

Conversational Eliciture*

Jonathan Cohen[†] and Andrew Kehler[‡]

It is an obvious point that the apparatus of model-theoretic semantics is not sufficient to predict the choice of a particular description of an object from among many semantically suitable ones [Suppes, 1973, 393].

Abstract

The sentence *The boss fired the employee who is always late* invites the defeasible inference that the speaker is attempting to convey that the lateness caused the firing (cf. *The boss fired the employee who is from Philadelphia*, which does not invite an analogous inference). We argue that such inferences cannot be understood in terms of familiar approaches to extrasemantic enrichment such as implicature, implicative, explicature, or species of local enrichment already in the literature. Rather, we propose that they arise from more basic cognitive strategies, grounded in processes of coherence establishment, that thinkers use to make sense of the world. Attention to such cases provides a richer and more varied landscape of extrasemantic enrichment than has been appreciated to date.

1 Pragmatic enrichment

Zipf [1949] famously posited two opposing desiderata in language design. The first is what he called the AUDITOR'S ECONOMY, which concerns the desideratum of EXPRESSIVENESS: Languages should be expressive enough that hearers can recover the speaker's message with minimal interpretive effort. The second is the SPEAKER'S ECONOMY, which amounts to the desideratum of EFFICIENCY: Languages should allow speakers to get their message across with minimal articulatory effort. One way that speakers manage to be economical while

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

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

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

remaining expressive is by designing their utterances to take advantage of the hearer's mental state and capacity for inference to communicate more than what is explicitly said. That is, speakers can, in certain cases, rely on their hearer's ability to perform PRAGMATIC ENRICHMENT to recover content that a speaker intends to convey; in such cases the speaker can avoid having to convey the content explicitly.

The source of such pragmatically-determined aspects of sentence meaning has occupied the attention of researchers interested in the semantics and pragmatics of language for many years, and became an industry of its own after the seminal work of Grice [1975]. The goal of theories of pragmatic enrichment is to offer accounts of the myriad of ways in which language interpreters come to identify elements of the message that a speaker intends to convey that go beyond the literal meaning associated with her linguistic contribution. Theories of enrichment have thereby historically engaged with several questions. First, theories typically provide a characterization of the type of enrichment at hand, as means for identifying those components of conveyed content that qualify as having resulted from that type of enrichment. Second, when applicable, they identify those aspects of the hearer's cognitive apparatus that a speaker exploits when constructing her utterance so as to convey the enrichment. Third, they characterize (as explicitly as possible) the nature of the inference processes associated with that cognitive apparatus that the hearer uses to recover the enrichments. Finally, they characterize the triggers (linguistic or not) responsible for the hearer's initiation of those inferential processes, when applicable.

Answers to these questions are necessary if one is to understand both the cognitive basis of a particular species of enrichment and the rationale for the classification of different types of enrichment into distinct categories. For instance, in positing his notion of IMPLICATURE (§2.1), Grice provides answers to all four of the aforementioned questions, describing a rational reconstruction of an inferential process that yields implicit conveyed propositions through a process triggered by the threat of communicative failure, in light of mutual assumptions about the interlocutors' rationality and cooperativity. While endorsing a Gricean framework at a general level, Bach differentiates his notion of IMPLICITURE from implicature by pointing to differences in both the types of communicative failure serving as the trigger and the nature of the inferences themselves (i.e., serving to narrow the meaning of an explicitly conveyed proposition rather than yielding a distinct proposition that sits alongside an explicitly conveyed one; see §2.2). Researchers investigating forms of local pragmatic strengthening (§2.3) have likewise argued for differentiations based on the local nature of the enrichments and other factors. Detailed theories of enrichment that engage with these questions are necessary if we are to understand the full range of respects with which languages have evolved to balance expressiveness with efficiency.

In this paper, we focus on a process by which extrasemantic content is conveyed that, we claim, fails to fit neatly into any of the types of pragmatic enrichment thus far described in the literature. To see the kinds of cases we have in mind, consider examples (1a-c):

- (1) 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*.
- c. The company fired *the manager who has a long history of corporate awards*.

In uttering example (1a), a speaker strongly invites the hearer to infer that she intends to communicate not only that the manager was embezzling *and* was fired, but that the embezzlement was the *reason* for the firing.¹ Note that this is merely a defeasible inference: (1a) could be followed with *The reason the manager was fired was because he was rude and always late*. In (1b), however, being hired in 2002 will normally not be understood to be the cause of the firing; here the relative clause is merely identificational. Example (1c) is a case that leads to a counter-to-expectation inference, leading the hearer to wonder why a manager with a positive history with the company would be fired.²

The differences in the interpretations of (1a-c) hinge on the choice of relative clause used in a referring expression. Similar inferences derive from other aspects of referring expressions as well, including the content of adjectival phrases. Example (2a), for instance, will generally be taken to indicate that the drugs *caused* the undergrad to fall off of the cliffs:

- (2) a. *The drug-addled undergrad* fell off of the Torrey Pines cliffs. (adapted from an example of Webber [1991])

¹For ease of exposition, in what follows we may at times speak in terms of utterances inviting addressees to infer some content X , when what addressees actually infer is that the speaker intends to communicate X , since understanding the content of a message does not require accepting/endorsing its content. It is this latter interpretation we intend. Thanks to Kent Bach for a discussion of this point. We further elaborate on the relationship between the speaker and hearer in making such inferences in §3.1.

²The existence of these inferences was confirmed by a Mechanical Turk passage completion experiment. Participants ($n = 17$) were provided with sentences like those in (1a-c), which included object-biased implicit causality verbs [Garvey et al., 1976, Caramazza et al., 1977, inter alia], and whose object NPs included relative clauses that were independently judged as encoding an explanation for the event described (type Expl, as in 1a), neutral (type NoExpl, as in 1b), or expectation-violating (type ViolExp, as in 1c). Participants supplied a follow-on sentence to complete the passage, and two annotators blind to the hypothesis coded the data. Contexts with implicit causality verbs in such studies have been previously shown to yield a substantial percentage of completions that provide an explanation (i.e., cause or reason) for the described event [Kehler et al., 2008]. But if participants inferred the relative clause to provide an explanation in the Expl condition, we predict that participants would provide fewer completions that provide explanations in this condition than the other two. This is indeed what we found: Participants wrote fewer explanations in the Expl condition (38.9%) compared to the NoExpl condition (75.4%; $p < .001$) and the ViolExp condition (65.2%; $p < .001$). See also Rohde et al. [2011] and Kehler and Rohde [2019] for experimental studies that utilize examples like (1a-b) to demonstrate how the inference of an explanation from a relative clause influences relative clause attachment and pronoun interpretation respectively.

- b. *The well-liked undergrad* fell off of the Torrey Pines cliffs.
- c. *The normally risk-averse undergrad* fell off of the Torrey Pines cliffs.

No such inference is typically drawn for (2b), however: presumably being well-liked was not a cause of the falling. And again, (2c) yields a counter-to-expectation inference, leading one to be surprised that a normally risk-averse undergrad would fall off of the cliffs.

Finally, extrasemantic inferences may result from the choice of nominal itself; consider (3a-b):

- (3) a. *A jogger* was hit by a car in Fresno last night. [Hobbs, 1990]
- b. *A farmer* was hit by a car in Fresno last night.

The use of the NP *a jogger* in example (3a) strongly invites the hearer to infer, but does not entail, that the the victim was jogging at the time of the accident, and hence is not merely someone who happens to jog a lot (but who was walking or bicycling when the accident occurred). In contrast, the analogous inference for (3b)—that the farmer was farming at the time of the accident—is much less naturally evoked, and indeed the utterance is completely felicitous without it.³

The property common to these examples is that a speaker’s decision to use a particular way of referring to an entity over other alternatives invites the hearer to draw inferences that, crucially, are not triggered by any syntactic relationship or other type of felicity requirement on linguistic material. For want of an appropriate term of art, we brand this phenomenon *ELICITURE*, a term intended to capture the fact that, by choosing a particular form of reference, a speaker elicits inferences on the part of the hearer that would not otherwise be drawn. In this paper, we will argue that the inferences under scrutiny differ in crucial ways from other forms of enrichment that have been posited in the literature with respect to the questions

³One might object that what we are construing as an interesting disanalogy between cases that invite an extrasemantic inference (e.g., 3a) and those that don’t (e.g., 3b) is more simply understood by noting that the nominal *a jogger* admits of a stage-level interpretation more readily than the corresponding nominal *a farmer*. However, we see this difference as a redescription of, and not an alternative explanation for, our observation. For an inference that the activity connected with the nominal was occurring at the time of the matrix event will necessarily imply temporariness for that activity—which just amounts to the stage-level interpretation. It is worth noting that these temporal properties are not inherent to the nominals themselves: The sentence *Fred talked to a jogger at the party last night* does not readily evoke the inference that the jogger was jogging at the time (Fred’s interlocutor is likely a hobbyist), which in turn entails an individual-level interpretation, whereas the sentence *A farmer flipped over his tractor yesterday afternoon* does trigger an extrasemantic inference analogous to (3a), yielding a stage-level interpretation. What cries out for explanation, then, is just why these nominals can be used with a stage-level interpretation in just those cases in which the extrasemantic inference is licensed. In what follows we’ll be attempting to answer this and similar questions in general pragmatic terms. (Thanks to Martin Schäfer and Chris Kennedy for pressing us on this point.)

outlined at the outset. Indeed, we emphasize that what is inferred in these cases does not result from presumptions about conversational norms that cooperative interlocutors are presumed to follow, nor do they involve the recovery of a value for an unsaturated (explicit or hidden) parameter, a strengthening of the meaning of a particular constituent, nor a filling out of an otherwise underspecified logical form. On the contrary, the inferred information in our examples —e.g., the information that, with respect to (1a), the embezzlement was the reason for the firing— is not a necessary component of, and appears to go well beyond, the logical form of the sentences that invite these inferences. Examples of eliciture reveal particular way in which a speaker’s choice of referential expression can be used to exploit the hearer’s world knowledge and capacity for inference as a way to make her linguistic contribution more efficient as a vehicle for communicating content. The question, however, is exactly how these inferences come to be drawn by the hearer. Whereas examples of this sort have been noted in the literature for some time (as the citations for examples (1)–(3) make clear), we aim to show that their place in, and ramifications for, pragmatic theory have not been fully appreciated.

Our strategy will be as follows. Below, we briefly describe some of the more familiar categories of pragmatic enrichment that have been proposed in the literature, and explain why elicitures constitute a distinct species (§2). We then follow with a detailed characterization of elicitures that answers the key questions to be addressed by a theory of enrichment as outlined above (§3). Finally, we revisit the place of elicitures in a general theory of pragmatics (§4).

2 Familiar approaches

A natural reaction to the examples in §1 is the thought that they can be assimilated to standard forms of extrasemantic content that have become familiar in the literature. Though understandable, we believe this strategy is unsuccessful—for, we aim to show, our examples reveal aspects of extrasemantic enrichment that go beyond familiar accounts.

2.1 Implicature

Grice (1975), of course, introduced the notion of implicature to describe the communication of extrasemantic content. Are the extrasemantic inferences in (1)–(3) the result of implicature? An initial examination might suggest that they are. After all, these inferences turn on the form of referring expression chosen by a speaker, and Grice himself noted that such choices can give rise to extrasemantic inferences that have the hallmarks of implicature. Considering the utterance *X is meeting a woman this evening*, he suggests that “a speaker would normally implicate that the person to be met was someone other than X’s wife, mother, sister, or perhaps even close platonic friend” (p. 56). Evidence that implicatures are at play in such cases is provided by the hallmark diagnostics of CANCELABILITY and

REINFORCEABILITY [Grice, 1975, Horn, 1984, Sadock, 1978, Hirschberg, 1991], illustrated respectively in (4a-b):

- (4) a. John is meeting a woman this evening, in fact she's his wife.
- b. John is meeting a woman this evening, but not his wife.

That is, the implicature that the woman X is meeting is not his wife is canceled without contradiction in (4a) and likewise reinforced without redundancy in (4b), establishing that the inference is in fact not part of what is said. Suggestively, the same tests succeed for the cases under scrutiny here, e.g. (1a):

- (5) a. The company fired the manager who was embezzling money, but the embezzlement isn't why he got fired.
- b. The company fired the manager who was embezzling money, and in fact the embezzlement is why he got fired.

Hence, these considerations naturally invite the suspicion that the extrasemantic inferences witnessed in examples (1a/2a/3a) are similarly the result of implicatures.⁴

⁴ Kronfeld [1990, p. 92] in fact offered a Gricean explanation of an extrasemantic contrast present in a conceptually similar but more restricted class of cases. Assume for illustrative purposes that Washington D.C. is both the city with the greatest diversity of languages spoken and also the murder capital of the world, and in which its mayor opens a meeting of the Linguistics Society of America with one of the utterances (6a-c):

- (6) a. *Washington D.C.* welcomes the meeting of the Linguistic Society of America.
- b. *The city with the greatest diversity of languages spoken* welcomes the meeting of the Linguistic Society of America.
- c. *The murder capital of the world* welcomes the meeting of the Linguistic Society of America.

(See Grice [1978, 114] for a similar example.) Kronfeld proposes to explain the difference in pragmatic effects in (6a-b) by claiming that a hearer will regard the referring expression used in the subject position of (6b) as a needlessly/irrelevantly prolix, and hence Quantity- and Manner-violating, version of that which appears in (6a), and attempt to restore conformity with that norm by establishing the 'conversational relevance' of the referring expression, by way of identifying an intensionally justified universal generalization that the utterance could be understood to express, such as that in (7a).

- (7) a. In view of *f*, any city with the greatest diversity of languages spoken must welcome the meeting of the Linguistic Society of America. [*f*=the fact that linguists like to encounter linguistic diversity]
- b. In view of *f*, any murder capital of the world must welcome the meeting of the Linguistic Society of America. [*f*=?]

By the same logic, the oddity of (6c) results from the hearer's inability to identify a plausible, suitable universal generalization of this sort (7b).

We see a number of problems with Kronfeld's implementation, many of which arise from the restriction of his analysis to simple predicational sentences with definite descriptions

To be sure, the phenomenon of eliciture can be brought under a Gricean umbrella if one takes a sufficiently broad view of what counts as Gricean enrichment as to include *any* form of meaning expansion that the speaker intends the hearer to recover: our cases clearly share many relevant features with other, previously documented types of enrichment.⁵ On the other hand, part of what makes Grice’s contribution so important is that he goes beyond this broad view of implicature to provide specific answers to the questions outlined in the introduction. And from this perspective, we can see that eliciture differs from implicature (and for that matter, implicature; see §2.2) in crucial respects that adequate theories of pragmatic enrichment must recognize.

As we understand him, Grice [1975, pp. 49–50] proposes that an enrichment counts as an implicature if it could be reached (in his words, “worked out”) by a hearer through a version of the following four step procedure. First, when the speaker utters a string, the listener recovers its literal semantic content *p* (by exercise of his semantic competence, plus perhaps mechanisms for disambiguation, and for assigning values to various context-dependent elements demanding contextual supplementation). Second, he notices that the expression by the speaker of the content *p* in the circumstances would violate one or more Gricean norms. Third, since he takes it that speakers ordinarily abide by these norms, he infers that the utterance must have as a secondary purpose the conveyance of some distinct, extragrammatic content *q* whose expression by the speaker is in compliance with Gricean norms. And fourth, he therefore treats the utterance as conveying the content *q* (in addition to the content *p*).

A crucial feature of this account is that, as presented, implicature rests in an important way on a threat of failure—specifically on the idea that, without the extragrammatic enrichment, the speaker’s utterance would violate assumptions about her conformity to rational and/or cooperative norms.⁶ In contrast, we argue

as subject terms. As our examples make clear, the phenomenon extends to many other constructions—including indefinites (e.g., in (3)), adjectival phrases occurring within definite descriptions (e.g., in (2)), and relative clauses (e.g., in (1))—and it is far from clear how his analysis can be made to extend to such cases. However, we will forgo elaboration of our arguments against his particular analysis in favor of ones that apply to any Gricean analysis of eliciture, which includes Kronfeld’s more limited treatment.

⁵Thus, our approach is broadly Gricean in just the way that Bach says his approach to the newly proposed category of implicature is:

What makes my approach Gricean is essentially this: in figuring out what a speaker means an addressee presumes that the speaker intends him to figure this out [Bach, 2010, p. 3].

Here, Bach emphasizes the crucial role of the recognition of communicative intentions in the conveyance of implicature, while at the same time arguing that implicature is fundamentally distinct from implicature. We follow a parallel strategy for eliciture here.

⁶ Several colleagues (Emma Borg, Alexander Dinges, Chris Kennedy) have suggested to us that Grice himself appears to have endorsed the idea that implicature need not rest on norm-violations, citing the first pair of examples he offers after characterizing the notion of conversational implicature: “Group A: Examples in which no maxim is violated, or at least

that elicitures neither rest on the threat of communicative failure nor engage assumptions about speaker rationality and/or cooperativity in the same manner. To illustrate our point, it will prove useful to step through possible analyses that appeal to one or more of Grice's four maxims. This notwithstanding, we wish to stress that the arguments ultimately apply more generally; they cannot be overcome via the inclusion of additional maxims beyond Grice's repertoire.⁷

We see no obvious strategy for pursuing an explanation of eliciture based on the maxims of Quality or Quantity. As noted in fn. 4, Kronfeld previously addressed a related set of cases by way of violations of the proximity submaxim of Manner, so we briefly consider this strategy first. Whereas we find his approach to the sorts of cases he addresses appealing, it is easy to see that proximity in the face of a salient alternative expression is not necessary to trigger the inference to an eliciture. For instance, consider a situation in which a speaker and hearer at a company are aware of another employee, and where, while they do not know his name, they share knowledge of various of his characteristics, including that he has red hair, a

in which it is not clear that any maxim is violated" (p. 32). (This suggestion is developed by [Dinges \[2015\]](#).)

However, our understanding of Grice on this point is that this characterization applies to these cases *after* the meaning of the utterance is modified so as to bring the utterance into norm-compliance. That is, as we read him, Grice's position is that there *would be* a clash between what is conveyed and the Relation maxim, *were it not that* ("B would be infringing ... unless ..." (p. 32)) the hearer can safely attribute further assumptions (which therefore gain the status of implicata) to the speaker: "In both examples, the speaker implicates that which he must be assumed to believe in order to preserve the assumption that he is observing the maxim of relation" (p. 32). This contrasts with what happens in examples in Groups B and C, where a maxim remains violated after the implicature is drawn—either because of a clash with another maxim (Group B) or because the maxim is flouted (Group C). On this characterization, therefore, implicatures do indeed arise from threatened norm violations. This contrasts with elicitures, which, we claim, need not rest on a threatened norm violation or other linguistic lapse.

⁷We again appeal to Bach on this point:

Although Grice presented [his maxims] in the form of guidelines for how to communicate successfully, I think they are better construed as presumptions about utterances, presumptions that we as listeners rely on and as speakers exploit. As listeners, we presume that the speaker is being cooperative (at least insofar as he is trying to make his communicative intention evident) and is speaking truthfully, informatively, relevantly, and otherwise appropriately. If an utterance superficially appears not to conform to any of these presumptions, the listener looks for a way of taking it so that it does conform. He does so partly on the supposition that he is intended to. As speakers, in trying to choose words to make our communicative intentions evident, we exploit the fact that our listeners presume these things [[Bach, 2006](#)].

As far as we can tell, Bach's characterization enjoys a broad consensus among (neo-Gricean) implicature theorists. We will argue that elicitures do not fit this characterization: the inferences in question are *not* triggered by any respect in which "an utterance superficially appears not to conform to any of these presumptions."

beard, glasses, and that he generally arrives late to work. In such a context, (8a) will generally be taken to convey the information that the employee's lateness was a cause of John's firing her, even though it is not more prolix than alternatives containing competing referring expressions, such as (8b), that do not evoke the inference. Further, example (8c) is more prolix than (8a), yet does not yield the inference that the employee's red hair, beard, and glasses were causes of John's firing her.⁸

- (8) a. John fired the employee who was always late.
- b. John fired the employee who has red hair.
- c. John fired the employee who has red hair, a beard, and glasses.

What the injunction against prolixity *does* predict is that (8c) may well be odd in our hypothesized situation if the referent is the only employee with red hair and the interlocutors are aware of this—i.e., in a situation in which the less prolix relative clause in (8b) would have sufficed. But as these examples show, this phenomenon is orthogonal to the question of under what conditions a cause will be inferred from the relative clause.⁹

The implicature theorist might answer that the extrasemantic inference at work in (8a) is a result of a violation of a Gricean norm of Relation/relevance rather than the norm of Manner/brevity. We see several reasons why such an analysis fails.

One worry (aside from the question of what relevance means; Grice's admonition to 'be relevant' is of notoriously little help) is that examples such as (8b-c) demonstrate that the function of allowing the addressee to identify the correct referent is sufficient to establish the relevance of a restrictive relative clause. That is, whereas (8b) may be odd in a context in which there is only one employee, it is fine if there is more than one (assuming that only one has red hair, of course). But the restrictive relative clause in example (8a) necessarily serves this identificational role as well—from a set of employees, it will pick out the one who is always late—and thus by virtue of that function it satisfies the same constraints on relevance that are operative in (8b-c). But with relevance thereby established by this identificational function, we are left without an explanation for why any additional inferences

⁸Indeed, a speaker might choose to utter (8c) in a situation in which (8a) would have succeeded in making the referent identifiable by the hearer, if she wanted to avoid the inference that the lateness caused the firing.

⁹A further reason for resisting treating elicitures as Manner/prolixity implicatures is that they fail to exhibit a hallmark feature of the latter—the characteristic narrowing of the denotational space of a constituent based on the existence of a less-prolix competing form that Horn [1984] dubs a division-of-labor effect. For instance, whereas it is the competition between the more prolix "John caused Bill to die" and the less prolix "John killed Bill" that results in the former's implicating indirect causation; there is no division-of-labor involved in elicitures at all: the difference between the referring expressions in (1a-c), for instance, only concerns the linguistic form used to refer to a particular referent.

of the sort drawn for (8a) would be necessary to satisfy the Relation/relevance maxim.¹⁰

What these considerations bring out is that eliciture lacks a central feature of implicature: the former, unlike the latter, need not arise from any threatened violation of rational or communicative norms (whether they be from Grice's inventory or otherwise). We take this point to reveal a crucial difference between eliciture and implicature, and therefore to provide a compelling argument against assimilating the former to the latter.

Moreover, the case becomes even stronger if one thinks of Grice's explication of implicature as providing a schematic outline of the psychological processes leading to extrasemantic expansion.¹¹ For one thing, any account that treats the enrichments of interest as triggered by a threat of communicative failure faces the difficulty that it's unclear what could serve as a trigger: it seems that *no* individual constituent is sufficient to bring about an eliciture on its own. Consideration of examples (1a–c) might initially suggest otherwise: at first blush, one might see these cases as showing that the relative clause (here, in the object NP) gives rise to the extrasemantic inferences when they obtain. However, examples (11a–d)—encountered, say, at the beginning of a newspaper story—make this position untenable.

¹⁰An anonymous reviewer asks whether the identificational function of restrictive RCs isn't as much a matter of inference as elicitures are, whereby identification and eliciture are just two different ways that a hearer could justify the speaker's inclusion of an RC. If this were the case, then the argument expressed above would have no force. However, we do not share this judgment, as we find restrictive RCs to be infelicitous in cases in which they are not required for referent identification, even if they carry an eliciture. Consider (9):

(9) The current CEO of IBM who embezzled money was just fired.

Setting aside situations in which IBM has more than one current CEO, we find example (9) to be felicitous only if the RC is interpreted as non-restrictive, as (10) necessarily is:

(10) The current CEO of IBM, who embezzled money, was just fired.

Examples like (10), of course, demonstrate that elicitures can involve contents evoked from non-restrictive RCs as well, and indeed in such cases it is up to the hearer to determine why the speaker chose to include an optional attributive RC. We focus on restrictive RCs precisely because their role in narrowing the domain of reference is obligatory, and hence any eliciture drawn is necessarily *in addition to* that function.

¹¹We put this claim conditionally because it is controversial whether to take Grice's own understanding of implicature as psychologically committed in this way. Thus, while Levinson [2000, e.g., ch. 1] appears to accept a psychological construal of implicature, Saul [2002] reads Grice as offering a mere rational reconstruction of the inference processes that underlie a class of communicated contents, i.e., an account with no aspirations to psychological reality. (This treatment of Grice is central to her defense of his views from Relevance Theorists' objections that the view is psychologically implausible.) Whatever one's views on this issue, however, we emphasize that the argument already provided (about the independence of eliciture from threatened communicative failure) stands, and provides a sufficient reason for not assimilating eliciture to implicature on its own.

- (11) a. A drunk pilot was arrested yesterday.
 b. A 53 year-old pilot was arrested yesterday.
 c. A drunk rapper was arrested yesterday.
 d. A drunk pilot was playing golf yesterday.

Sentence (11a) will typically give rise to a rich mental picture of a pilot who was arrested because he was flying (or perhaps preparing to fly) while inebriated. This picture arises despite the fact that (11a) could be used to truthfully describe a situation in which a pilot, who happened to be drinking a fair bit on his day off, got arrested at home for cheating on his taxes. The variants in (11b–d), on the other hand, invite no analogous inferences, despite the fact that (11b) contains the same head noun and verb phrase as (11a), (11c) contains the same adjective and verb phrase as (11a), and (11d) contains the same head noun and adjective as (11a). The reason for this is intuitively clear: the basis for drawing the eliciture depends on a causal rule that requires all three constituent meanings provided by (11a) to be instantiated, that is, that a pilot—crucially, when flying or preparing to do so—can be arrested for being inebriated. The speaker who utters example (11a) therefore takes advantage of this knowledge being in the common ground to convey her message in a particularly efficient way.

Finally, if these examples show that eliciture cannot be seen as arising from an externally-triggered search for relevance, reflection suggests that such a picture implausibly reverses the order of explanation in any case: it is the operation of the hearer’s cognitive capacities for drawing connections between contents that leads to a recognition of a relation of relevance (and to the enrichment itself), and not the reverse. For instance, when we interpret (1a), it is not that our understanding of the semantic contents of the matrix verb and the relative clause first provides a signal (*Relation holds!*) to the effect that those contents stand in a relevance relation, and which a search is required to identify (*hmm, firing and embezzling money; those two things are somehow related to one another, now I need to go figure out how*). Nor is it the case that interpreting (11a) results from the explicit consideration of ways in which the meanings of all three constituents might be related (presumably to occur only after performing three distinct pairwise searches that attempt to relate each of the three combinations of two constituents, which as (11b–d) illustrate, will turn up empty). Instead, it is our cognitive apparatus, in the face of learning of the firing and the embezzlement in (1a), and the drunkenness, piloting, and arrest as an ensemble in (11a), evokes the potential causal relationship automatically. Hence, an appeal to the Relation maxim here is superfluous: the only possible trigger for the inference is the very machinery that we have for establishing relevance itself, as it can only be this very machinery that presents us with the candidate inference in the first place. We will return to this issue in §3.2. For our present purpose of resisting the assimilation of eliciture to Gricean implicature, what is important is that an appeal to relevance in such cases puts the cart before the horse: It is not a violation of the Maxim of Relation that starts the inference process; instead the

inference process is already running, picking up the additional relevance relation when suggested by the content of the utterances.

All told, then, we conclude that eliciture lacks hallmarks central to the notion of implicature. Eliciture is a distinct species of extrasemantic expansion, and one whose understanding is poorly served by assimilating it to implicature.

2.2 Implicature

Even if it is true, as we have contended, that elicitures cannot be seen to result from Gricean norm violations, one might still hope to account for them in terms of other familiar tools that have been proposed to account for various forms of extrasemantic inference (sometimes as supplements to, and sometimes as replacements for, Gricean tools). For example, Bach [1994] urges that expansions like those indicated in (12) should not be understood as implicatures, because they involve the inference of content “built out of what is said”, rather than “additional propositions external to what is said” as in implicature (p. 141). From this he concludes that there is a level of extrasemantic expansion (“implicature”) that sits between what is said and what is implicated.

- (12) a. Donald is too crazy. [to be a serious contender for President]
b. I haven’t had breakfast. [today]

Bach distinguishes two types of implicatures. The first type, *completions*, involve “the filling in of a propositional radical”, which occurs “because the utterance is semantically underdeterminate and completion is required” [Bach, 1994, p. 126], as in (12a). The second type, *expansions*, involve “the fleshing out of the minimal proposition expressible by an utterance” (p. 126), which occurs as a result of a hearer recognizing “that a speaker cannot be plausibly taken, and therefore does not intend to be taken, to mean what he is saying” (p. 136), as in (12b).

Eliciture differs from implicature in two important respects. First, like implicature, implicature is driven by a characteristic kind of breakdown. Specifically, implicature arises when the output that semantics generates without those forms of expansion is insufficient to fix (in the case of completions) or to fix appropriately (in the case of expansions) truth-conditional content. Elicitures, on the other hand, do not require the types of triggers necessary for these types of expansion: they do not appear to depend on any kind of incompleteness or infelicity, and are not the result of a grammatical shortcoming, a norm violation, lack of relevance to the speaker’s goals, nor any question on the part of the hearer about the plausibility of the unenriched meaning of the utterance. Indeed, sentences that give rise to elicitures are perfectly felicitous even if the inferences under consideration are not drawn.

The second dissimilarity is that implicatures are restricted to inferences that constitute developments of the logical form of an utterance. Elicitures, in contrast, are not properly characterized as developments of a single logical form; they instead, like implicatures, represent inferences to additional propositions that are

external to the logical form of the utterance itself: in (1a) the proposition denoted by *the embezzlement was the reason for the firing* is inferred alongside *the company fired the manager who was embezzling money*, in (2b) the proposition denoted by *the drug-induced confusion caused the fall* is inferred alongside *the drug-addled undergrad fell off of the Torrey Pines cliffs*, and in (3a) the proposition denoted by *the victim was jogging at the time of the accident* is inferred alongside *a jogger was hit by a car in Fresno last night*. It would appear, then, that elicatures fail to meet the defining characteristics of implicature.

2.3 Local pragmatic strengthening

There is yet another class of extragrammatic enrichments —henceforth, LOCAL PRAGMATIC STRENGTHENINGS— that is potentially relevant to our examples. The hallmark of this class is that its members involve a strengthening (or, more generally, modification) in meaning of a particular constituent or implicit relationship between constituents. Once again, one can ask whether the inferences we are highlighting can be understood by subsuming them under this heading. We argue that they cannot.

An early articulation of the local pragmatic strengthening idea is offered by Geis and Zwicky [1971], who introduce cases of “invited inference” such as the conditional perfection inference in (13a) and the strengthening of the conjunction *and* in (14a):

- (13) a. I’ll give you five dollars if you mow the lawn.
 b. I’ll give you five dollars if and only if you mow the lawn.
- (14) a. Martha observed the children at play and smiled with pleasure.
 b. Martha observed the children at play and as a result smiled with pleasure.

They observe that the typical interpretation of the conditional in (13a) is as a biconditional, i.e. a strengthening from *if* to *if and only if* per (13b). Similarly, the typical interpretation of the conjunction in (14a) involves strengthening to a causal relation, per (14b).

Later neo-Griceans have sometimes held [contrary to Geis and Zwicky, 1971, p. 565] that such cases should be thought of as instances of Gricean implicature. Thus, Horn [2004] proposes to treat Geis’s and Zwicky’s cases as instances of what he calls R-implicatures, assimilating them to the domain of the second part of the quantity maxim (“say no more than you must”). The reasoning goes that these expressions have come to have stereotypical meanings, so to spell out *if* as *if and only if* or *and* as *and as a result* is to provide more information than is necessary. Similarly, Levinson [2000], who advocates a three-way division among Q-, M-, and I-Principles, classifies these cases as resulting from I-inferences. His I-Principle 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”). In this category he collects conditional strengthening (13), “conjunction buttressing” (14), inference to stereotypes (*secretary* ⇒ female secretary), and negative strengthening (*I don’t like Alice* ⇒ I positively dislike Alice).^{12 13}

Recanati [1993, 2004, 2010] offers a rather different view —one he explicitly labels ‘non-Gricean’— that analyzes these varieties of local enrichment by a handful of optional/non-mandated processes of pragmatic enrichment that partially overlap with those falling under Levinson’s Enrichment Rule. Recanati’s processes include “free enrichment”/“specification” (“making the interpretation of some expression in the sentence contextually more specific” [Recanati, 2004, 24]; e.g., the bridging inference in (15a)), “loosening” (“a condition of application packed into the concept literally expressed by a predicate is contextually dropped so that the application of the predicate is widened” (p. 26); e.g., in (15b)), and “semantic transfer” (“the output is neither an enriched nor an impoverished version of the concept literally expressed by the input expression ... [but] a different concept altogether, bearing a systematic relation to it” (p. 26); e.g., in (15c)).

- (15) a. Free enrichment: *Mary took out her key and opened the door.* ⇒ Mary opened the door with the key she had taken out.

¹² Levinson also includes a number of other phenomena in this category, making for quite a heterogeneous lot. Some examples:

- (a) Bridging: *John unpacked the picnic. The beer was warm.* ⇒ The beer was part of the picnic.
 (b) Preferred local coreference: *John came in and he sat down.* ⇒ John came in and John sat down.
 (c) Noun-noun compounds: *The oil compressor gauge* ⇒ The gauge that measures the state of the compressor that compresses the oil.
 (d) Possessive interpretation: *Wendy’s children* ⇒ Those children to whom she is a parent.

These are clearly different from the phenomena discussed above in that they are linguistically mandated forms of strengthening. For instance, the definite *the beer* in (a) is not otherwise felicitous if the addressee cannot uniquely identify the referent. Likewise, one has not understood *he* in (b) if its referent is not recovered, *oil compressor gauge* in (c) without recovering some semantic relationship between the nominals, or *Wendy’s children* in (d) without recovering the relationship between the possessive and the head noun.

¹³At one point Levinson seems to acknowledge the need for further inferential machinery beyond that supplied by Gricean mechanisms themselves: in his full spelling out of his I-Principle (p. 114), Levinson gives four specific subcorollaries of his Recipient’s corollary, including the instruction to “Assume the richest temporal, casual and referential connection between described situation or events, consistent with what is taken for granted”. Though we are unsure what Levinson has in mind here (he doesn’t cite examples like ours, and offers no proposal about the source of such inferences), it may be that he would regard elicitures as falling under this heading.

- b. Loosening: *The ATM swallowed my credit card.* ⇒ The (inanimate) ATM took my credit card and didn't give it back.
- c. Transfer: *I am parked out back.* ⇒ The car I own/drive is parked out back.

One way in which Recanati differs from the other theorists mentioned in this section is in holding that the processes responsible for the enrichments in cases like (13)–(14) also apply during the process of semantic composition (Recanati [2004, 34–36]; Recanati [2010, ch. 1]), as in (16a–b):

- (16) a. *He eats/wears rabbit.* ⇒ the meat/fur from a rabbit.
- b. *The driver/policeman stopped the car* ⇒ by applying the brake/by directing the driver to apply the brake. [adapted from Rumelhart, 1979, 78]

The enrichments in these cases still have the property of being local, however, in the sense that they arise from the combination of the meanings of sibling constituents in the syntax.

In contrast, our examples differ from the cases addressed by both Geis & Zwicky and Recanati in their unrestrictedly non-local character. For one thing, as we have seen, eliciture-based inferences include many for which it is difficult to locate any lexical/constructional trigger smaller than the whole utterance (e.g., 3a, 11a). Hence, such cases lack the specific lexical or constructional triggers characteristic of Geis's and Zwicky's examples: e.g., the result reading for 'and', the biconditional for 'if...then', etc.¹⁴ Second, unlike Recanati's examples, the meanings whose association give rise to the enriched interpretation are often not those of sibling constituents in syntax, and hence cannot be triggered via semantic composition [cf. Recanati's rejection of free enrichment at the "topmost level" Recanati, 2010, p. 23]. Third, the inferred extrasemantic content generated by eliciture is propositional, rather than an alternative interpretation of any subpropositional constituent as in the examples addressed here. And finally, as we mentioned in §2.2, the propositional output of eliciture is not a replacement of the unenriched meaning of the utterance, but a supplement that sits alongside the latter (e.g., *The company fired the manager who was embezzling money* ⇒ the embezzlement was the reason for the firing). Consequently, eliciture cannot be analyzed simply as a strengthening of a constituent meaning to a stereotypical interpretation (or other local modification and replacement).

2.4 Relevance Theory

We take the foregoing discussion to have shown that the phenomenon of eliciture is conceptually distinct from standard forms of extrasemantic content that have become familiar within the Gricean tradition and its nearby offshoots. We now want to consider eliciture from the perspective of Relevance Theory (henceforth

¹⁴Reasons for doubting the applicability of Horn's and Levinson's treatment of examples (13)–(14) to cases of eliciture, i.e., by appeal to implicature, were discussed in §2.1.

RT) —perhaps the leading alternative to the Gricean framework for understanding extragrammatic expansion— which aims to treat a surprisingly wide variety of interpretive phenomena (including all pragmatic enrichment) in a uniform way [Sperber and Wilson, 1986, Wilson and Sperber, 2004, *inter alia*]. Here we'll briefly outline RT before considering its application to eliciture, and then argue that (despite important similarities between the view and our own positive account, presented in §3) eliciture ultimately resists explanation in Relevance Theoretic terms.

Broadly speaking, RT treats linguistic interpretation as a matter of inferring speaker meaning from the evidence of heard utterances and other assumptions operative in the context. The key technical notion in the theory is that of the RELEVANCE of an input—a ratio of its number of contextual effects (*viz.*, conclusions derivable from that input and background assumptions together but not from either alone) to the processing costs of obtaining it in that context [Sperber and Wilson, 1986, p. 125]. On the RT picture, heard utterances guide a hearer's recovery of speaker meaning because they raise expectations of their own relevance, as per what Sperber and Wilson call the COMMUNICATIVE PRINCIPLE OF RELEVANCE: "Every utterance (or other act of overt communication) communicates the presumption of its own optimal relevance" [Sperber and Wilson, 1986, p. 260]. Given that, for them, speakers guide their communicative activity in accord with the principle of relevance, hearers should interpret utterances by disambiguating, resolving, and expanding them in the most relevant way—*i.e.*, by pairing them with specific meanings that facilitate the most inferences with the least effort. Thus, on this view,

The hearer should take the linguistically encoded sentence meaning; following a path of least effort, he should enrich it at the explicit level and complement it at the implicit level until the resulting interpretation meets his expectations of relevance [Wilson and Sperber, 2004, p. 258].

This conception of interpretation is ultimately enshrined in the theory's single interpretive procedure, the RELEVANCE-THEORETIC COMPREHENSION PROCEDURE [Wilson and Sperber, 2004, p. 259]:

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

Applied to an instance of eliciture like (1a), the proposal would presumably be that interpreting the utterance as communicating the information that the embezzling caused the firing represents a gain in relevance over an interpretation that omits this information, in so far as the former interpretation adds important contextual effects at relatively low processing costs.

Within its broad approach to linguistic interpretation, RT makes room for two specific forms of pragmatic expansion: EXPLICATURE, a type of expansion

aimed at the development of the logical form of the utterance (“An assumption communicated by an utterance U is *explicit* if and only if it is a development of a logical form encoded by U”), and *IMPLICATURE*, defined as any enrichment that is not explicature (“Any assumption communicated, but not explicitly so, is implicitly communicated: it is an *implicature*”) [Sperber and Wilson, 1986, p. 182].¹⁵ The fact that explicatures constitute developments of logical forms entails what has become a canonical diagnostic used to distinguish explicature from implicature—the capacity for explicature to affect, or “intrude on,” truth value. For instance, RTists argue that explicatures lead to inclusion of the bracketed material in such examples as (12a–b, repeated as 18a–b), and that this happens *prior* to semantic evaluation, as a step on the way to fixing the truth-conditional contents expressed by utterances.

- (18) a. Donald is too crazy. [to be a serious contender for President]
 b. I haven’t had breakfast. [today]

Thus, Carston [2002] suggests that the pragmatic strengthening resulting in the conveyance of the bracketed material in (19a–b) must be treated as explicature rather than implicature (i.e., as making a contribution prior to semantic evaluation), on pain of our being unable to account for the possible truth of (19c) (cf. Cohen [1971, p. 58]; Wilson [1975, p. 151]).¹⁶

- (19) a. Annie got married and had a baby. [in that order]
 b. Annie had a baby and got married. [in that order]
 c. If Annie got married and had a baby, her parents would be happy, but if she had a baby and got married, they’d be very unhappy.

This behavior is inconsistent with implicatures in RT, which by definition are propositions that are inferred in addition to, and hence do not affect the truth value of, the proposition denoted by the linguistically coded material and its explicatures.¹⁷

¹⁵The term ‘explicature’ is used variably in the literature to refer either to the entirety of what is explicitly communicated by an utterance, or only the components of this meaning that result from a development of the logical form that the utterance encodes; we follow the latter usage here.

¹⁶Here we follow Carston (and the consensus) in treating the conveyed/bracketed material in (19a–b) (*mutatis mutandis*, the conveyed information at issue in our example (1a)) as being supplied extrasemantically rather than semantically (and so as potential instances of pragmatic intrusion into semantics). Of course, as King and Stanley [2005] point out, that treatment presupposes that there is no adequate semantic explanation of the origin of the relevant information; and they are surely correct that, at least in principle, this presupposition is open to question. We believe the standard treatment is defensible in these cases, e.g. because the information appears to pass standard tests for extrasemantic status such as being cancelable and reinforceable.

¹⁷It should be stressed that whereas both Grice’s account and RT distinguish implicatures from the propositions that are semantically encoded by utterances, RT’s division between

It is worth noting that, partly because it has the aim of accounting for linguistic understanding quite generally speaking, RT's approach to the topic of extrasemantic expansion differs in important respects from the approaches we have considered so far—respects to which we are largely sympathetic (and that are shared significantly by our own approach to eliciture; cf. §3). For one thing, and unlike the views discussed in §§2.1–2.3, RT does not treat enrichments as special phenomena whose interpretation is triggered by communicative failure. On the contrary, for RT, *any* utterance is incomplete and hence in need of enrichment if it is to satisfy the presumption of relevance [cf. the comprehension procedure for interpreting heard utterances outlined by Wilson and Sperber, 2004, p. 261ff]. The view therefore construes enrichments as natural by-products of normal interpretative processes. For another, proponents of RT view those normal interpretive processes as instances of general cognitive capacities that happen to be directed on linguistic inputs, rather than applications of language-specific procedures. (This commitment falls out of their Cognitive Principle of Relevance: “Human cognition tends to be geared to the maximisation of relevance”; p. 254.) Moreover, they hold that utterance-construction is guided by the aim of deliberately exploiting comprehenders' not-specifically-linguistic inferential/interpretive capacities, so that the latter not only direct interpretation but feed utterance planning as well (“I may be able to produce a stimulus which is likely to attract your attention, to prompt the retrieval of certain contextual assumptions and to point you towards an intended conclusion”; p. 254).

Despite the areas of agreement between our account and RT, we believe that elicitures are problematic for RT in at least two significant respects. The first is that elicitures are not classifiable into either of the two types of enrichment that the theory posits. On the one hand, elicitures appear to exhibit the kind of pragmatic intrusion that, as we saw, is compatible only with explicature. To see this, consider (20), spoken by a union leader at the union's monthly membership meeting:

- (20) If the company fires an employee who comes in late, a union complaint will be lodged.

It seems that this sentence could be judged true even under a scenario in which an employee came in late one day, was promptly fired for having embezzled money, and yet the union did nothing. That is, the fact that he wasn't fired *because* he was late causes the event to not satisfy the antecedent of the conditional in (20). As we have seen, the fact that eliciture exhibits this sort of intrusion on truth-conditional meaning is incompatible with treating it as a form of implicature within RT.¹⁸

explicature and implicature fails to line up with Grice's distinction between what is said and what is implicated. Specifically, a variety of enrichments that Grice would characterize as implicatures are argued to be explicatures in RT (e.g., (19)).

¹⁸Note that this point applies to both STRONG implicature and WEAK implicature in RT. Wilson and Sperber [2004] say:

A proposition ... is STRONGLY IMPLICATED ... if its recovery is essential in order to arrive at an interpretation that satisfies the expectations of relevance

On the other hand, eliciture lacks a defining characteristic of explicature as well. As we have discussed, the core of RT's distinction between explicature and implicature lies in the fact that the former but not the latter are developments (completions/expansions) of the logical form of an utterance. However, as we observed in §2.2, this is not true of eliciture, which often results in inferences to additional propositions entirely external to the logical form of the utterance itself. It would appear, then, that elicitures fit neatly into neither of the two categories of pragmatic expansion allowed by RT.

A second and more general problem that elicitures present for RT has to do with the nature of the comprehension procedure posited by the view (17). One of our central arguments against invoking Grice's notion of relevance as a source of explanation for our cases (see §2.1) applies equally here, specifically with respect to examples like (8a–c): given that relative clauses are deemed relevant if they do nothing more than restrict the domain of reference for the NP they modify (8b), we see no explanation for why inference would go any further in cases like (8a), in light of the hearer's instruction to stop upon the meeting of relevance expectations per (17b). After all, the vast majority of restrictive relative clauses found in discourse serve only this basic identificational function; those that convey an eliciture (which, crucially, do so *in addition to*, and not instead of, serving the basic function) are atypical. It is therefore hard to see how an interpreter's expectations of relevance would not be met by establishing that an RC felicitously serves the basic function alone.¹⁹

raised by the utterance itself. It is WEAKLY IMPLICATED if its recovery helps with the construction of an interpretation that is relevant in the expected way, but is not itself essential because the utterance suggests a range of similar possible implicatures, any one of which would do [Wilson and Sperber, 2004, p. 269].

Examples of weak implicatures include metaphors —e.g., *John has a square mind*— which convey meaning beyond their (generally false) literal meanings, but which are otherwise less determinate than typical (strong) implicatures. Whereas elicitures are more similar in their behavior to strong implicatures —the candidate enrichments are clear and determinate— neither type can allow for the sort of intrusion into truth conditional content witnessed in (20).

¹⁹We will argue in §3.2 that the inferential processes that endow hearers with the ability to draw elicitures in cases like (1a) are precisely those that hearers utilize to establish the coherence of intersentential discourses such as (21).

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

But there is an important difference between the inter- and intra-sentential cases that amounts to a further difficulty for RT: in the intersentential case, the hearer *must* draw some sort of relevancy relation ('coherence relation') between the clauses, or else the sentences will not constitute a coherent discourse. We believe that RT's characterization of the interpreter's need for linguistic material to meet his expectation of relevance aptly characterizes the interpretation of intersentential cases (as witnessed by the fact that interpreters will, within limits and according to certain principles, infer additional content needed to meet that expectation). However, as we will see, this is precisely the quality that elicitures lack (see

One might respond to this objection by proposing that perhaps elicitures result from an automatic/low-cost inferential process that, per the maximization of relevance, is outweighed by the added cognitive effect associated with it. After all, as Sperber notes,

The claim that the human cognitive system tends to allocate resources to the processing of available inputs according to their expected relevance is at the basis of relevance theory [Sperber, 2005, pp. 67–68].

However, as Levinson [1987, 1989] and others have pointed out, the problem with this claim is that the processor can have no way of knowing what the added cognitive effects (if any) would be of a particular enrichment without doing the work to calculate the candidate enrichment beforehand:

This would clearly involve comparison of different interpretations; but obviously to compare interpretations would involve computing each of them, thus additively increasing C[osts], so R[elevance] predicts it will be cheaper to pick any one interpretation at random! (Sperber and Wilson are aware of this difficulty, and suggest heuristics must exist that allow estimation of costs without computation (131); and so they must, and must be described and shown to work, if the theory of R is to survive) [Levinson, 1989].

As Levinson [1987] notes, however, RT offers no detailed hypotheses regarding such heuristics; but without such concrete hypotheses, the view's explanations for enrichments will remain objectionably *post hoc*.

Even setting this point aside, it seems doubtful that such heuristics, regardless of the form they take, would naturally cover the cases of eliciture addressed here. As we have noted, only a small minority of RCs found in natural language usage appear to support the sort of causal eliciture we see in (1a); certainly cases like (1b) are by far the typical case. Given this, it is hard to see how a heuristic would recommend expending the effort to attempt the inference of an eliciture, when very few instances will pay back the effort expended. We therefore take the theory to predict, on the basis of the relevance-theoretic comprehension principle, that interpretation will stop without attempting to seek enrichments of the sort witnessed in (1a).²⁰

the discussion of examples (27) and (28) in §3.2). It is unclear to us how this important distinction between the intersentential and intrasentential cases could be captured in RT, as one would expect the principles of the theory to apply with equal force in the two scenarios.

²⁰Another source of doubt arises from Sperber and Wilson's characterization of a positive cognitive effect:

When is an input relevant? Intuitively, an input (a sight, a sound, an utterance, a memory) is relevant to an individual when it connects with background information he has available to yield conclusions that matter to him: say, by answering a question he had in mind, improving his knowledge on a certain topic, settling a doubt, confirming a suspicion, or correcting a mistaken

We take the arguments of this section to establish that an adequate explanation of eliciture within RT is just as unlikely as an account framed in terms of the categories of the broad Gricean/neo-Gricean tradition we considered in §§2.1–2.3. Having made these arguments in some detail above, we are tempted by the following (necessarily more programmatic) diagnosis: RT struggles with eliciture because it attempts to explain this and all other forms of extrasemantic expansion in natural language, and the full array of inferences that result, as arising from a single explanatory source (viz., the mandate to maximize relevance). As we have taken pains in this paper to stress, the properties of the various enrichments we have surveyed differ along various dimensions, and are heterogeneous in their character. The history of linguistic research has taught us that language is highly complex at all levels of interpretation, and there is little reason to expect that the landscape of pragmatics would be significantly simpler. (We'll return to this theme in §4.) In our view, therefore, the argumentative burden rests firmly with those who would posit a uniform principle governing the establishment of such a wide range of inferences to extrasemantic content. In what follows, we offer our account of the cognitive apparatus that underlies the conveyance of eliciture. We intend for eliciture to take its place alongside the other types of extrasemantic expansion discussed in §2.1–2.3, and not as providing a new overarching principle intended to cover the breadth of phenomena addressed by RT.

3 Eliciture and inference

We have argued that eliciture resists reduction to other mechanisms of extrasemantic expansion. In this section, we'd like to suggest an alternative picture of eliciture and the kinds of extrasemantic expansion involved. We'll claim that this picture is in some respects more minimal than standard accounts of other types of extrasemantic enrichment, and connect it with an understanding of the cognitive strategies underpinning discourse processing.

impression. In relevance-theoretic terms, an input is relevant to an individual when its processing in a context of available assumptions yields a POSITIVE COGNITIVE EFFECT [Wilson and Sperber, 2004, p. 251].

As far as we can tell, it is not true that a hearer's propensity to draw an eliciture depends on his own goals and needs. On the contrary, a hearer who is utterly unconcerned with the reason for the manager's firing is as likely to draw the inference as a much more invested hearer. Indeed, as we noted in fn. 2, even experimental participants—who are completely detached from the situations conveyed by the experimental stimuli—are prone to draw elicitures. This is expected on our hypothesis that relevance-establishing machinery naturally serves up the inference regardless of the goals and needs of a particular hearer (see §3).

3.1 Linguistic enrichment and cognitive enrichment

On our picture, elicitures result from a speaker constructing her utterances so as to exploit a particular set of her hearer's general cognitive (not specifically linguistic) strategies for making sense of the world.

It is a familiar thought—and one that is very plausibly independent of any details of our linguistic capacities in particular—that cognitive agents make ampliative inferences that take them to conclusions beyond what is immediately justified by perception or other forms of direct evidence. Among other things, they also draw inferences that their observations invite, and incorporate the conclusions of these inferences into their evolving world models [cf., *inter alia*, [Rauschenberger and Yantis, 2001](#), [Hubbard, 1995](#), [Strickland and Keil, 2011](#)].

To see this point, consider an utterly unremarkable situation in which someone sees a chronically tardy department store employee show up late for work again, and soon thereafter witnesses the employee being fired. A reasonable cognitive agent might infer that the firing was due to the lateness, even if she had no first-hand evidence. Now consider a situation in which our department store employee shows up late for work again, and our agent subsequently sees a customer asking the employee where the automotive department is. In this case, our cognitive agent is unlikely to infer any relationship between the customer's question and the employee's lateness. Why? Because our world knowledge does not support a causal connection between the events. The lesson appears to be that when an inference of this sort suggests itself as in the firing scenario, a thinker will likely draw it, at least provisionally. However, the world remains perfectly coherent when no such inference between eventualities presents itself, as in the customer question scenario.

Our suggestion is that the very same strategies for making sense of the non-linguistically presented world should show up in the course of interpreting discourses that describe such situations linguistically. That is, when a hearer interprets (22),

(22) The boss fired the employee who came in late again.

he will reasonably associate the firing with the lateness in the same way in which they are associated in the (non-linguistic) case considered above, in which the events were perceived. Likewise, our addressee will presumably not draw a causal inference for the variant in (23):

(23) A customer asked the employee who came in late again where the automotive department is.

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

In some respects, then, the inferences underpinning elicitures are of the most pedestrian sort. They are simply the inferences cognizers make in order to organize and bring coherence to their evolving world models.

That said, and notwithstanding our arguments in §2, we do not wish to claim that elicitures have nothing in common with Gricean implicatures or inferences grounded in other types of pragmatic reasoning. Quite the contrary. The common property that these inferences share is that they result from the hearer asking not only *what* the speaker said, but also *why*. In the case of certain Gricean implicatures, the inference results from asking why the speaker made a particular statement instead of something else that might have been possible. For instance, the possible implicatures associated with *Some students got an A*—that either it is not the case that *all* students got an A (quantity implicature), or that the speaker does not know whether all students got an A (ignorance implicature)—can be seen to result from the consideration of possible reasons why the speaker chose to say *some* instead of *all*. In the case of indirect speech acts, interpretation depends on understanding why a particular utterance was made in light of the speaker's communicative goals. That is, the addressee who recognizes the statement *It's cold in here* as an indirect request to close the window has done so by asking *why* the utterance was made by the speaker in light of (his estimation of) her beliefs, desires, and intentions, and not just *what* the utterance denotes.

Theorists have held that the inference procedures that give rise to such enrichments require an appeal to MUTUAL KNOWLEDGE: a speaker can only be said to be communicating extragrammatic content if the knowledge required to support the inference to that content is intended by the speaker to be taken into account by the addressee, if the addressee likewise recognizes that the speaker intends for it to be taken into account, and so on. Likewise, when interpreting (3a), it stands to reason that a hearer would ask why the speaker is describing the referent as a jogger, as opposed to other possible descriptions (*a man, a Stanford student, etc.*). A likely justification is that the speaker wishes to communicate that the victim was jogging at the time, which is a proposition that the hearer can readily accommodate as long as it is consistent with his beliefs about the world. Again, as with other types of pragmatic enrichment, mutual knowledge about plausible relationships (here, between jogging and accidents) is necessary for this reasoning to go through. But whereas the speaker, in choosing this form of reference, is counting on the hearer to ask why that particular form was chosen, she is not relying on the kinds of inference procedures that underlie Gricean implicatures or the plan-based reasoning about intentions upon which indirect speech act recognition is based. Instead, she relies on the fact that the addressee, upon hearing (3a), will draw a connection between the jogging and the accident because world knowledge suggests that they can be so related, and takes advantage of this fact to communicate her message in a particularly economical way.²¹

²¹The claim that intention recognition is crucial for the inference to elicitures predicts that, in certain contexts, a hearer might fail to draw an eliciture that would otherwise ordinarily be drawn. For instance, if in a particular situation the referring expression *the employee who is always late* is the only form that can achieve successful reference and this fact is clear to the conversational participants, the existence of an alternate explanation for the speaker's choice of RC may likely weaken the hearer's inclination to draw the eliciture. Another case,

3.2 Inference and coherence relations

The idea that eliciture is grounded in general cognitive mechanisms for understanding the world invites us to ask just what the relevant cognitive mechanisms are and how they operate. We propose that the mechanisms should be understood as a class of operations that connect individual mental states into a coherent whole by virtue of relations between their contents.

Of course, the idea that mentation proceeds in terms of links between contents is a familiar one in the history of thought about cognition, with roots stretching at least to Epicurus (*Diogenes Laertius* vii. §52, x. §32), Plato (the doctrine of anamnesis in the *Phaedo*), and Aristotle (On Memory and Recollection 2, 451b16–22), and coming to prominence with 18th century thinkers including Hobbes (*Human Nature*, iv, (4.15)), Berkeley (*New Theory of Vision*, §25), and Hartley (*Observations on Man*). However, perhaps the most influential articulation of the idea comes from Hume, who holds that ideas juxtaposed in thought are always connected by one of three types of associative links: “Resemblance, Contiguity in time or place, and Cause or Effect” (*Enquiry Concerning Human Understanding*, §III).

Though we do not wish to commit to a Humean or associationist picture of human psychology, we are sympathetic to the more general idea that cognition can be understood in terms of relations between the contents of thoughts entertained in a thinker’s temporally extended cognitive sequence. More specifically, we are attracted to the view that discourse understanding is often fruitfully explained in terms of hearers’ attempts to bring coherence to heard clauses by subsuming the latter under semantic relations—henceforth, COHERENCE RELATIONS.²² For instance, on its most natural interpretation, the sentences in example (24) are connected by an Explanation relation [Hobbs, 1990], whereby the state described in the second sentence is interpreted as a cause of or reason for the event expressed in the first sentence.

(24) John is traveling to Paris. He wants to visit his family.

Establishing an Explanation relation will typically require additional inferences; upon hearing (24), for instance, a hearer will normally infer that John’s family is or will be in Paris, and that John intends to visit his family during his trip. Even though the passage never asserts that John has this intention, inferring that he does is required in order to treat the proposition expressed by the second sentence as a reason for the proposition expressed by the first.

There is a clear parallel between these intersentential inferences and those arising in the examples we have discussed in this paper; the inference drawn

suggested by an anonymous reviewer, is where one friend often tells another stories about work colleagues, and the phrase *the employee who is always late* has, over time, become a conventionalized way of referring to one of them. We see no way that elicitures in such cases could be blocked without appeal to intention-based reasoning.

²²For more detail on this program, see Kehler [2002], which draws on Hobbs’s [1979, 1990] theory of discourse coherence, and taxonomizes coherence relations in terms that mirror Hume’s catalog of associative relations.

for (25), for instance, mirrors the one that will be drawn to establish the coherence of (26):²³

(25) John fired the employee who was always late.

(26) John fired the employee. He was always late.

As we briefly mentioned in fn. 19, the crucial difference between these interclausal inferences and the intraclausal ones is that there is a requirement that the hearer be able to draw *some* coherence relation among the clauses in (26) in order to justify their juxtaposition in the discourse. Specifically, as argued by many theories of discourse coherence [e.g., Hobbs, 1990, Kehler, 2002, Asher and Lascarides, 2003], coherent discourses require their clauses to be structurally related in much the same way that grammatical sentences require their syntactic constituents to be; discourses thus admit of hierarchical structures in the way that sentences do. The fundamental difference is that whereas the constraints on when adjacent constituents can combine to form larger ones in building a sentence structure are dictated by grammatical rules, in the realm of discourse ‘grammar’ it is the ability to infer coherence relations that determines when clauses (or collections thereof at higher levels of structure) can be combined. As such, it is not enough for addressees to merely interpret the two clauses in (26) as independent statements about John [Hobbs, 1979]. In contrast, the fact there is no coherence criterion that applies among constituents within a clause means that the same inference for (25) is merely invited; hence, felicity is not at stake at this level.

This characterization predicts that we will find cases in which multiple predications about an entity will be infelicitous when expressed as a discourse yet felicitous when expressed within a sentence. To see that this prediction is borne out, consider (27):

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

Outside of an exceptional context, (27) lacks coherence.²⁴ In contrast, a version that expresses the same content intraclausally is fine:

(28) The employee who likes plums broke his leg.

Whereas the inclusion of the plum-liking in a relative clause in (28) presupposes that it will help the hearer identify the referent, no further coherence-driven inferences are required at the sentence level, in contrast to (27).

²³See Hobbs [2010] and Pagin [2014] for similar observations. These works will be discussed briefly in fn. 28.

²⁴The reader might object by offering a context that would make (27) coherent, e.g., one in which the employee made an unsuccessful attempt to climb a plum tree. Indeed, it appears that it is always possible to come up with a context that will make any pair of seemingly unrelated sentences coherent. But as Hobbs [1979] points out, the very fact that addressees are driven to identify such contexts shows that establishing coherence between sentences is a necessary component of discourse interpretation.

Another coherence relation that we already have seen is Denial, in which an expected cause-effect relation is denied, as in (29):

(29) Snodgrass is honest, even though he's a politician.

Importantly, Denial presupposes the very cause-effect relation that is being denied in the particular instance being described; here the addressee would rightfully assume that the speaker believes that politicians are normally dishonest. (The addressee could even felicitously reject this presupposition —*Hey, I think most politicians are honest!*— even though nothing in what (29) asserts contradicts that sentiment.) And, as expected, we see the same pattern between intrasentential and intersentential cases:

- (30) a. The company fired the manager who has a long history of corporate awards.
b. The company fired the manager (even though) he has a long history of corporate awards.
- (31) a. The company fired the manager who drives a blue Lexus.
b. # The company fired the manager (even though) he drives a blue Lexus.

Passages (30a-b) both give rise to the counter-to-expectation inference between getting fired and having a long history of corporate awards. Passages (31a-b) again illustrate the fact that such inferences are unnecessary at the intrasentential level: whereas (31a) is impeccable, (31b) leaves the addressee searching for a coherence relation (causal or otherwise) between the two clauses.

Example (3a), repeated below as (32), is among the cases we have considered that are not causal in nature:

(32) A jogger was hit by a car in Fresno last night.

Here the jogging is not inferred to have caused the accident, but only as what the victim was doing at the time the accident occurred. Such relationships are instances of the Occasion coherence relation [Hobbs, 1990]. Occasion relations characterize spatiotemporally related series of events that are connected by intermediate states of affairs in the world. Consider:

(33) The employee went to the store. He bought a bottle of scotch for the office party.

Although there is no causality at play here, an inference is necessary to establish Occasion on the most natural interpretation of (33), specifically that the scotch was bought at the store.²⁵ We can compare (33) to its intrasentential variant in (34):

²⁵It is of course possible to construct a context that favors a different coherence construal for passages like (33). For instance, if (33) was given as an answer to the question *What things did the employee do today?*, it is possible to interpret the two clauses as independent events. The operative coherence relation in this case would be Parallel.

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

Again, passage (34) invites the same inference as (33). However, unlike (33), (34) remains felicitous if the inference is not drawn. As such, we again see a contrast in acceptability between variants in which Occasion is not supported:

- (35) a. The employee who went to UC San Diego for grad school bought a bottle of scotch for the office party.
b. # The employee went to UC San Diego for grad school. He bought a bottle of scotch for the office party.

Our proposal, then, is that eliciture is grounded in the same cognitive mechanisms that are already needed to explain the establishment of coherence relationships across utterances in a discourse. The main difference is that whereas identifying *some* relationship is required when establishing relevance across clauses, there is no mandate to do so between constituent meanings within a clause.

4 Taking stock: Eliciture, coherence, and pragmatic theory

Having presented an account of eliciture in terms of intrasentential coherence in §3, we are now in a position to take stock, ask what we can and cannot show, and draw lessons for pragmatic theory.

4.1 Intersentential coherence and pragmatic theory

One such lesson concerns the relationship between *intersentential* coherence and the familiar categories of pragmatic enrichment discussed in §2. We have argued (§2) that eliciture is not assimilable to these familiar categories. Moreover, we have argued (§3.2) that the inferential mechanism responsible for drawing elicitures is the same one that underlies the process of establishing intersentential coherence. Taking these thoughts together would seem to cast similarly strong doubt on the possibility of understanding cases involving intersentential coherence in terms of the familiar categories. This is notable because there is a long tradition in pragmatics of proposals that take this form.

For example, consider Levinson's attempt to explain as instances of implicature the class of cases he dubs "conjunction buttressing," as in Geis's and Zwicky's (14a), repeated below as (36):

- (36) Martha observed the children at play and smiled with pleasure.

As we pointed out in fn. 13, for Levinson the strengthening of the meaning of *and* to *and as a result* (36b) is an example of an I-implicature. However, as Levinson himself (pp. 122–127) and others note, the same enrichments we see in (36) can occur without the conjunction:

(37) Martha observed the children at play. She smiled with pleasure.

There cannot be “conjunction buttressing” if there is no conjunction to buttress. It is true that *and* does serve a function relevant to coherence establishment, in that it is only compatible with certain coherence relations (it disallows Explanation, for instance—the sentence *Fred slipped and he stepped on a plum* cannot mean that Fred slipped because he stepped on a plum). But 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.

Such strengthenings have been treated as examples of implicature as well. Specifically, Bach [2010] includes cases like (36) among those involving implicature—particularly as expansions—on analogy with (12a), repeated below as (38).

(38) I haven’t had breakfast. [today]

However, it is hard to see how Bach’s characterization of expansion —“the fleshing out of the minimal proposition expressible by an utterance” as necessarily triggered by a hearer recognizing “that a speaker cannot be plausibly taken, and therefore does not intend to be taken, to mean what he is saying”— would apply to many examples. Consider:

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

Note that there are two salient construals for (39). The first is an (unenriched) Parallel construal: Trump and his Vice-President were each overcome with negative emotions, possibly (but not necessarily) due to the same external stimulus. The second is an (enriched) Result construal: Trump, being unsympathetic to Pence’s show of emotion, threw a tantrum as a result. But if the second reading is only triggered by recognizing “that a speaker cannot be plausibly taken” to intend to express the unenriched construal, we have no explanation for why that construal is not only a plausible, but indeed salient, interpretation for the utterance.²⁶

The crucial point is that cases that involve the apparent strengthening of a conjunction meaning represent only a narrow subset of the full range of inferences that result from the establishment of discourse coherence—inferences that generally go well beyond what can be characterized as developments to the logical form of an utterance, as we can see from reflecting on (40a–c):

- (40) a. John is traveling to Paris. He wants to visit his family. (=24)
(Explanation)
- b. The employee went to the store. He bought a bottle of scotch for the office party. (=33) (Occasion)
- c. Country singers sometimes sell more albums than your typical pop star. Taylor Swift’s latest album sold 9 million copies. (Exemplification)

²⁶The more recent analysis of coherence and pragmatic enrichment of Pagin [2014], discussed briefly in fn. 28, suffers from the opposite problem: it predicts that such cases will always be strengthened to a causal interpretation.

As we noted in §3.2, there is more to establishing coherence in (40a) than the inference that John is traveling to Paris *because* he wants to see his family. Here the addressee will not be content to merely record the fact that *some* causal relationship holds. Additional inferences are necessary to support this causal interpretation, such as that his family is or will be in Paris, and that John intends to see them during this particular trip. Needless to say, upon hearing (40a) an addressee would be quite surprised to find out that John had no intention of seeing his family during his trip, even though the passage never asserts that he does so intend. Similarly, example (40b) supports an Occasion relation. Here the addressee will not merely infer temporal progression among the events, but also infer that the scotch was purchased at the store. An addressee presumably makes this enrichment effortlessly, and perhaps without even noticing that the enriched content is not stated explicitly. Finally, (40c) is an example of Exemplification. Let us suppose the addressee doesn't know much about the music industry nor who Taylor Swift is. Assuming that the addressee recognizes and accepts the speaker's intention to trigger the recognition of Exemplification, reading this passage will then teach him a few things that are not explicitly stated: that Taylor Swift is a country singer and that 9 million copies of an album is more than what a typical pop star would sell. These inferences are required if the proposition expressed by the second sentence is to be understood as providing an example of the proposition expressed by the first. All of the inferences discussed for these three examples are cases of pragmatic enrichment, and all have coherence establishment as their source. Specifically, they are inferences drawn in order to meet constraints imposed by the operative coherence relation, and hence go well beyond inferring that a particular conjunction should be given a 'strengthened' —e.g., temporal or causal— interpretation.

If, as we believe, these considerations suggest that attempts to understand intersentential coherence in terms of the inventory of familiar pragmatic tools are unlikely to succeed, this failure is unsurprising given the arguments of this paper. For if the inferential mechanisms at work in the interpretation of eliciture rest in processes of coherence establishment operative in both intersentential and intrasentential settings, then it is predictable that cases involving intersentential coherence should be just as resistant to explanation in terms of these familiar tools as cases of intrasentential coherence (*viz.*, eliciture).

4.2 Conversational Eliciture: Advances and Limitations

We close with an assessment of what we have and have not accomplished in our discussion of eliciture. We began the paper with the goal of describing a process by which extrasemantic content is conveyed that fails to fit neatly into any of the types of pragmatic enrichment thus far described in the literature. We accomplished this goal through an analysis of the properties of eliciture with respect to four questions with which theories of enrichment have historically engaged (§1). First, in §§1-2 we provided a characterization of the type of enrichment that results from the inference of an eliciture, emphasizing its non-local character and distinguishing it from other, more familiar, types of enrichment. Second, in §3 we identified those aspects of

the hearer’s cognitive apparatus that a speaker exploits when constructing her utterance so as to convey eliciture, specifically those capacities responsible for establishing the coherence of intersentential discourse and of the world more generally. Third, in §3 we characterized the nature of the inference processes associated with that machinery for establishing coherence, that is, through the satisfaction of constraints that are associated with particular types of coherence relations. Finally, in §2.1 we argued that unlike many other types of enrichment, the inference to eliciture is not triggered by any violation of communicative norms or other felicity requirement on linguistic material, but instead results from more basic and constantly operative cognitive machinery that serves up candidate inferences. This result leaves us with a heterogeneous view of pragmatic enrichment, in which eliciture takes its own place among a broad inventory of enrichment mechanisms that speakers can rely on to make their linguistic contributions efficient while still remaining expressive.

Having accomplished this, we want to be clear that there are limitations to our account as well which, in turn, plausibly reflect principled limitations on the possibility of systematic theorizing in pragmatics more generally. Specifically, as with the other accounts of various types of enrichment surveyed in §2, the account of §3 does not provide specific predictions about arbitrary cases and hence cannot be regarded as a total, closed-form, predictive theory of the extrasemantic enrichments it targets. But this shouldn’t be too surprising. Indeed, lessons from the theory of coherence establishment across sentences suggest that an adequate total theory of extrasemantic expansion would have to be sensitive to an apparently unlimited range of facts that could turn out to be relevant to making the right predictions about cases.²⁷

It would need to predict, for example, that (41) doesn’t give rise to a causal eliciture in ordinary contexts, but that it will in a context in which John is a protective parent with a Bieber-fan daughter who also works for his company. It would need to predict that, though an utterance of (42) strongly invites the inference that the runner was running at the time of the accident if the conversants take the victim to be an occasional/hobbyist runner, the invitation to that inference is much weaker if they take the victim to be a professional runner—or, alternatively, in a case where the string is uttered by one runner to another at a meeting of their running club.

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

(42) A runner was hit by a car in Fresno last night.

And so on. Of course, our point is not that it is hard to imagine a proposal that takes the specific considerations mentioned into account—particularly after they have been highlighted. It is, rather, that it is hard to imagine a general proposal

²⁷As Sperber and Wilson [1986] and others have noted, much the same is likely to be true of the cases Grice would classify as particularized implicatures. Hence eliciture is far from alone in this respect.

—one that extends systematically to all cases— that takes into account (in advance) whatever specific considerations turn out to be relevant to whether or not such extrasemantic inferences go through.

What our analysis does do is explain an interesting class of expansions *relative to the establishment of coherence*, which distinguishes it from the other types of enrichment surveyed in §2. For example, our account is able to predict that the causal eliciture will be drawn from (1a) if the hearer is able to construe the components of that utterance as standing in an Explanation relation, and that an occasion eliciture will be drawn from (3a) if he is able to construe the components of that utterance as standing in an Occasion relation. Indeed, it allows similarly conditional predictions about the hard cases raised above—viz., that a causal eliciture is drawn from (41) given appropriate (presupposition-introducing, hence) Explanation-supporting stipulations about John’s protectiveness toward his Bieber fan daughter who works at his company, and that an Occasion eliciture is not drawn from (42) given stipulations about the context that block or remove from the common ground the general presuppositions required for the establishment of Occasion. A direct consequence of this dependence of eliciture on the framework for establishing coherence relations is that any account of the latter should also advance considerably our understanding of the former. As it stands, however, because there is no complete theory of coherence establishment on offer, and our analysis proposes to understand eliciture in terms of coherence establishment, there is necessarily no complete theory to offer for eliciture either.²⁸

²⁸ Some theorists who have addressed pragmatic enrichment via coherence establishment have taken a more optimistic view of the prospects for a complete theory of this sort.

One such account is the INTERPRETATION AS ABDUCTION framework of Hobbs et al. [1993], according to which interpreting sentences in discourse is viewed as the process of inference to the best explanation of why the sentences would be true. Although the best explanations will prove as much information deductively as possible, typically the proof procedure will require the assumption of information that cannot be established deductively from context and world knowledge. These assumptions yield pragmatic enrichments, such as those that result from establishing coherence. Of course, the space of possible analyses will be quite large (indeed, presumably infinite); hence Hobbs [2010] points to a set of theoretical virtues as a metric for determining which explanation is best. Thus, he says, cognition prefers short proofs over long ones, salient axioms over nonsalient ones, minimal sets of assumptions, and exploitation of implicit redundancies. But none of these virtues are well-understood, and certainly none are specified to the degree necessary to make predictions about specific cases. Further, all are implementation-dependent; what proofs are short or long, how many assumptions will be required for those proofs, how much particular assumptions ‘cost’ to make, and so forth, will depend on the design decisions made during the logical encoding of the knowledge base.

Similar concerns apply to Pagin’s (2014) analysis of coherence establishment and pragmatic enrichment. Pagin defines five categories of coherence that are ordered by increasing strength: VACUITY (no relation), CONTIGUITY (connected in time and/or space), RESEMBLANCE (parallelism), POSSIBILITY (influence, e.g., enablement), and NECESSITY (causality), with the idea being that free pragmatic enrichment occurs when it serves to raise the degree of coherence with respect to this ordering. The problem is that, since it is

5 Conclusion

The balance between expressivity and efficiency proposed by Zipf predicts that speakers may, when appropriate, design their utterances to take advantage of the hearer’s mental state and capacity for inference to communicate more than what they explicitly say. The variety of ways in which they accomplish this has been the focus of research on pragmatic enrichment for many years, starting with the seminal work on implicature of Grice [1975], and following numerous others that include Horn’s [1972] neo-Gricean revision to Grice’s system (which, famously, engages Zipf’s opposition directly), Bach’s implicature, various forms of local pragmatic strengthening proposed by Levinson, Recanati, and others, and of course Relevance Theory as an alternative to the Gricean picture. Various central questions—including the character of the particular enrichments at hand, the aspects of the hearer’s cognitive apparatus that a speaker exploits when constructing her utterance so as to convey the enrichment, the details of the inference processes used by the hearer to recover the enrichment, and the triggers responsible for the hearer’s initiation of those inferential processes—have been at the heart of the debate as the field moves toward a broader understanding of the ways in which extragrammatic content is conveyed.

We have argued for the existence of a previously unidentified type of enrichment—eliciture—that takes its place against this rich landscape. Whereas, like implicature, our analysis engages intention recognition at the highest level, we have taken pains to argue (§2) that eliciture differs from other types of enrichment in key respects. Further, we contend that consideration of eliciture, and the cognitive mechanisms we claim underpin its operation, highlights connections between language understanding and other mental processes (§3), and holds significant lessons about the prospects for systematic theories of extragrammatic expansion (§4). For all these reasons, we believe there is much to be gained from examining eliciture and consideration of its special properties.²⁹

always possible to establish coherence between clauses if enough contextual information is accommodated, a theory needs to tell us what the limits are. Pagin offers two principles intended to answer this need: PLAUSIBILITY and SIMPLICITY. We will not belabor the point here, but we doubt that defining these principles with enough precision to make predictions about arbitrary cases will prove any easier than it will for Hobbs’s theoretical virtues.

Whereas the many degrees of freedom in both these accounts no doubt leave one able to construct a story for particular examples post hoc, we see no reason to be optimistic that these stories can be motivated independently of the examples they purportedly explain, nor that the amalgamation of the analyses required for larger sets of examples would result in a consistent knowledge base. Indeed, lessons learned from the evolution of artificial intelligence research suggest just the opposite.

²⁹For discussions and comments that have greatly improved the paper we are grateful to Kent Bach, Betty Birner, Emma Borg, Ivano Caponigro, Michael Glanzberg, Jerry Hobbs, Larry Horn, Chris Kennedy, Max Kölbel, Line Mikkelsen, John MacFarlane, Peter Pagin, Paul Pietroski, Chris Potts, François Recanati, Martin Schäfer, Josef Stern, Isidora Stojanovic, Una Stojnic, Gregory Ward, and several anonymous referees; thanks also to

References

- Nicholas Asher and Alex Lascarides. *Logics of Conversation*. Cambridge University Press, Cambridge, 2003.
- Kent Bach. Conversational implicature. *Mind and Language*, 9(2):124–162, 1994.
- Kent Bach. The top 10 misconceptions about implicature. In *Drawing the Boundaries of Meaning: Neo-Gricean Studies in Pragmatics and Semantics in Honor of Laurence R. Horn*, pages 21–30. John Benjamins, Amsterdam/Philadelphia, 2006.
- Kent Bach. Implicature vs explicature: What’s the difference? In Belén Soria and Esther Romero, editors, *Explicit Communication: Robyn Carston’s Pragmatics*, pages 126–137. Palgrave Macmillan, Basingstoke, 2010.
- Alfonso Caramazza, Ellen Grober, Catherine Garvey, and Jack Yates. Comprehension of anaphoric pronouns. *Journal of Verbal Learning and Verbal Behaviour*, 16: 601–609, 1977.
- Robyn Carston. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Wiley-Blackwell, Oxford, 2002.
- L. Jonathan Cohen. Some remarks on Grice’s views about the logical particles of natural language. In Yehoshua Bar-Hillel, editor, *Pragmatics of natural languages*, pages 50–68. Reidel, Dordrecht, 1971.
- Alexander Dinges. Innocent implicatures. *Journal of Pragmatics*, 87:54–63, 2015.
- Catherine Garvey, Alfonso Caramazza, and Jack Yates. Factors underlying assignment of pronoun antecedents. *Cognition*, 3:227–243, 1976.
- Michael Geis and Arnold Zwicky. On invited inferences. *Linguistic Inquiry*, 2(4): 561–566, 1971.

the SemanticsBabble discussion group at the University of California, San Diego, and to audiences at the University of California, Merced; University of California, Davis; University of Chicago; Florida International University; Institut Jean Nicod; Rutgers University; Stanford University; University College London; and the University of Wisconsin, Madison; and to conference audiences at the Investigating Semantics conference (Bochum, 2013); the California Semantics and Pragmatics workshops (UC Berkeley, 2013; UC Santa Cruz, 2016); the Conceptual Structure, Discourse, and Language meeting (Santa Barbara, 2014); the Textlink Action Conference (Louvain-la-Neuve, 2015); the Cognition and Language Workshop (Stanford University, 2015); the Society for Philosophy and Psychology (Duke University, 2015; UC San Diego, 2019); the Semantics, Pragmatics and Language Philosophy Workshop on Semantics & Cognition (UC Santa Cruz, 2016), the Maryland Mayfest (University of Maryland, 2016); the Joint SEMdial and SIGdial Meeting on Discourse and Dialog (Saarland University, 2017); the Making Sense of Discourse conference (Utrecht University, 2018); and the At-Issue, Scope, and Coherence workshop (University of Cologne, 2018) who heard and provided valuable feedback on earlier versions of this material.

- H. Paul Grice. Logic and conversation. In Peter Cole and Jerry L. Morgan, editors, *Syntax and Semantics*, volume 3, pages 41–58. Academic Press, New York, 1975.
- H. Paul Grice. Further notes on logic and conversation. In Peter Cole, editor, *Syntax and Semantics: Pragmatics*, volume 7, pages 113–127. Academic Press, New York, 1978.
- Julia Hirschberg. *A Theory of Scalar Implicature*. Garland, New York, 1991.
- Jerry R. Hobbs. Coherence and coreference. *Cognitive Science*, 3:67–90, 1979.
- Jerry R. Hobbs. *Literature and Cognition*. CSLI Lecture Notes 21, Stanford, CA, 1990.
- Jerry R. Hobbs. Clause-internal coherence. In P. Kühnlein, Anton Benz, and Candace Sidner, editors, *Constraints in Discourse 2*, pages 15–34. John Benjamins, Amsterdam, 2010.
- Jerry R. Hobbs, Mark E. Stickel, Douglas E. Appelt, and Paul Martin. Interpretation as abduction. *Artificial Intelligence*, 63:69–142, 1993.
- Laurence R. Horn. *On the Semantic Properties of Logical Operators in English*. PhD thesis, UCLA, 1972.
- Laurence R. Horn. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicatures. In Deborah Schiffrin, editor, *Meaning, Form, and Use in Context*, pages 11–42. Georgetown University Press, 1984.
- Laurence R. Horn. Implicature. In Laurence R. Horn and Gregory Ward, editors, *Handbook of Pragmatics*. Basil Blackwell, 2004.
- Timothy L. Hubbard. Environmental invariants in the representation of motion: Implied dynamics and representational momentum, gravity, friction, and centripetal force. *Psychonomic Bulletin & Review*, 2(3):322–338, 1995.
- Andrew Kehler. *Coherence, Reference, and the Theory of Grammar*. CSLI Publications, 2002.
- Andrew Kehler and Hannah Rohde. Prominence and coherence in a Bayesian theory of pronoun interpretation. *Journal of Pragmatics*, 154:63–78, 2019.
- Andrew Kehler, Laura Kertz, Hannah Rohde, and Jeffrey L. Elman. Coherence and coreference revisited. *Journal of Semantics (Special Issue on Processing Meaning)*, 25(1):1–44, 2008.
- Jeffrey C. King and Jason Stanley. Semantics, pragmatics, and the role of semantic content. In Zoltan Szabo, editor, *Semantics Versus Pragmatics*, pages 111–164. Oxford University Press, 2005.

- Alistair Knott and Robert Dale. Using linguistic phenomena to motivate a set of coherence relations. *Discourse Processes*, 18(1):35–62, 1994.
- Amichai Kronfeld. *Reference and Computation: An Essay in Applied Philosophy of Language*. Cambridge University Press, Cambridge, 1990.
- Stephen C. Levinson. Implicature explicated? *Behavioral and Brain Sciences*, 10: 722–723, 1987.
- Stephen C. Levinson. A review of Relevance. *Journal of Linguistics*, 25:455–472, 1989.
- Stephen C. Levinson. *Presumptive Meanings*. MIT Press, Cambridge, MA, 2000.
- Peter Pagin. Pragmatic enrichment as coherence raising. *Philosophical Studies*, 168: 59–100, 2014.
- Robert Rauschenberger and Steven Yantis. Masking unveils pre-amodal completion representation in visual search. *Nature*, 410(6826):369–372, 2001.
- François Recanati. *Direct Reference: From Language to Thought*. Blackwell, Cambridge, Massachusetts, 1993.
- François Recanati. *Literal Meaning*. Cambridge University Press, Cambridge, 2004.
- François Recanati. *Truth-Conditional Pragmatics*. Oxford University Press, Oxford, 2010.
- Hannah Rohde, Roger Levy, and Andrew Kehler. Anticipating explanations in relative clause processing. *Cognition*, 118:339–358, 2011.
- David Rumelhart. Some problems with the notion of literal meanings. In A. Ortony, editor, *Metaphor and Thought*, pages 71–82. Cambridge University Press, Cambridge, 1979.
- Jerrold M. Sadock. On testing for conversational implicature. In Peter Cole, editor, *Pragmatics*, pages 281–297. Academic Press, New York, 1978.
- Jennifer M. Saul. What is said and psychological reality: Grice’s project and relevance theorists’ criticisms. *Linguistics and Philosophy*, 25(3):347–372, 2002. doi: 10.1023/A:1015221313887.
- Dan Sperber and Deirdre Wilson. *Relevance: Communication and Cognition*. Harvard University Press, Cambridge, MA, 1986.
- Daniel Sperber. Modularity and relevance. In Peter Carruthers, Stephen Laurence, and Stephen Stich, editors, *The innate mind: Structure and contents: Volume 1*, pages 53–68. Oxford University Press, Oxford, 2005.

- Brent Strickland and Frank Keil. Event completion: Event based inferences distort memory in a matter of seconds. *Cognition*, 121(3), 2011.
- Patrick Suppes. Semantics of context-free fragments of natural languages. In Jaakko Hintikka, Patrick Suppes, and J.M.E. Moravcsik, editors, *Approaches to Natural Language*, pages 370–394. Reidel, Dordrecht, 1973.
- Bonnie Lynn Webber. Discourse modelling: Life at the bottom. In *AAAI Fall Symposium Series on Discourse Structure in Natural Language Understanding and Generation*, pages 146–151, Asilomar, CA, 1991. American Association for Artificial Intelligence.
- Deirdre Wilson. *Presupposition and Non-Truth-Conditional Semantics*. Academic Press, New York, 1975.
- Deirdre Wilson and Dan Sperber. Relevance theory. In Laurence R. Horn and Gregory Ward, editors, *Handbook of Pragmatics*. Basil Blackwell, 2004.
- George Kingsley Zipf. *Human Behavior and the Principle of Least Effort*. Addison Wesley, Cambridge, 1949.