rightarrow *female secretary*), negative strengthening (*I don’t like Alice* \Rightarrow *I positively dislike Alice*), conditional strengthening (*if* \Rightarrow *if and only if*) and conjunction buttressing (*and* \Rightarrow *and as a result*). Although of these we will only discuss conjunction buttressing, we are disinclined to view any of these phenomena as involving implicature, and are confident that L&S would agree.

member (let's call him Andy) uttered (1) to his colleague (let's call him Jonathan) one afternoon:

- (1) We should stay on campus and work this evening. Our paper for the Lepore & Stone volume is overdue.

Andy's utterances are likely to get Jonathan pondering the question of whether he is able work with Andy on the paper that night, and he would be quite surprised (and perhaps angry) if, after agreeing to rearrange his schedule, it turned out that Andy's only intention was to sink a couple beers at the campus bar. Now, clearly Andy never said that he wanted to work on the paper with Jonathan that evening. It was merely an inference that Jonathan drew.

On the envisaged scenario, Jonathan has a right to be angry nonetheless. For his attempt to bring coherence to (1) will undoubtedly lead him to infer an EXPLANATION relation, whereby the second sentence is understood to describe the cause or reason for the event described in the first. Importantly, the additional inference involved in such examples goes beyond merely assuming that some causal relationship exists. Instead, the relationship needs to be established using the context and the interlocutors' shared knowledge and capacity for inference. A natural assumption to make when interpreting (1), for instance, is that Andy intends to work on the paper that evening; with this assumption one can see how the late paper could provide a reason for wanting to stay late. Note that if this assumption conflicted with the hearer's knowledge, he would be well within his rights to question it, e.g., with a response such as *I thought your opinion was that our L&S paper is pretty much done*. Our hearer in this case isn't countering anything that is entailed by (1), but instead a pragmatic enrichment that must be drawn to establish its coherence on a natural and salient interpretation.

As L&S note, it has been commonly assumed that such inferences are the result of implicature. Grice himself, for instance, famously argued that implicature explains how the conjunction *and* can be treated as having a single meaning (i.e., that of logical conjunction) even though it has the appearance of being associated with other, more specific meanings (Grice:81).² He considers the minimal pair in (2).

- (2) a. He took off his trousers and went to bed.

²Grice's paper appeared in 1981, but an earlier version of it had circulated since 1970.

- b. He went to bed and took off his trousers.

In typical contexts, the events described in each of examples (2a–b) will be understood to have occurred in the order in which they are presented; hence they receive different construals. Noting the lack of theoretical parsimony associated with multiplying meanings of *and* (here, to include *and then*), he suggests instead that the effect results from compliance with the Maxim of Manner, specifically the submaxim imploring the speaker to “be orderly”. Hence, on this analysis, the ordering of events is an extragrammatical enrichment, which for him means an implicature.

However, this analysis runs into immediate problems. To see this, note that (3a) is easily read as conveying nothing about the relative order between the two events.

- (3) a. (For Sue’s baby shower,) Mary bought a stroller and Nancy crocheted a baby blanket.
b. Paul went to the liquor store. He ran out of scotch.

Likewise, on a natural construal of (3b), the depletion of Paul’s scotch occurs before the trip to the liquor store. The obvious question for the Gricean analysis is why, in light of the fact that it is possible to construe (3a–b) as describing the events in temporal order of occurrence, such enrichments are not drawn.

Luckily, there is a satisfying explanation of these and related facts in terms of coherence and coherence-driven enrichment. Specifically, the idea is that (3a) is an instance of a PARALLEL coherence relation, in which the utterances each instantiate a common, more general theme, or put another way, provide a partial answer to a common question under discussion (Roberts 1998). The fact that no temporal ordering is conveyed results from the fact that the Parallel relation doesn’t require one: the two clauses in (3a) can provide answers to the question *Who did what for Sue’s baby shower?* without providing a temporal order between the events. Similarly, (3b) is a canonical case of an Explanation relation as defined above. Here the hearer will infer that the depletion of scotch is the reason for, and hence occurred prior to, the trip to the store.³ Finally, Grice’s examples (2a–b)

³Of relevance to the Gricean account is the fact that the speaker could have followed the Maxim of Manner by employing a different coherence relation – RESULT – which is like Explanation but in which the cause is described before the effect (*Paul ran out of scotch and went to the liquor store*). Counter to what we take to be a prediction of the Gricean analysis, the speaker’s decision to put the effect before the cause in (3b), and hence be *unorderly*, carries no hint of uncooperativity.

are each most naturally construed to participate in an OCCASION relation, in which the events are understood to be spatio-temporally contiguous. The constraints associated with Occasion thus yield not only a forward sequence of events but other enrichments as well: for instance, in the case of (2b), that the man being referred to took off his trousers *while in bed*.

Subsequent (neo-)Griceans have been similarly quick to treat other cases of coherence-driven enrichment as implicatures. For example, Levinson (2000), who labels cases such as (2) as examples of “conjunction but-tressing”, advocates a three-way division among Q-, M-, and I-Principles. He classifies these cases as resulting from inferences that arise from his I-Principle, which comes in two parts: a Speaker’s maxim of minimization (“Say as little as necessary; that is, produce the minimal linguistic information sufficient to achieve your communicational ends (bearing Q in mind)”), and a Recipient’s corollary termed the Enrichment Rule (“Amplify the information content of the speaker’s utterance, by finding the most *specific* interpretation, up to what you judge to be the speaker’s m-intended point”), the latter which includes a subcorollary instructing hearers to “Assume the richest temporal, casual and referential connection between described situation or events, consistent with what is taken for granted.” On this analysis, the speaker is justified to conjoin the clauses with only *and*, leaving it to the hearer to enrich its meaning to the causal one.

But such Gricean explanations are problematic for several reasons.

A first is that Levinson’s phrase “consistent with what is taken for granted” is certainly too narrow to apply to the full range of cases. For example, consider (4), for which a cause-effect reading is not only available but natural.

(4) Fred got bitten by a mosquito, and contracted the Zika virus.

Once again, the availability of this interpretation is captured in terms of coherence: world knowledge that tells us that the fact that mosquito bites *can* cause Zika is enough for (4) to be taken as expressing a Result relation between the clauses. In contrast, we take it the proposed Gricean account would predict this reading of (4) only if it is presumed to be typical (indeed, taken for granted) that people who get bitten by mosquitoes contract Zika. But the interpretation at issue is easily evoked even for interlocutors who know that an exceedingly small portion of bite victims contract the disease.

Second, Levinson’s picture, as we understand it, predicts that unnecessary prolixity should trigger M(anner)-implicature, in the way that, say, *Mary got the machine to stop* implicates that the stoppage was achieved

through atypical means – an inference triggered from the speaker’s avoidance of the the less prolix *Mary stopped the machine*. However, we find that including a more specific connective in cases like (4) — *Fred got bitten by a mosquito, and as a result, contracted the Zika virus* — does not trigger a similar Gricean inference to atypical or otherwise indirect causation. Indeed, and contrary to the predictions of the Gricean view, here a speaker’s inclusion of a more prolix alternative form does not come across as uncooperative, even if she could have conveyed the same meaning without it.

Finally, as Levinson himself (pp. 122–127) and others have noted, the same enrichments we see in (4) regularly occur without the conjunction:

(5) Fred got bitten by a mosquito. He contracted the Zika virus.

Obviously there cannot be “conjunction buttressing” if there is no conjunction to buttress. Surely we want our story for the operative enrichments in (4) to apply to (5) as well, but to say that in describing the two events in (5) the speaker has somehow *implicated* a causal relationship seems to us to strain credulity. What is true is that *and* serves a function relevant to coherence establishment, in that it is only compatible with certain coherence relations: it disallows Explanation, for instance, as sentence (6a) cannot typically be interpreted to mean what (6b) expresses.⁴

(6) a. Fred slipped and he stepped on a banana peel.

b. Fred slipped because he stepped on a banana peel.

But that clearly points to conventional properties of the meaning of *and*. So to speak of drawing a causal relation between clauses as an enrichment of a conjunction meaning gets things back to front: conjunction meanings influence coherence establishment, not the other way around.

3 Interpretation as ambiguity resolution

For these reasons (among others) we are extremely sympathetic to L&S’s charge that the Gricean program overreaches in its attempt to explain such interpretive phenomena (and many others) in terms of implicature. But why do L&S go on to hold that these and other cases that are standardly

⁴The fact that the meaning of (6a) cannot be enriched to that of (6b) is interesting in light of the neo-Gricean account that allows strengthening to stereotypical interpretations. Why can such enrichment not occur here? The relationship seems perfectly stereotypical.

understood as extrasemantic enrichment are better understood as disambiguation between conventionally/grammatically specified alternatives?

Their principal support for this claim comes from a single, high-level argument form turning on Grice's criterion of detachability. As an example, they discuss speech acts such as the use of (7), in response to a waiter's question "what would you like to order?", to make an indirect request:

(7) Can I have the French Toast? (Lepore and Stone 2015, p. 92, ex. 25)

Their case that the capacity to understand (7) as a way of making polite requests must be understood as a consequence of its linguistic form comes from two central claims: (i) that interpreters organize discourses by construing their constituent sentences as standing in relations of coherence (here they cite Asher and Lascarides 2003; Hobbs 1990; Kehler 2002); and (ii) that linguistic competence requires knowing both what coherence relations are available and which of them are associated with which linguistic forms.

We hope it is clear that the combination of (i) and (ii) would indeed license L&S's radically deflationary attitude toward extrasemantic enrichment. Moreover, and as we have discussed above, we agree with L&S's view that much linguistic understanding is mediated by the establishment and recognition of coherence relations, and so are prepared to concede (i) happily. But why should we believe (ii)?

L&S's case for (ii) rests on intra- and inter-linguistic applications of Mill's Methods: they show that coherence and conversational role can vary with shifts in conventions about linguistic form even when rationality and truth-conditional content remain fixed, and so conclude that coherence and conversational role are determined by the former rather than the latter. Thus, on the intralinguistic side, they reason that if the representation of coherence were a result of general, rational processes not specifically and conventionally tied to particular linguistic forms, then replacing an expression with a paraphrase or clause reordering that preserves truth-conditional content but not form should not significantly alter conversational role. But this prediction fails for examples like (8), which, though it is a near content-match with (7), is hard to hear as an indirect request:

(8) Am I able to have the French Toast? (Lepore and Stone 2015, p. 101, ex. 96)

Analogously, on the interlinguistic side, they suggest that if coherence representation were mediated only by language-independent rational considerations, then one should expect coherence to operate in the same

way on truth-conditionally-equivalent sentences in different languages (assuming equally rational conversants). Again, they observe that this prediction fails for pairs like (10a–b): where (10a) can be used by speakers of English to offer a beer to a hearer, its Polish translation, (10b), cannot be so used by speakers of Polish (Wierzbicka 1985).⁵

- (10) a. Would you like a beer? (Lepore and Stone 2015, p. 102, ex. 97)
b. Miałbys ochotę na pivo? (Lepore and Stone 2015, p. 102, ex. 98)

L&S make the same case for the sort of coherence-driven enrichments described in §2. Specifically, they offer the contrast between (11a), in which the clauses are in the simple past tense, and the variant in (11b), in which the clauses are in the present perfect.

- (11) a. Oil prices doubled and demand for consumer goods plunged.
b. Oil prices have doubled and demand for consumer goods has plunged. (Lepore and Stone 2015, p. 117, ex. 129)

In typical contexts, passage (11a) will be construed to express a Result relation, according to which the doubling caused the plunging.⁶ In contrast, however, a Result interpretation is far less inevitable for (11b); indeed the preferred interpretation appears to be one in which a Parallel relation is operative. L&S observe correctly that this pattern of facts is problematic for the Gricean analysis, since the two versions seem to be truth-conditionally equivalent: as long as both events occurred before the speech time, both (11a-b) will be true. That we get different construals

⁵This isn't L&S's only argument for (ii). They also argue (100–102) that the indirect speech act of requesting is grammatically licensed by noting (following Horn 1989; Lakoff 1973; Sadock 1974) that (presumably by convention) *please* felicitously combines with marked and unmarked requests (cf. (9a–b)), and that it can be felicitously appended to (7), as in (9c):

- (9) a. # I'm thirsty please. (Lepore and Stone 2015, p. 101, ex. 94–96)
b. I'd like a drink please.
c. Can I have the French Toast please?

For reasons of space, we'll ignore this apparently subsidiary consideration for L&S in favor of responding to the argument on which they place the most weight.

⁶L&S actually classify it as a weaker Narrative (i.e., Occasion) relation, according to which the events are related by contingency rather than causality. The boundary between the two relations can be a fine one. Since both impose a requirement for forward movement of time, this difference will not concern us here.

thus violates the non-detachability criterion associated with conversational implicature. We find this criticism persuasive, and we take it as another reason to doubt the scope and adequacy of Gricean explanations.

As there is no middle ground between implicature and ambiguity resolution on their view, the arguments against treating such enrichments as the result of implicature entails that coherence establishment is simply an ambiguity resolution problem (see their §6.2), one for which speakers and hearers bring a variety of conventional cues to bear in negotiating discourse construals. Support for that view is provided by the fact that, as dictated by coherence theory, the process of establishing coherence is mandated: hearers have to infer *some* type of relevancy relation between adjacent clauses within a discourse segment as part of discourse comprehension. So as long as the inventory of coherence relations is finite, the process can be justifiably viewed as a problem of disambiguating among the possible relations. Indeed, applying this process iteratively to larger discourse segments yields a mechanism for discourse parsing that is in some ways analogous to sentence parsing, the latter of which is broadly agreed to be a disambiguation process.

4 Conversational eliciture

Whether or not one finds the ambiguity view of intersentential coherence establishment compelling, we believe that there is a highly related class of cases – ones for which comprehension recruits the same interpretative machinery – that is nonetheless much less naturally cast in this way. These cases, which we have labeled ELICITURES (Cohen and Kehler 2016), are those in which coherence establishment processes apply optionally within clauses, thereby generating pragmatic enrichments that are not linguistically mandated. We can see a simple example of eliciture in (12a), which strongly invites us to infer, but does not entail, that the speaker not only intends to communicate that the company fired the employee *and* the employee was embezzling money, but that the company fired the employee *because* the employee was embezzling money.

- (12) a. The company fired *the manager who was embezzling money*. (Rohde et al. 2011)
- b. The company fired *the manager who was hired in 2002*.

Note that this is merely a defeasible inference: (12a) could be followed with *The reason the manager was fired was because he was rude and always*

late. In (12b), on the other hand, being hired in 2002 will normally not be understood to be the cause of the firing; here the relative clause (RC) is merely identificational.

According to the analysis we offer in Cohen and Kehler (2016), elicitures result from the speaker's exploitation, by way of her particular choices of referring expressions, of the same types of cognitive machinery that hearers use to establish discourse coherence between clauses.⁷ That is, the relationship inferred for (12a) is unmistakably parallel to that which underlies the establishment of the Explanation relation for the intersentential variant in (13):

(13) The company fired the manager. He was embezzling money.

Whereas intrasentential cases like (12a) and intersentential ones like (13) differ in that only the latter case mandates that some coherence relation be established between the relevant propositions, we see no reason to think that the inference process itself – including the world knowledge that the hearer brings to bear – is any different.

While we are confident that L&S would agree that implicature is not the source of the causal inference in example (12a), we are likewise confident that they would disagree about what the source is: we still consider it to result from a form of pragmatic enrichment, whereas they are forced to the view that it is a result of ambiguity resolution. How would an ambiguity story go for such cases?

We assume that the ambiguity view would start at the same place as the eliciture view, with the observation that the RCs in (12a-b) each participate in a standard modification relationship with the NP to which they attach (in these cases, restricting the domain of reference for that NP), as produced straightforwardly by a standard compositional semantics. On the eliciture view, of course, that is all there is to say on the semantics side. For L&S to capture the possibility of additional content resulting from establishing coherence, we suspect that they would follow the approach outlined at the end of §3, specifically by articulating a set of coherence relations that can hold. On this story, in addition to deriving the standard meaning of an RC based on compositional principles, a search would be triggered to find an appropriate coherence relation between two propositions: one expressed by the matrix sentence (*The company fired the manager*), and one derived from the relative clause and the NP to which it attaches (*The manager was*

⁷And for that matter, the same cognitive machinery they use to understand non-linguistic situations that they encounter in the world. See §5.

embezzling money). The same world knowledge and inference process that the eliciture view relies on for enrichment is then used to disambiguate the operative coherence relation, which in the case of (12a) would be Explanation.

When we consider a larger set of cases, however, we see this picture grow increasingly complicated and, in our view, completely untenable. Consider first cases like (12b) for which, like the great many sentences containing RCs, there is no eliciture. To account for such examples, the existence of a relation of “No-Relation” would have to be posited, one which just happens to be the overwhelmingly typical case. (Note that this move creates a discontinuity with the theory of intersentential coherence, for which the lack of a No-Relation relation is crucial for accounting for infelicitous discourses.) So at this stage, we have a disambiguation problem for every RC one encounters between two relations, the most common of which is the lack of a relation.

The complexity multiplies when we discover that elicitures are not limited to relationships between the proposition derived from an RC and the one denoted by the matrix. For instance, they can relate a proposition derived from an adjectival and the one denoted by the matrix sentence. Consider (14a-b):

- (14) a. *The drug-addled undergrad* fell off of the Torrey Pines cliffs. (adapted from an example of Webber (1991))
- b. *The well-liked undergrad* fell off of the Torrey Pines cliffs.

Sentences (14a-b) exhibit the same distinction between Explanation and No-Relation that (12a-b) does. Specifically, (14a) sees the inference of a relationship between the proposition denoted by the matrix (*the undergrad fell off of the Torrey Pines cliffs*) and one derived from an adjectival and the nominal that it modifies (*the undergrad was drug-addled*). So if the ambiguity treatment of (12a-b) requires positing a two-way ambiguity for the relationship between RCs and matrix sentences, the same must be done for adjectivals, where again No-Relation will by far be the most common outcome.

Extending this inquiry in the obvious way, it doesn't take long to see that elicitures can relate propositions that are derived from any two constituents, and hence need not even involve the proposition denoted by the matrix. Consider (15a):

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

Example (15a) invites the eliciture that the drinking led to the accident. Here the eliciture results from establishing 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. Thus, to maintain an ambiguity view of elicitures, one would have to posit a coherence relation between every pair of constituents from which propositions can be derived in any given sentence, triggering a disambiguation process for each, one that will again result in No-Relation in a large majority of cases.

That there would be an explicit search for such relations, with the requisite invocation of machinery for disambiguation on such a broad scale, seems highly implausible to us. But it actually gets worse. Consider (16):

(16) The drunk pilot was arrested.

Example (16) gives rise to a rich picture of a pilot who was arrested because he was flying (or perhaps preparing to fly) while inebriated. The enrichments that give rise to this picture result despite the fact that (16) could be used to describe a situation in which a pilot, who happened to be drinking a fair bit on his day off, got arrested for cheating on his taxes. On the pragmatic enrichment account, the inferences result from world knowledge that tells us that a pilot – crucially, when flying or preparing to do so – can be arrested for being inebriated; the speaker who utters example (16) therefore takes advantage of this knowledge being in the common ground to convey her message in a particularly efficient way. The ambiguity account, on the other hand, now needs to disambiguate a relation that involves propositions derived from *three* constituents: one derived from the combination of the adjectival and the nominal it modifies (*the pilot was drunk*), one derived from the nominal itself (*the pilot was flying or preparing to*), and one derived from the matrix (*the pilot was arrested*). Examples (17a-c) reveal that variants of (16) which lack any one of these three propositions do not give rise to the same eliciture:

- (17) a. The drunk pilot was hitting on a stewardess at the hotel bar.
- b. The drunk person was arrested.
- c. The tall pilot was arrested.

It is difficult to see how examples like (16) could receive a compelling treatment on a view in which possible interpretations are conventionally prespecified and disambiguated amongst. Not only would (16) require a search for a coherence relation involving three propositions, but one in which these propositions are utilized in a particular configuration to form a two-place relation: here it is the *conjunction* of two propositions (*the pilot was drunk* and *the pilot was flying or preparing to*) that play the role of the cause, with the third proposition (*the pilot was arrested*) providing the effect.⁸ Hence, the ambiguity view must now allow for the conventionally-provided possibility of a two-place Explanation relation that combines propositions derived from multiple constituents to fill one of its argument positions.

We could continue to pursue cases of greater complexity, but hope that it has become clear that the ambiguity view is highly implausible as an account of eliciture, and at the same time fails to offer any explanatory advantages over the pragmatic enrichment view. It seems clear to us that such enrichments do not result from a ‘search’ for an interpretation, but instead are triggered by associations that our cognitive apparatus is built to recognize automatically; ones that are served up linguistically by virtue of the particular expressions that a speaker chooses to employ. That is, the only plausible trigger for these inferences is machinery that we have that is capable of recognizing such associations; machinery that is already running as we interpret not only language but indeed the world around us.

5 Further costs of the ambiguity view

Having offered an initial argument against L&S’s ambiguity resolution view based on the properties of eliciture, we now want to argue that there are number of further theoretical costs that make the view untenable. Specifically, we find that the view gives rise to an unwelcome proliferation of ambiguities, is badly underconstrained, non-compositional, that it threatens to make the semantics of natural languages unlearnable, and that it rests on a surprising coincidence between the content of linguistic

⁸Note that this situation is distinct from one in which multiple constituents give rise to multiple elicitures:

(18) The drunk kid who got into a car accident is in the hospital.

Here it is natural to infer both that the drinking led to the accident, and that the accident in turn led to the hospital stay. The situation with (16) is different in that there is only one Explanation relation being established.

conventions and the standard menu of coherence relations. We'll take these points in turn.

To begin, if eliciture-amenable sentences are ambiguous between interpretations in which the relevant elicitures are drawn and interpretations in which they are not, then there will be turn out to be *a lot* more ambiguity in natural languages than one might have expected. After all, as [Hobbs \(1979\)](#) points out, it is always possible to build contexts supporting coherence relations between otherwise apparently unrelated discourse elements: even those in a paradigm incoherent discourse like (19) can be brought into coherence, e.g., on the supposition that the employee injured himself in an unsuccessful attempt to climb a plum tree with the goal of attaining plums.

(19) # The employee broke his leg. He likes plums. (variant of example from [Knott and Dale \(1994\)](#))

Similarly, though one might not initially expect a causal eliciture between the event of firing and the RC in the object NP in (20), the eliciture is naturally evoked in a context in which John is a protective parent with a Bieber-fan daughter who also works for his company.

(20) John fired the employee who looks like Justin Bieber.

The point here is simple: if coherence can be established this freely for more or less arbitrary discourses by manipulating the non-linguistic background, and if, per L&S, elicitures are diagnostic of linguistic ambiguity, then it would seem that more or less every discourse in any natural language will turn out to be (very many-ways) ambiguous (cf. [Szabó 2016](#), p. 168).

Additionally, the ambiguity resolution view is seriously underconstrained in its current form. If the view is to avoid the charge of post-hoc storytelling that L&S (appropriately) level at Griceans, we need a detailed account of just which particular elements of linguistic form allow for which ranges of specific disambiguations, and a description of just how disambiguation derives its output from context, world knowledge, and linguistic knowledge. It won't suffice merely to say that linguistic forms turn out to be ambiguous just when, and in just the ways in which, elicitures or other phenomena classically treated as pragmatic enrichment are drawn in light of context and world knowledge. Until the account is provided in far greater detail than L&S have given so far, it's hard to know what the view predicts about cases, and consequently difficult to evaluate that view empirically.

In the case of eliciture, this concern is reinforced by the observation that, in many examples, the enriched content cannot be seen to originate in any single constituent in linguistic form, but only from the combined occurrence of multiple constituents not directly related in the syntax. To see this point, consider (21a–c):

- (21) a. The teacher met with *the student who looks like Harry Potter*.
- b. The Hollywood studio representative met with *the student who goes to the local community college*.
- c. The Hollywood studio representative met with *the student who looks like Harry Potter*.

We take it that the causal/explanatory eliciture here is much more strongly evoked by (21c) than by either (21a) or (21b): it is (21c), rather than (21a) or (21b), that evokes an image of the representative looking for the next child movie star. Crucially, the eliciture in question cannot be tied to the occurrence of any one constituent: example (21a) fails to evoke the inference even though it features the same RC as (21c), and (21b) fails to evoke the inference even though it features the same subject noun phrase as (21c). The eliciture arises only when both constituents occur together, as in (21c).

It is a consequence of this observation that, at least in many cases, L&S's ambiguities must be choices of whole propositional meanings for complete sentential linguistic forms rather than smaller units. But this means that the information L&S envisage treating as conventionally encoded pairings of form and meaning won't, in general, be compositionally determined from the meanings assigned to subsentential components and syntax. In effect, the conventional pairings at issue will be specifiable only as a look-up table (with the full power of a Turing Machine) connecting whole sentential forms with whole propositional meanings. (This point reinforces our concern that the view will be underconstrained, and therefore open to charges of post-hocery.)

And, indeed, this point encourages the suspicion that the ambiguity resolution view places severe burdens on language learning. The view requires that linguistic knowledge specify, as possible disambiguations, all the elicitures that could be possibly drawn from sentences/discourses. But, as we have seen, one can nearly always evoke elicitures from a given discourse passage, provided the context is suitable. This suggests that, for L&S, linguistic knowledge will have to predict the full observed variation

in inferential behavior of discourses in an unpredictable range of contexts. Moreover, and as we have also seen, elicitors of sentences are, in general, not compositional: they are not predictable from the interpretations of sub-sentential constituents and their syntactic configurations. Consequently, linguistic knowledge will have to specify the range of possible effects of each sentence as a separate, unstructured list. If we assume an infinitude of sentences in each natural language, then it becomes difficult to see how finite learners could attain this sort of linguistic knowledge, as required by the ambiguity resolution view.

Finally, the ambiguity resolution view is committed to a surprising coincidence between the content of linguistic convention and the standard coherence relations. We can bring this point into relief by contrasting what the ambiguity resolution theorist will say about linguistic examples like those we have already seen and structurally analogous non-linguistic cases. On the non-linguistic side, consider the following two contrasting situations. In the first, someone sees a chronically tardy employee show up late for work again, and soon thereafter witnesses the employee being fired; in the second, she sees a chronically tardy employee show up late for work again, and soon thereafter witnesses the employee being asked by a customer where the automotive department is located. A reasonable cognitive agent might infer that the firing was due to the lateness in the first situation, but is unlikely to infer any relationship between the customer's question and the employee's lateness in the second. Presumably this is because world knowledge supports the possibility of a causal/explanatory connection in the first situation but not the second. With this pair in mind, we can now consider the interpretation of linguistically expressed reports of the very same situations, such as (22a–b):

- (22) a. The boss fired the employee who came in late again.
b. A customer asked the employee who came in late again where the automotive department is.

Just as in the non-linguistic cases, a reasonable agent is likely to draw a causal elicitor in interpreting (22a), but not in interpreting (22b). Considering these cases together, it's hard to avoid the conclusion that the non-linguistic pair and the linguistic pair are related in analogous ways, and that we should aim for a theory that treats both pairs in terms of a common species of cognitive machinery.⁹

⁹This is a species of connection that Grice (1975, p. 28) himself famously emphasizes. Perhaps needless to say, it is also a species of connection that comes for free on a theory that

We presume an ambiguity resolution theorist will agree, and will attempt to capture the apparent structural similarity on display by saying that the very same inferential strategies enlisted in the non-linguistic cases are deployed in the linguistic cases in the service of disambiguation. Specifically, she will claim that linguistic convention dictates that a relative clause such as that in (22a–b) (or adjectival, or what have you) can express a cause or not, and that whatever inferential procedures license or fail to license a causal interpretation in the non-linguistic cases are in linguistic cases like (22a–b) used to choose between the available disambiguations made available.

But now we should ask: exactly why is it that convention makes available an Explanation-involving disambiguation, in particular, whenever there is a relative clause? Why *this* specific interpretive alternative for *this* specific linguistic form? As L&S correctly emphasize, linguistic and other conventions are deeply contingent: hence, linguistic convention might have paired with RCs either no coherence relation or entirely different coherence relations as potential disambiguations. Why, then, does linguistic convention make the particular pairings it does? The ambiguity resolution theorist cannot answer this question by appeal to the inferential machinery that both sides think are at work, and that she understands as serving disambiguation. After all, her view is that that that inferential machinery enters the interpretive process only after convention has already made available the relevant interpretive alternatives. And, of course, the conventions at issue are language-specific, hence, not shared with or explicable in terms of other aspects of our mental processing.

Moreover, the coincidence runs deeper than just this one particular form-content pairing of RCs with the Explanation relation. Indeed, even this single linguistic form is associated with a range of distinct elicitures. Thus, (23a) plausibly invites an interpretation that denies, rather than affirms, a causal relation between the matrix verb *fired* and the RC; (23b) plausibly invites an interpretation involving an Occasion relation, where the subject is understood to have bought the scotch at the store mentioned in the RC; and, as we have seen, vastly most RCs invite no coherence-driven interpretation (or, if you like, invite interpretation in terms of the No-Relation relation).

- (23) a. The boss fired the employee who had won many corporate rewards.

treats elicitures as extrasemantic expansions driven by the very inferential mechanisms we use in understanding the non-linguistically presented world.

- b. The employee who went to the liquor store bought a bottle of scotch.

It would seem, then, that the ambiguity resolution theorist is committed to holding that linguistic convention — which she invokes to explain linguistic but not non-linguistic cases, despite the striking similarities between the two — just happens to make available as possible disambiguations for RCs the same sorts of standard coherence relations (Explanation, Denial, Occasion, No-Relation) that come for free on an enrichment account. That is, to our minds, a surprising coincidence — and one that, as far as we can tell, the ambiguity resolution theorist lacks resources to explain.

6 Ambiguity resolution reconsidered: Tense and event structure

At this point we have offered several arguments against the ambiguity view, resting on a variety of conceptual and empirical grounds. In constructing our arguments, a recurring frustration we have had with L&S's own argumentation is that while they repeatedly point to a role for convention for the phenomena they address, they typically do not offer detailed linguistic analyses of these phenomena, such that the reader can clearly see how convention and disambiguation combine to cover the same explanatory ground as accounts based on pragmatic enrichment.

Perhaps the treatment for which they paint the clearest picture is in their analysis of how the temporal ordering of events described in a discourse are recovered. Recall from §3 that one of L&S's central arguments against Gricean treatments of such orderings concerns the contrast between (11a) and (11b), repeated below as (24a) and (24b) respectively.

- (24) a. Oil prices doubled and demand for consumer goods plunged.
- b. Oil prices have doubled and demand for consumer goods has plunged. ([Lepore and Stone 2015](#), p. 117, ex. 129)

Recall that in typical contexts, passage (24a) will be construed to express a Result relation, according to which the doubling caused (and hence preceded) the plunging, whereas (24b) will be construed as a Parallel relation, in which no such causal (nor temporal) relationship is inferred. L&S argue that this difference is problematic for the Gricean analysis, since

the two versions appear to have the same truth conditions, violating the non-detachability criterion on implicatures.

L&S take the failure of the Gricean analysis to show that convention and disambiguation, by themselves, account for temporal interpretation: “it’s logical form that settles whether a sentence has a narrative reading or another possible interpretation” (p. 116). In this section, we take a closer look at how L&S’s appeal to convention can potentially explain the difference witnessed for (24a-b). We ultimately conclude, however, that convention and disambiguation are incapable of doing all the work that L&S assign to them; such an analysis still requires a mechanism for pragmatic enrichment. We then follow in §7 by sketching a middle ground account, one that recognizes robust contributions from both convention and pragmatic enrichment (and their interaction), using the interpretation of tense and event structure as a testing ground for theorizing.

L&S lay out their preferred account (henceforth, ‘the anaphoric analysis’ of tense) in their §§7.2–7.3, which we briefly summarize here. They list three possible temporal interpretations between events described by successive clauses in the simple past: simultaneity, backward movement, and forward movement. Simply put, they capture these possibilities in terms of a similarly three-way anaphoric ambiguity, whereby the simple tense can refer to “the time of the previous event,” a time “immediately leading up to the time of the previously mentioned event,” or a time “immediately following some previously mentioned event” (pp. 121–122). On this proposal, the options for temporal relations are specified by the conventions associated with tense; at best, coherence establishment merely plays a role in disambiguating among these options.

An immediate question that arises is what advantages the anaphoric analysis brings to the table, as it seems largely redundant with explanatory tools (such as coherence establishment) that L&S already recognize. If coherence establishment and other reasoning processes they already accept deliver the right interpretative possibilities (as we aim to show below), it’s hard to see the point of enlisting convention (and reconstruing coherence establishment as a mechanism for disambiguation) to cover the very same explanatory ground.¹⁰

¹⁰This is a slight oversimplification, as there is a difference between the analyses with respect to cases in which a Parallel relation is operative. L&S’s treatment of a simultaneity reading as a distinct interpretative possibility is intended to capture examples like (25).

(25) John played the piano. Mary played the kazoo. (Webber 1988)

The main problem we see, however, is that the anaphoric account cannot cover the same ground that coherence establishment does: pragmatic enrichment will still be required to enrich temporal interpretations beyond those specified by tense on the anaphoric account. To see this, consider (26):

- (26) a. Fred was taken to the hospital.
b. He got bitten by a mosquito and contracted the Zika virus.

On the most accessible understanding of (26), we know three things about the ordering of events: the biting occurred before the hospital trip, the contraction of Zika occurred before the hospital trip, and the contraction of Zika occurred after the biting. These three facts can be plausibly recovered straightforwardly in terms of the establishment of a Result relation between the two events described in (26b), which explains the ordering between the biting and the contraction, and then the inference to an Explanation relation between these events and the trip to the hospital, which explains the other two temporal orderings. Crucially, this explanation does not depend in any way on treating tense as anaphoric.

We don't see that these three temporal relationships can all be recovered by the anaphoric account by itself, however. The fact that the biting occurs before the hospital trip can be recovered by anaphorically resolving the past tense associated with *got bitten* to the interval preceding the hospital trip. The issue is with the other two relationships. Specifically, there seem to be two relevant interpretive possibilities for the simple past associated with *contracted* from which we must pick. On the one hand, if it is anaphorically resolved to the interval preceding the hospital trip, this correctly predicts that the contraction occurred before the trip, but does not capture that the contraction occurred after the biting. On the other hand, if the tense is resolved to the interval immediately following the biting, we can capture that the contraction occurred after the biting, but not that it occurred prior to the hospital trip. Either way, then, the disambiguation of tense only yields one relation, and coherence-driven enrichment will still be required

They note, correctly, that coherence does not deliver this reading; as we've already indicated, Parallel typically does not order events. On the other hand, this ordering can easily be seen as an inference from the assumption that John and Mary are playing together, and even then the inference isn't inevitable (John and Mary could have participated in different songs at the same concert). More to the point, recall that examples like (3a-b) do not impose an order among the events, and we find this to be true of L&S's own (11a) as well. L&S do not explain how such cases can be captured with the three interpretive possibilities they specify.

to obtain the other. And, in so far as these enrichments must be understood as going beyond what is specified by convention and disambiguation (for reasons discussed in §3), it follows that the anaphoric analysis is incapable of eliminating the need for pragmatic enrichment in particular.¹¹

We therefore conclude that the tools provided by the ambiguity resolution view are both unnecessary and insufficient for accounting for the facts about tense and event structure.

7 Tense and event structure redux

So far we have argued against both Gricean views and L&S's ambiguity resolution view, claiming that neither of these extremal positions adequately accounts for language understanding. In light of these considerations, we strongly suspect that the best way forward will involve a mixed account — one that makes room for substantive contributions from both convention and pragmatic enrichment. In this section, we'll investigate how such an analysis might look for examples like (24a-b).

We can start by first asking what the conventional properties of tense — as they pertain to temporal interpretation and the mental representation of event structure — tell us about cases like (24a-b), and then ask where pragmatic enrichment might take over. We can begin by making an uncontroversial observation: that the semantic properties of the tenses used in a passage constrain the types of coherence relations that can be inferred. For instance, whereas (27a) admits of both Explanation and Result interpretations, (27b) only has the Explanation interpretation.

(27) a. John slipped. He spilt a bucket of water. (Lascarides and Asher 1993)

b. John slipped. He had spilt a bucket of water.

L&S, as we have argued, would seek to explain these facts solely by way of the conventional properties of tense. We propose to explain them, on the other hand, with an account that appeals to an interaction between what is conventionally encoded by tense and processes of coherence

¹¹In making their case, L&S, following Partee and others, point to certain analogies between tense and pronominal reference. There is much to say here, but for now we simply point out that there are significant disanalogies as well. For instance, one might wonder why it's acceptable to begin a discourse with a sentence in the simple past when no antecedent is available, as in (26a). This isn't the case with antecedentless pronouns — replacing *Fred* with *He* in (26a) results in infelicity.

establishment.¹² Specifically, we propose that, on the one hand, the past perfect associated with (27b) is anaphoric: it specifies that the event time is ordered prior to an anaphorically-identified reference time, which in this case, will be the event time associated with the slipping. Coherence establishment is then constrained to infer coherence relations consistent with that ordering, which rules out Result. On the other hand, the simple past on our analysis is not anaphoric — instead it merely orders the event time prior to the speech time (cf. Reichenbach 1947). What this means is that the simple pasts in (27a) do not place any constraints on the ordering between two events. Coherence establishment is then free to choose any order, with Result imposing forward movement of time and Explanation backward movement. It would thus seem that the fixing of temporal interpretation follows coherence establishment in this case rather than the other way around. The analysis of these two cases thus instantiates a general picture in which the temporal properties associated with tense may constrain the ordering among events (perhaps only partially), and then coherence establishment, while adhering to those constraints, may further enrich the temporal relationships that are ultimately conveyed.

It is likewise clear that coherence establishment is sensitive not only to tense, but to conventions that pertain to the mental representation of event structure as well. Consider the difference between the perfective and imperfective forms in (28a-b):

- (28) a. Andy handed the corkscrew to Jonathan. Jonathan opened the wine.
- b. ?? Andy was handing the corkscrew to Jonathan. Jonathan opened the wine.

Whereas (28a) is a perfectly coherent Occasion relation, (28b) is odd. The reason is evident: even though the handing event occurred in the past, it is described in (28b) as if it is in process, which is to say the focus is on the ongoing development of the event rather than its consequences. This creates a problem for the inference to Occasion, which requires a salient consequent state for the previous event to serve as the presumed initial state for the subsequent one (Hobbs 1990; Kehler 2002); with no salient consequent state provided by (28b), incoherence results. Note that interpretation could have been such that hearers would simply accommodate the fact that the handing event had successfully completed and hence interpret

¹²This account is spelled out in greater detail in Kehler (2002, ch. 7).

(28b) much like (28a), but that's not how it works. It matters where the hearer's focus resides within event structure in his mental model of the discourse.

With these observations in hand, let us now return to (11a-b). Whereas (11a-b) may have the same truth conditions, it is well-known that the simple past and present perfect are not fully interchangeable. Consider the first clauses of (11a-b), with the possible follow-ons given in (29a-b):

- (29) a. Oil prices doubled (but then promptly retreated soon afterward).
- b. Oil prices have doubled (?? but then promptly retreated soon afterward).

Unlike (29a), (29b) sounds odd with a continuation that makes it clear that the state of affairs that resulted from the doubling — i.e., prices that are twice as high than at a salient prior time — is no longer true at the speech time. The analysis of [Moens and Steedman \(1988\)](#) that L&S appeal to captures this through an interaction between times (particularly event, speech, and reference times, per Reichenbach) and event structure. By situating the reference and speech times associated with the present perfect in the consequent state of event structure, they capture the intuition that, in L&S's words, the purpose of the present perfect is to “reference particular consequences of an event located indefinitely in the past, and to present those consequences as still holding in the present”.

In light of the effect we saw for (28a-b), it is perhaps not surprising that the component of event structure that a particular choice of tense and aspect places in focus would affect coherence in (11a-b) as well. Here, the meaning of the present perfect — by placing focus on the fact that the result state of the first event continues to hold at the speech time — may disrupt the establishment of a Result relation in cases in which the effect is described as holding before the speech time, since Result orders the relevant times in forward progression. If this is the case, then establishing a Result relation when the first clause is in the present perfect should only be possible if the event described in the second clause not only occurred subsequent to the first event, but not prior to the speech time as well.

The predictions associated with this conjecture are easy enough to test; consider (30a-d):

- (30) a. Oil prices doubled and so I'm going to start taking the train to work (next week).

- b. Oil prices doubled and so I started taking the train to work (last week).
- c. Oil prices have doubled and so I'm going to start taking the train to work (next week).
- d. ?? Oil prices have doubled and so I started taking the train to work (last week).

Whereas both (30a-b) are fine in the simple past, the present perfect versions in (30c-d) differ. Specifically, (30d) is odd because the focus on present circumstances that arises from the first clause is incompatible with moving back to the past in the second clause. A continuation that talks about a result that will happen after the speech time, on the other hand, is fine, per (30c).

This explains why the inference to Result in (11b) would be disrupted as well, since the initial event is similarly described with the present perfect, and the second described as having occurred in the past. Indeed, our explanation makes a specific prediction: that the problem with (11b) is due to the first clause being in the present perfect, and not the second. This prediction is confirmed by the status of the following two variants: the version with the first clause in the present perfect and the second in the simple past is odd on a Result interpretation (31a), whereas the version with the first clause in the simple past and the second in the present perfect is fine (31b).

- (31) a. ?? Oil prices have doubled and demand for consumer goods plunged.
- b. Oil prices doubled and demand for consumer goods has plunged.

This analysis thus shows how convention can play a greater role in determining construals than Grice envisaged, while at the same time maintaining a role for pragmatic enrichment. Needless to say, there remains much to say about these particular examples and the explanations on offer. Still, we hope we have made the case that L&S's criticisms of Gricean accounts — compelling as they are — do not, by themselves, license the conclusion that there is no important role to be played by pragmatic enrichment. Rather, we take these considerations to show that there exist promising explanations of the phenomena under consideration that give substantive roles to both convention and pragmatic enrichment, but where both of these components and their relationship will have to be understood at a more detailed level than Grice typically pursued.

8 Conclusion

We find L&S's case against the explanatory excesses of the Gricean program persuasive. They have performed an important service for the field in showing the shortcomings of both Griceans' radically minimal conception of what is said, and their equally radical treatment in terms of the single notion of implicature of the diverse range of phenomena going beyond this minimal notion of what is said.

On the other hand, we find that L&S err in the opposite direction by propounding an overunified analysis of their own. While we accept their conclusion that convention plays a greater role than has typically been appreciated, and that it fixes far more at the level of what is said than Griceans allow, we do not believe that all of the interpretive effects L&S hope to explain are adequately accounted for in terms of ambiguity resolution between conventionally specified alternatives.

At the end of the day, we believe that an adequate account of language understanding will have to recognize both extrasemantic expansion (including but not limited to implicature) and a robust role for convention in fixing interpretive alternatives. Indeed, we find it unsurprising that both forces should figure centrally in language use. After all, a language that did not take advantage of context and its users' knowledge stores and capacity for inference would thereby pass up significant opportunities for improved communicative efficiency with respect to the speaker. And at the same time, a language that failed to conventionalize certain common relationships that would otherwise be left to inference would likewise ignore significant opportunities for improved communicative efficiency with respect to the hearer. We expect that the investigation of the precise ways in which these forces are balanced within languages will become a rich source of progress in the field.¹³

References

- Asher, Nicholas and Alex Lascarides (2003), *Logics of Conversation*, Cambridge University Press, Cambridge.
- Bach, Kent (1994), "Conversational Implicature," *Mind and Language*, 9, 2, pp. 124-162.
- Cohen, Jonathan and Andrew Kehler (2016), "Conversational Eliciture."

¹³We are grateful to participants in the 2017 Philosophy Desert Workshop, and especially to Will Starr, for helpful discussion of an earlier version of this paper.

- Grice, H. Paul (1975), "Logic and Conversation," in *Syntax and Semantics*, ed. by Peter Cole and Jerry L. Morgan, Academic Press, New York, vol. 3, pp. 41-58.
- Hobbs, Jerry R. (1979), "Coherence and Coreference," *Cognitive Science*, 3, pp. 67-90.
- (1990), *Literature and Cognition*, CSLI Lecture Notes 21, Stanford, CA.
- Horn, Laurence R. (1989), *A Natural History of Negation*, University of Chicago Press, Chicago.
- (2016), "Conventional Wisdom Reconsidered," *Inquiry*, 59, 2, pp. 145-162.
- Kehler, Andrew (2002), *Coherence, Reference, and the Theory of Grammar*, CSLI Publications.
- Knott, Alistair and Robert Dale (1994), "Using Linguistic Phenomena to Motivate a Set of Coherence Relations," *Discourse Processes*, 18, 1, pp. 35-62.
- Lakoff, George (1973), "Some Thoughts on Transderivational Constraints," in *Issues in Linguistics: Papers in Honor of Henry and Renee Kahane*, ed. by Braj B. Kachru, Robert B. Lees, Yakov Malkiel, Angelina Pietrangeli, and Sol Sporta, University of Illinois Press, Urbana, pp. 442-452.
- Lascarides, Alex and Nicholas Asher (1993), "Temporal Interpretation, Discourse Relations, and Common Sense Entailment," *Linguistics and Philosophy*, 16, 5, pp. 437-493.
- Lepore, Ernie and Matthew Stone (2015), *Imagination and Convention: Distinguishing Grammar and Inference in Language*, Oxford University Press, Oxford.
- Levinson, Stephen C. (2000), *Presumptive Meanings*, MIT Press, Cambridge, MA.
- Moens, Marc and Mark Steedman (1988), "Temporal Ontology and Temporal Reference," *Computational Linguistics*, 14, 2, pp. 15-28.
- Reichenbach, Hans (1947), *Elements of Symbolic Logic*, Macmillan, New York.
- Roberts, Craige (1998), "Information Structure in Discourse: Towards and Integrated Formal Theory of Pragmatics," in *OSU Working Papers in Linguistics 49: Papers in Semantics (1996)*, ed. by J. H. Yoon and Andreas Kathol, Revised version as of Sept. 1998, Ohio State University, pp. 91-136.
- Rohde, Hannah, Roger Levy, and Andrew Kehler (2011), "Anticipating Explanations in Relative Clause Processing," *Cognition*, 118, pp. 339-358.
- Sadock, Jerrold M. (1974), *Towards a Linguistic Theory of Speech Acts*, Academic Press, New York.

- Sperber, Dan and Deirdre Wilson (1986), *Relevance: Communication and Cognition*, Harvard University Press, Cambridge, MA.
- Szabó, Zoltán Gendler (2016), "In Defense of Indirect Communication," *Inquiry*, 59, 2, pp. 163-174.
- Webber, Bonnie Lynn (1988), "Tense as Discourse Anaphor," *Computational Linguistics*, 14, 2, pp. 61-73.
- (1991), "Discourse Modelling: Life at the Bottom," in *AAAI Fall Symposium Series on Discourse Structure in Natural Language Understanding and Generation*, American Association for Artificial Intelligence, Asilomar, CA, pp. 146-151.
- Wierzbicka, Anna (1985), "Different cultures, different languages, different speech acts: Polish vs. English," *Journal of Pragmatics*, 9 (2-3 1985), pp. 145-178.