

Pragmatics and the Lexicon

Florian Schwarz & Jérémy Zehr

June 2020

1 Introduction

Pragmatics, as the study of meaning in language use, is concerned with the ways that contextual information affects the overall conveyed meaning of an utterance. For example, whether (1) is uttered out of the blue or in response to a preceding claim that Nina swims influences whether one is more or less likely to interpret it as communicating that Lola does *not* swim.

(1) Lola runs.

While to some extent the term *run* itself may be seem to be associated with an inherent opposition to *swim*, this opposition clearly is not part of its core meaning—after all, one can perfectly well claim that Lola runs and swims without contradiction. Furthermore, as we just saw, the presence and salience of this opposition is clearly modulated by context (here, a preceding utterance—or its absence).

The present chapter reviews how contextually driven inferences, broadly construed, interact with the encoding of meaning at the lexical level. Our starting point is a fairly standard view of divvying up meaning into semantic and pragmatic components: semantics is primarily concerned with what we alluded to as “core” meaning above, while pragmatics is concerned with contextual reasoning. In light of this view on the semantics-pragmatics interface, a central research question is whether (and to what extent) a given ingredient of meaning associated with the use of a particular expression in context should be seen as directly encoded in the lexicon, or whether (and to what extent) it is derived via general reasoning in context.

Perhaps unsurprisingly, the answer to this question is oftentimes far from clear-cut, and the details of just which bits are conventionally encoded and which are pragmatically inferred are delicate and controversial. Our discussion will focus on key phenomena that serve as case studies to illustrate the task of spelling out the division of labor between semantics and pragmatics in full detail. First, we turn to phenomena that crucially are construed as relating to scales, namely scalar implicatures and scalar adjectives. Underlying the notion of scale is the generic process of comparison; as we will see in our discussion of these two phenomena, context and lexical encoding both fundamentally contribute to the identification of the objects being compared, and to the structures they define in that process.

The next case study turns to presuppositions, which introduce another set of issues with regards to the source of ingredients of meaning. On the

one hand, presuppositions are inherently contextual, as they integrate the conversation's participants' beliefs in the discourse; besides, while traditional theoretical accounts of meaning model with a certain accuracy how conventionally encoded meaning is computed, they fail to model characteristic behaviors of presuppositions. On the other hand, presuppositions seem to be directly tied with particular lexical items, suggesting a conventional association. Presuppositions thus constitute a perfect example of the interaction between context and conventionally encoded content at the lexical level. But before turning to these detailed case studies, the following section reviews the theoretical background and assumptions that form the basis for the underlying view of the semantics-pragmatics interface in more detail.

2 Background: Constructing meanings

A core notion that is the basis of our discussion is that the overall meaning intuitively conveyed by uttering a given expression is a conglomerate of distinct ingredients that can be separated out based on their properties, and which may have different sources. As sketched for (1) above, some aspects of meaning are thought to be conventionally associated with particular lexical items (e.g., that Lola runs), while others (e.g., that Lola doesn't swim) can enter the picture via contextual reasoning that enriches the conventionally encoded meaning. Putting aside further complications for a moment

(of which there are many), the former constitute the realm of semantics, whereas the latter fall into the purview of pragmatics.

Let us explicate the key terms here in a bit more detail: conventional encoding of meaning refers to the arbitrary links any language relies on between expressions in the lexicon and their meaning, which have to be learned. The notion of context includes both the physical utterance context as well as the linguistic and discourse context, including assumptions about the communicative purposes and issues at stake. Reasoning about utterances in context encompasses considering their ‘core’ meaning in relation to aspects of the context, which then can lead to inferring additional meaning beyond the core parts to enrich what is taken to be the speaker’s intended message.

As we noted above, a—and perhaps THE—central theoretical question for this perspective on the semantics-pragmatics interface is to determine which aspects of meaning have to be seen as conventionally rooted, and thus encoded in the lexical entries of particular expressions, and which can be accounted for without such encoding. Two guiding principles that commonly underlie deliberations in this regard, though often just implicitly, are the following: First, meanings should only be conventionally encoded in lexical entries if accounting for related phenomena forces us to, i.e. positing conventional encoding is subject to Occam’s razor. Secondly, if a given bit of meaning is present in some contexts but not in others, that provides an argument against encoding it conventionally. In other words, conven-

tional aspects of meaning are those that are stably present across contexts, and which cannot be explained in terms of general reasoning. In much of this chapter, we will be concerned with phenomena that seem *prima facie* pragmatic, but which have been argued to require conventional encoding of at least some related ingredients to adequately capture the empirical facts. We will complement our theoretical discussions with insights from empirical investigations, which systematically reveal a more gradual picture than expected under a radical convention-vs-context split approach. We will see that many ingredients of meaning, while failing to display a stable presence across contexts, cannot seem to be fully accounted for without assuming further conventional encoding at the lexical level after all.

2.1 The interplay of content and context: basic illustrations

The need for integration of conventional content and contextual information is pervasive. In many instances, both are needed even in just accounting for the interpretation of individual expressions. Consider indexicals, such as first and second person pronouns: what is the meaning of, say, ‘I’? In one sense, it crucially has a stable meaning that systematically contributes a referent to the interpretation of utterances it occurs in. But just who the contributed referent is is of course not at all constant across utterances, as ‘I’ will pick out the speaker of whatever utterance we are

considering. The standard solution is to assume a conventionally encoded meaning that crucially incorporates contextual information: The referent introduced by ‘I’ is conventionally fixed to be the speaker of the utterance, where information about interlocutor-roles, such as speaker (as well as addressee, for 2nd person), is assumed to be directly supplied by (a suitable representation of) the context.¹

Another class of expressions that can, on certain interpretations, draw on the context to fill in their lexically encoded content are third person pronouns, such as ‘she’. These differ from purely indexical pronouns in that their semantic values can depend on another expression in the linguistic context, as (2) illustrates:

- (2) a. Lola said that she runs.
b. Every girl said that she runs.

Identifying what ‘she’ stands for here requires consideration of the semantic value(s) of another expression in the sentence. Besides the ‘co-referential’ and ‘bound’ readings in (2), ‘she’ also has an indexical reading, where it refers to a single contextually salient individual.²

Various other expressions display parallel types of context-sensitivity, including a wide range of phenomena such as domain restriction for quan-

¹Parallel effects can, of course, be observed for various expressions relating to other aspects of the context, e.g., the time (‘now’) and place (‘here’) of utterance.

²Some languages distinguish these readings by using different forms. Note that there also are languages that allow certain bound readings of their first- or second-person pronouns, making the distinctions we’re drawing here for illustration more subtle.

tificational expressions (where, say, ‘every student’ can be contextually understood to effectively mean, e.g., ‘every student in this class’; Westerstahl 1984, von Stechow 1994) and standards of comparison for gradable adjectives like ‘tall’ (which can differ when we’re talking about preschoolers vs. basketball players; see Section 4.2.2 below), and even the semantics of tense (e.g., when ‘I didn’t turn off the stove’ seems to refer to some particular contextually salient time where this happened; Partee 1973).

While the types of cases considered so far come with many intricacies of their own, it is fairly clear what aspects of their interpretation come from context and what aspects have to be conventionally encoded, at least at a basic level (there’s an abundance of technical issues for any full-fledged formal theoretical account integrating contextual information into the interpretation process). Our main focus in the following will be phenomena that are best discussed in the context of the meanings of entire sentences, and the contributions that individual words make to them.

2.2 Compositional semantics for sentence meanings

In order to talk about pragmatic aspects of sentence meanings, we need at least a rough sketch of the general framework for deriving sentence meanings from the lexical meanings assumed for individual expressions. In the framework we’re adopting, sentence meanings are construed in terms of truth conditions. As the first sentence of an influential formal semantics textbook states: ‘To know the meaning of a sentence is to know its truth

conditions' (Heim & Kratzer 1998, p.1). We hasten to note that this does not preclude the existence of other dimensions of 'meaning', which we focus on below. It simply acknowledges as basic the capacity of speakers to differentiate between situations in the world that fit the description provided by a sentence and those that do not.

The truth conditions of a sentence derive from the semantic contributions of the words in the sentence: 'Lola runs' and 'Lola swims' are true in different situations because 'run' and 'swim' describe different activities. The core task of formal semantics is to provide a theory of how the meanings of sentences (and more generally, any complex, multi-word expression) can be derived using the fundamental principle of compositionality (Frege 1892): the meanings of complex expressions are a function of the meanings of their parts and the syntactic structure they occur in. For example, the semantics of predicates like 'run' is modeled as a function from individuals to truth-values, e.g., 'run' will return true if the individual of whom it is predicated is a runner (more sophisticated variants will instead characterize this in terms of the individual having to be the agent of a running event).

2.3 Gricean reasoning

With such a basic semantic framework in place, we can now consider what is conveyed by utterances of sentences in context more broadly, and turn to our central question: which aspects of the overall conveyed meaning have their source in the conventionally encoded content of the expressions in-

volved, and which are derived from contextual information and domain general reasoning about the utterance? To start with a simple case where it's easy to discern and illustrate the sources of different aspects of the meaning that's conveyed, consider the following (adapted from Grice):

(3) Context: A driver stops to talk to a pedestrian.

A: Excuse me, I'm out of gas.

B: There's a gas station around the corner.

Among the things conveyed by B are the literal meaning that a gas station is located around the corner, as well as additional notions about it being functional and open for business (or at least B not being aware of any facts to the contrary). Intuitively, this seems to ride on the idea that if B thought the gas station was not in operation or closed at this time, they should have said so or simply not mentioned that gas station at all.

The most influential account of these types of inferences is due to seminal work by Grice (1975). It is based on a view of conversation as a cooperative endeavor amongst rational interlocutors. His overarching Cooperative Principle, which assumes discourse participants to be cooperating in pursuing a common cause, is more concretely reflected in the Gricean Maxims: the maxim of quality, requiring speakers to be truthful; the maxim of quantity, requiring speakers to provide appropriate amounts of information; the maxim of relevance, requiring utterances to relate to the current discourse goals; and the maxim of manner, requiring appropriately

brief, transparent, and well-ordered utterances. These maxims, which are intended as descriptive characterizations of principles that actually guide interlocutors (rather than prescriptive rules), commonly interact with one another, so that pragmatic inferences often arise by striving for a balance between them. For example, if B in (3) knew that the local gas station was closed, the answer would still have been true (and thus aligned with the maxim of quality), but the mere existence of a gas station is not enough to make B's contribution relevant. The pragmatic interpretation of the existential utterance can be explained in terms of Relevance and Manner: B takes the existence of a gas station to be relevant to the discourse situation, so it must be that the gas station is open and serves gas, but chose to not make that explicit for the sake of brevity.

In addition to these highly contextual inferences that are not very closely tied to particular lexical entries, a host of other inferences are quite generally associated with connectives and other logical expressions. For illustration, take the common temporal inference associated with 'and':

(4) John stepped out of the house and changed out of his pajamas.

The described circumstances seem unusual, given standard societal conventions, because what seems to be conveyed is that the two events occurred in the order described, and John thus changed clothes out on the street. This effect would be puzzling if all we have to work with is the conjunction of statement logic, which is fully symmetric, so that 'p and

q' cannot possibly have a different meaning than 'q and p.' Rather than giving up on the logical analysis of 'and' for purposes of meaning composition, Grice's solution supplements it with a system for deriving additional pragmatic inferences that account for its actual use. In particular, it attributes the order effect to reasoning about the maxims, as adherence to the maxim of manner will generally give rise to the conclusion that a narrated sequence of events occurred in the order they were introduced in (barring any indications to the contrary). In the context of our discussion, a Gricean view allows for a minimal conventionally encoded content for "and" (the logical conjunction) while reasoning about its use in context can add to its overall meaning contribution.

For historical context, it is worth noting that Grice's account of cases like these played a pivotal role in reconciling two opposite approaches in the philosophy of language at the time: one analyzing meaning in natural language using logic, the other, ordinary language philosophy, focusing on language use. Thereby, Grice defended logical approaches to semantic meaning, while independently accounting for systematic aspects of language use. This general approach created the basic framework of encoding certain aspects of meaning in the lexical entries of particular expressions (such as the logical core of conjunction), while deriving other aspects (such as the order effect above) in independent ways, typically in terms of reasoning about language use in context. In the next section, we turn to the perhaps most influential, and most studied, phenomenon in this realm,

namely that of scalar implicatures. These are of particular interest to our current discussion for at least two reasons. First, scalar implicatures embody the Gricean idea that the conventional content of a lexical entry can include less than what might seem to be its overall contribution to meaning at first sight. Second, scalar implicatures fundamentally involve comparing linguistic expressions; as we will see, this may require conventional association between lexical entries in some cases.

3 Scalar implicatures

Both the first case study on scalar implicatures in the present section and the second one in the next section on scalar adjectives are fundamentally concerned with the notion of scales and the extent to which they may require lexical association of particular expressions with such scales. Formally speaking, a scale integrates several abstract objects (e.g., linguistic expressions or degrees) into an ordered set. Association with a scale allows for comparison with other elements of the scale; when these elements themselves are lexical entries, this brings into play the notion of an alternative. While there is a fair amount of—perhaps under-appreciated—overlap in issues related to the role of scales in scalar implicatures and scalar adjectives, and potentially fruitful connections can be drawn between them, we will largely present each set of phenomena in its own right, while highlighting potential parallels when discussing scalar adjectives.

3.1 The basic phenomenon

Scalar implicatures have particular relevance for the central concerns of this chapter, as they are standardly seen as a pragmatic phenomenon involving competition between lexical entries, and thus invoking so-called alternatives. As the name suggests, these are taken to be ordered on a *scale*; we will elaborate this notion below. Following much of the discussion in the literature, let us illustrate scalar implicatures by looking at the quantifier ‘some’:

(5) Some of the guests already left.

Uttering this sentence intuitively conveys that there also are guests that are still there, i.e., that not all guests left. This is commonly seen as the result of reasoning about the alternative possibility of making the statement in (6):

(6) All guests already left.

In principle, one could also consider the lexical entry for ‘some’ to encode a meaning equivalent to that expressed by ‘some but not all.’ However, there are many environments where these are not intuitively equivalent:

- (7) a. If some of the guests already left, the party was a failure.
b. If some but not all of the guests already left, the party was a failure.

Clearly, the possibility considered in the ‘if’-clause of the ‘some’ version includes the case where all of the guests already left (in which case the party would have been an even greater failure, if anything). In contrast, the ‘some but not all’ version allows for the (contextually implausible) possibility that if all of the guests left, the party somehow wasn’t a failure. Another indication that the ‘but not all’ part of the meaning conveyed by (5) is not encoded in the literal meaning derived from its lexical parts is that it can be explicitly cancelled without giving rise to a sense of contradiction (8a)—in contrast to cancelling literal meaning, as in (8b).

- (8) Some of the guests already left.
- a. ✓ In fact, everybody’s gone.
 - b. # In fact, nobody’s gone.

This contrast suggests that *some* inherently negates *none* as part of its lexical meaning, but the negation of *all* cannot be encoded on an equal footing, for otherwise (8a) should give rise to the same sense of contradiction as (8b). Rather, as we already alluded to, the negation of *all* obtains as a scalar implicature. Scalar implicatures arise when speakers do *not* make a more informative claim using an alternative sentence. Let us walk through our current example. Both (5) and its *all* alternative (6) are true if all of the guests left, but only (5) is still true if only a subset of the guests left. In other words, the situations that make (6) true are a proper subset of those that make (5) true. So from the point of view of a speaker

who knows that, in fact, all of the guests left, using either (5) or its *all* alternative would count as a (technically) true statement, but from the point of view of a hearer ignorant of the facts, (5) would not be as *informative*, because its truth is not enough to conclude with certainty that all of the guests left. Therefore, an informed and cooperative speaker is expected to use (5), and not the more informative *all* alternative, only if they do *not* think that all of the guests left. That explains why, upon interpreting (5), we typically conclude that not all of the guests left.

3.2 Constraining Alternatives

3.2.1 The symmetry problem

In light of the key role of reasoning over potential alternative utterances in the discussion above, it comes as no surprise that one of the central issues in the literature is just what alternatives are to be considered in the process. That this has to be constrained has been clear from early on, due to what has become known as the ‘symmetry problem’: if not only (6), but also (9) were alternatives to (5), the desired derivation of the standard pragmatic interpretation of (5) would be lost.

(6) All guests already left.

(9) Some but not all guests already left.

After all, (9), too, is a more informative statement than (5) (assuming, as before, a ‘some and possibly all’ interpretation of ‘some’). So the

hearer should apply the same reasoning to this statement as to (6), and conclude that (9), too, is false, i.e. that either all guests already left, or none left yet. But in conjunction with the literal meaning of (5) (that some and possibly all guests left) this is inconsistent with the parallel inference based on (6) (namely, that not all guests left). Since there is no obvious way to choose between these inferences, this at best leaves us with a very weak overall interpretation, namely that some and possibly all guests left, while the speaker lacks evidence to assert either that all guests left or that only some guests left. But the standard observed interpretation is that the negation of (6) is inferred, whereas (9) does not seem to enter the reasoning at all. The conclusion is that we need a systematic way of excluding (9) from being considered as an alternative when reasoning about (5). And more generally, we need a systematic account of which alternatives are taken into consideration when evaluating a given sentence.

3.2.2 Alternatives and the lexicon

Seminal work by Horn (1972, and subsequent work) proposed that expressions like ‘some’ are associated with scales that provide the relevant alternatives. The perhaps most basic version of this proposal is that it is part of speakers’ lexical knowledge that certain words are associated with scales, like $\langle \textit{some}, \textit{all} \rangle$, and that the alternatives needed for the reasoning above are to be found on those scales. On this view, truth conditional meaning is not the only type of information that is stored in lexical entries—certain

properties that play an essential role at the *pragmatic* level can also be encoded lexically. This notion of ‘Horn-scales’ forms the basis of much discussion in the literature, although the theoretical underpinnings are far from clear, as highlighted by recent work homing in on the question of how precisely alternatives for a given expression are selected.

While some early work (Gazdar 1979, a.o.) put forward the notion that Horn scales are essentially ‘just given to us’, this can hardly be the whole story: claiming that alternatives are retrieved from memorized scales that are lexically associated with a given expression immediately raises the question of *how* those associations came to form, and why those but not others. Thus, more systematic considerations on why certain expressions but not others are to be included on a scale for a given lexical entry seem to be called for. One necessary condition on alternatives was already implicit in the above: they must be related by strength in terms of an entailment relation. The reasoning based on the maxim of quantity spelled out above crucially considers logically stronger statements. But ‘some but not all’ is logically stronger than ‘some’, as well. So there must be more to the *sufficient* conditions for being an alternative.

One straightforward option to consider is that only lexical entries are candidates for being alternatives. This would rule out all more complex expressions, including ‘some but not all.’ However, there are cases where more complex expressions do seem to enter the process of scalar reasoning, e.g., when the alternatives are explicitly present in the local discourse

context and in so-called downward entailing environments (such as the restrictor of *every*; Ladusaw 1980):

(10) It was warm yesterday, and it is a little bit more than warm today.

(Matsumoto 1995, p. 44)

(11) Every day on which it was a little bit more than warm we went swimming.

(based on parallel examples in Katzir

(2007*a*))

In both cases, ‘warm’ and ‘a little bit more than warm’ seem to be considered as alternatives, as reflected in the inferences that yesterday was *not* a little bit more than warm (10) and that we did *not* (necessarily) go swimming every day that was (merely) warm, respectively (11). Katzir (2007*b*) argues that parallel examples do not exist for ‘some but not all’,³ and that we thus need to differentiate the two cases. Matsumoto (1995), expanding on Horn (1972, 1989), formulates a monotonicity constraint on scales, such that all scalemates have to be positively or negatively scalar (for a formal definition of the relevant notions, see Sevi 2005). (Note that this differentiates between (9) and (10).) However, as Katzir (2007*a*) points out, the monotonicity constraint need not be met for cases like (13), where the non-monotonic expressions *exactly three semanticists* and *John but not Mary* seem to give rise to alternative-based reasoning:

³The resulting sentences are strange, as Katzir notes as well, so one may want to be cautious in drawing too firm a conclusion here; see Katzir’s paper for discussion and his take on this.

(12) I doubt that exactly three semanticists will sit in the audience.

(13) Everyone who loves John but not Mary is an idiot.

No doubt is conveyed here about the possibility of there being more than three semanticists (as encoded in the commonly assumed meaning of ‘three’ that is semantically equivalent to ‘at least three’); and loving John is not claimed to make one an idiot, as long as one also loves Mary. This suggests that *three* and *John* can indeed serve as alternatives to *exactly three* and *John but not Mary*. Based on these examples and some further considerations, (Katzir 2007a) proposes a formal definition of alternatives that, in addition to lexical substitutions, allows substitutions with constituents in the discourse context (making ‘a little bit more than warm’ an alternative for ‘warm’ in (10)), as well as with less complex substitutions (bringing ‘warm’ into play as an alternative in (11)). Several other authors have further extended complexity-based proposals for constructing structural alternatives, (Katzir 2007b, Fox & Katzir 2011, Trinh & Haida 2015, a.o.). However, in any variant, issues of both under- and overgeneration remain, as reviewed by Breheny et al. (2018). For example, as noted by Swanson (2010), not even all lexical alternatives seem to be considered in the relevant reasoning processes, in particular where meanings of the ‘some but not all’ kind seem to be lexically encoded, as is arguably the case for ‘intermittently’:

(14) a. The heater sometimes squeaks.

- b. The heater intermittently squeaks.

Complexity-based approaches (without anything like the monotonicity constraint added in) predict (14b) to be an alternative to (14a), but there is no inference that the negation of (14b) holds—to the contrary, the observed inference (based on the negation of alternatives such as ‘constantly’) is precisely that (14b) DOES hold. Such lexical instances of the symmetry problem thus pose a serious challenge to these structural approaches to alternatives.

3.2.3 Alternatives and context

Another possible response to examples involving more complex alternatives is to maintain a notion of scales as involving only lexical alternatives and explain other examples in a different way, i.e., consider lexical and contextual alternatives as distinct phenomena with distinct explanations. Whatever the best explanation for these may be, it is clear that the range of contextual alternatives is much broader than what we have considered so far. Consider the following examples of what has traditionally been considered a *particularized* (as opposed to *generalized*, or more specifically *scalar*) implicature (see Hirschberg 1991, for extensive discussion on this)

(15) A: Bill is a swimmer. What does Lola do?

B: She’s a runner.

(16) It was warm yesterday and it is warm and sunny with gusts of wind

today.

(Katzir 2007*a*, p. 687)

As discussed in the introduction, B seems to convey both that Lola (identified as the referent of ‘she’) is a runner and that she’s not a swimmer in (15). But of course this is completely driven by the specific context at hand, and in no way generally associated with the lexical entry for ‘swimmer’. Similarly (16) suggests that yesterday it was not sunny with gusts of wind, which again is not part of the lexical meaning of *warm*. Examples like these have often been characterized as a distinct case of their own, given that they depend on very specific aspects of the context. But note that (at least certain variants of) structural approaches will analyze these examples in the same way as more standard ‘scalar’ implicatures, as the mechanism for generating alternative structures can draw on contextually present expressions, making *warm and sunny with gusts of wind* an alternative of *warm* in the context of (16). (Also note that they may need additional constraints to rule out the lexical alternative ‘swimmer’ when it is not present in the context). Hirschberg (1991) put forth an alternative approach extending a more lexical, scale-based approach, by allowing for contextually constructed ad hoc scales.

Another illustration of what can easily be seen as an ad hoc scale based on context and world knowledge is the following:

(17) *Context: A left Los Angeles and is on their way to Seattle*

A: I passed Sacramento.

In the given context, passing Portland entails passing Sacramento in a way that passing Phoenix does not, and thus Portland qualifies as an alternative to Sacramento on the *ad hoc* scale $\langle Los\ Angeles, Sacramento, Portland, Seattle \rangle$. Accordingly, A's statement conveys that they have not yet passed Portland (while not conveying anything about passing Phoenix). Examples like these suggest that the alternatives considered in implicature reasoning need not stand in a semantic entailment relationship - purely contextual entailment (resulting from the combination of the semantic meanings at play with contextual assumptions) suffices for triggering the relevant comparisons with regards to the strength of the statements.

3.2.4 Taking stock

Our discussion here has focused on perspectives that all share the assumption that the lexical entries of expressions like 'some' do not directly encode the stronger 'some but not all' meaning, but rather have a semantics amounting to the more general 'some and possibly all' meaning.⁴ Instead, they assume that the strengthened meaning comes about by considering alternative statements.⁵ The main dimension of divergence that we focused

⁴Other accounts, e.g., the defaultist view (Levinson 2000), don't necessarily share this assumption, as they assume a more direct conventional association between the lexical entry and the 'some but not all' meaning, while still allowing for a 'some and possibly all' interpretation, perhaps via ambiguity. We are unable to go into further details on these for reasons of space.

⁵Note that this is also true of so-called grammatical approaches, which don't see the implicature-generating process itself as pragmatic, but rather as coming about through an exhaustivity operator akin to 'only' in the structure. Again, space reasons preclude us from discussing these in more detail; see Sauerland (2012) for an introduc-

on here concerns the different ways in which these alternatives are selected, and how they relate to the lexicon. On traditional lexical approaches, the alternatives are directly hard-wired in the lexicon in the form of Horn-scales. Other approaches have a looser connection to the lexicon, as they merely see it as one source of alternative expressions for a more general mechanism, which also can take into account contextually salient alternatives. This extends the coverage to more context-based cases of particularized implicatures. Extensions of lexical approaches that try to capture the more contextual cases as well also loosen the connection to the lexicon by introducing the notion of ad hoc scales.

Taking a step back, a key question that is at play in this area is precisely whether there are any lexically associated alternatives that have a special status distinct from those that are in play due to salience in the context. Most current accounts try to work with broad enough notions to accommodate both in one system, but it is also perfectly possible that there are more fundamental differences between the two types of cases meriting a separate treatment. The answer to that question more or less determines the extent to which this key ingredient for a central type of pragmatic reasoning requires specific lexical encoding of information that goes beyond core truth-conditional meaning. As things stand, this remains a genuinely open question that will need to be resolved in future work.

tory overview.

3.3 Connections to experimental work

On standard theoretical accounts, implicatures have two properties that have historically put them at the forefront of the emerging subfield on the experimental study of meaning: (i) they are a *secondary* type of meaning, taken to build on conventional meaning (the primary meaning); (ii) they are *optional*, in that contradicting them does not come with the red flag associated with contradicting conventional meaning. This can easily be mapped onto a cognitive model where conventional meaning is computed first and then, in an optional secondary step taking place in real time, contextual information is integrated into the interpretation process to derive the implicature interpretation (where warranted). There is a by now rich and extensive body of literature investigating the processing and acquisition of implicatures (see Chemla & Singh 2014*a,b*, for an overview). Consistent with the simple cognitive model just laid out, early results were argued to point towards delays for implicature interpretation in both domains, with relatively slower access to implicature meanings in online processing and later availability in acquisition (Bott & Noveck 2004). These results have been commonly interpreted as supporting the view that literal (e.g., ‘some and possibly all’) meanings have a primary status, presumably due to being lexically rooted, while implicature (e.g., ‘some but not all’) meanings have to be derived through additional pragmatic reasoning which takes time in online processing, and which has to be mastered separately by children in the acquisition process. The empirical picture has

since grown substantially more complex, with various studies showing rapid access to implicature meanings in other contextual setups (e.g. Grodner et al. 2010) and earlier diagnoses of acquisition ages based on richer contextual support (e.g. Papafragou et al. 2018). While we cannot review this literature in greater depth here, it is worth noting that one key factor that has been discussed for variation in both processing speed and acquisition age is the contextual accessibility of the alternatives. For example, Katsos & Bishop (2011) argue that the apparent delays in acquisition are not necessarily due to an inability to carry out the relevant implicature reasoning, but rather should be attributed to issues with accessing the (correct) set of alternatives (also see Barner et al. 2011, Skordos & Papafragou 2016). This interpretation moves the focus from mastering an independent high-level reasoning module to a mechanism that young language learners need to leverage when structuring the lexicon in order arrive at mature in-context meanings.

The experimental work referenced so far, however, has usually focused on a very limited set of scales (e.g. *<some,all>*). Another recent line of work has been looking at variation in the implicature effects across a wider range of expressions. For example, van Tiel et al. (2014) drew a list of 43 pairs of alternatives and proceeded to experimentally measure how likely an implicature is to obtain for each pair. They identified two factors presumably affecting adults' access to alternatives: semantic distance (e.g. how much more intense is *hot* as compared to *warm*) and bounded-

ness (e.g. *certain* representing an end-point as compared to *likely*). Those factors however do not define clear categories, as van Tiel et al. observed a very gradual variation in inference rates across the 43 pairs of alternatives they tested. This failure to observe clear-cut categories does not immediately support the idea of a radical split between lexically-given and contextually-determined alternatives. At the same time, it suggests the possibility of a model that integrates contextual factors at the lexical level: it could be that speakers conventionally encode lexical entries on a conceptual space that includes semantic distance and boundedness, and that these properties are leveraged, and maybe further contextually weighed, in the process of computing scalar implicatures.

Finally, some work has begun to investigate the extent to which the availability of alternatives can be modulated through manipulation of the presence of potential alternatives in the experimental context. Kim (2016) shows that varying the frequency of occurrences of *all* within an experiment has a substantial impact on the frequency of implicature interpretations of *some*, suggesting that the salience of an alternative in context directly modulates the salience of the implicature. With regards to symmetric alternatives, on the other hand, Uppili (2018) finds that including, e.g., *only some* in an experiment leads to higher rates of implicature interpretations of *some*, i.e., participants assimilate *some* to *only some*, rather than considering it as an alternative that *some* might contrast with.

As the empirical picture becomes more fleshed out, it is important to

integrate insights from the experimental and theoretical literatures further. This will not only help determine which theoretical take on alternatives best aligns with empirical findings, but it will also help test new hypotheses that arise from more refined theoretical proposals. Ultimately, this will get us closer to answering the question of how much information needs to be lexically encoded to account for the range of observed implicatures.

4 Scale structures and adjective meanings

The previous section highlighted how the notion of alternatives raises the question of whether accounting for scalar implicatures requires positing certain ingredients to be lexically encoded. We saw that while Horn scales have traditionally been thought of as ‘given to us,’ recent proposals have worked towards deriving rather than stipulating alternative sets, with some success but also with some limitations. At the end of the day, it remains a genuinely open question at this point whether at least some lexical entries are associated with (something like) a conventionally encoded Horn scale. In this section, we will focus on another notion of scale traditionally taken to be lexically given, which in part (but only in part) can be directly related to the notion of Horn scales.

As detailed below, so-called *gradable* adjectives, like *warm* and *hot*, crucially invoke *degree scales* as their key semantic ingredient. The case of *warm* and *hot* in particular indicates a connection between the two notions

of scales: both adjectives refer to the same degree scale (of temperature) and differ only in their respective thresholds on that scale; and they also are standardly seen (on the relevant accounts) as forming a Horn scale, $\langle \textit{warm}, \textit{hot} \rangle$, to account for the *not hot* implicature of *warm*. The central role of the very same degree scale in the lexical entries with the different thresholds opens up the possibility that all the information that is needed to encode this Horn scale is already independently present in the lexicon. This specific deflationist option, of course, can only apply to alternatives that consist of gradable adjectives; the issues in the previous section, however, are mostly concerned with expressions of a different nature, and considerations about what content gradable adjectives lexically encode will therefore not settle all those issues by any means. In light of space constraints, we are unable to discuss the relation between the two notions of scales in more detail, even though the topic promises further insight into the lexical root of some pragmatic phenomena.⁶ We will, however, critically review the traditional lexicalist view whereby gradable adjectives are mentally represented in the form of degree scales. While the ultimate conclusions here may be different from those warranted for scalar implicatures in general, our discussion highlights parallels between issues in the two domains, specifically concerning the tension between the (at least

⁶We refer the interested readers to Chanchaochai & Zehr (2019) who argue for an alternative-based analysis of certain Thai degree constructions, thus fleshing out the connection between the two types of scales further to account for a specific new set of linguistic data. On a similar note, Beltrama (2018) offers an alternative-based account of emphatic *simply* for extreme adjectives.

apparent) need for lexical encoding and the desire for general mechanisms that can be applied across the board.

4.1 Gradable vs. non-gradable adjectives

It has long been noted that two classes of adjectives can be distinguished, namely the gradable vs. non-gradable ones, based on their distributional properties (Sapir 1944). *Gradable* adjectives naturally appear in a number of so-called *degree constructions* whereas *non-gradable* adjectives generally cannot be used in these. For example, the adjective *cheap* is usually classified as gradable (18) whereas the adjective *free* (in the *free of cost* sense) is not (19) (question marks indicate that the sentences are degraded).

- (18) a. Attending college is cheaper in France than in the United States
b. How cheap is attending college in France?
c. Attending college is very cheap in France
- (19) a. ?? Attending college is freer in France than in the United States
b. ?? How free is attending college in France?
c. ?? Attending college is very free in France

In light of the main theme of the present chapter, we can now ask whether gradability is encoded lexically and, if so, what type of information must be stored in the lexicon in order to capture the semantic properties of adjectives with regards to gradability. Intuitively, the meaning of

comparative uses of gradable adjectives, as in *cheaper*, amounts to comparing the degree to which two entities have a certain property. The comparative morpheme (*more/-er*) introduces the element of comparison. But according to one prominent approach, the degree element is already present in the adjective meaning itself, so that *cheap* on its own denotes a relation that maps individuals onto degrees on the relevant scale, here involving cost (Cresswell 1976, Kennedy & McNally 2005, a.o.). This is fundamentally different from the meaning of non-gradable adjectives, which have standard predicate denotations (corresponding to sets of individuals). On this type of view, the gradability properties of a given adjective are fundamentally lexical, as they are part and parcel of its core semantic make-up. In other words, the minimal contrast between (18) and (19) is taken to indicate that the gradability of *cheap* and the non-gradability of *free* inherently follow from their meanings, leaving little room for contextual factors to play a role in an adjective's distributional pattern in this regard.

At the same time, as we saw with implicatures, there is a general inclination to account for aspects of meaning as deriving from more general principles whenever possible, rather than simply stipulating them in the lexicon. And so one could also entertain the notion that speakers need not encode the (non-)gradability of *cheap* vs. *free* in their mental lexicon, and that the contrastive patterns in (18) vs. (19) obtain from general pragmatic considerations rather than a difference in the type of meaning involved in the two cases. For example, speakers could store both *cheap* and *free*

as simple predicates relating to cost, but reason that the very concept of something being *free* leaves no room for further ordering of cost, as otherwise required by degree constructions like those exemplified in (18). A potential line of reasoning supporting such a pragmatic view of gradability could point to the apparent fluidity of the gradability split. While the adjective *pregnant* is standardly considered a non-gradable adjective, it can still naturally appear in degree constructions in everyday conversations, as in (20).⁷

- (20) a. Do you think I could be *more pregnant* than I thought?
b. We just found out we're pregnant, find out *how pregnant* I am!
c. This is what it's like to be *very pregnant*

If one were to strictly maintain that each adjective is stored in the mental lexicon as either a relation involving degrees on a scale or a binary property that can be modeled as a set, these data introduce some uncertainty as to whether *pregnant* should be encoded as non-gradable, following the tradition in the literature and its fundamental either-or semantics, or as gradable, in an effort to account for data like (20). A possible resolution of this issue from the perspective of the traditional analysis of *pregnant* as non-gradable is to enrich the interpretative machinery with an optional mechanism that can coerce non-gradable expressions into gradable

⁷The examples in (20) were found online at, respectively, <https://www.mumsnet.com/Talk/pregnancy/710893-Do-you-think-I-could-be-more-pregnant-than-I>; <https://www.youtube.com/watch?v=e2DkWskS61U> and <https://www.scarymommy.com/what-its-like-to-be-very-pregnant/>.

ones by introducing contextually salient and plausible scales (e.g., of time passed pregnant in the cases above).⁸ One can then maintain the split in lexical representations, where gradable adjectives involve degrees on lexically defined scales, whereas non-gradable adjectives require a modification via coercion of their core semantics in context to account for (more or less exceptional) gradable uses.⁹ By defending a lexicalist approach, coercion-based proposals are able to address the tension in how gradability manifests itself: as a linguistic phenomenon giving rise to specific grammatical constructions (recall the contrast between (18) and (19)) but, at the same time, with a determining role of extra-linguistic factors (recall how (20) came with somewhat atypical conceptions of pregnancy). Those proposals render the linguistic aspect of gradability as clear-cut lexical categories that appear in specific grammatical constructions, and its extra-linguistic aspect as a coercion operation that provides leeway to redraw those category boundaries when contextually motivated.

⁸Coercion is a type of operation often introduced to account for a range of linguistic phenomena, of which gradability is only one. See Lauwers & Willems (2011) for an overview.

⁹Note that this is not dissimilar to the discussion of ad-hoc scales for scalar implicatures, a potential parallel that is worth exploring in future research.

4.2 Relative vs. absolute adjectives

4.2.1 Bounded vs. open scales

Among gradable adjectives, we can further distinguish between so-called *relative* vs. *absolute* gradable adjectives (Unger 1975, Kennedy & McNally 2005). For example, *dry* is a gradable adjective allowing comparatives etc. (21), patterning with *cheap* (and contrasting with non-gradable *free*), but it also patterns with the non-gradable adjective *free* in some respects (22).

- (21) a. Californian summers are drier than Arizonian summers
b. How dry are Californian summers?
c. Californian summers are very dry
- (22) a. ?? Attending college in France is completely cheap
b. Attending college in France is completely free
c. The soil is completely dry

This is because unlike *cheap*, which is a relative adjective, *dry* can target an endpoint, i.e., it is an absolute adjective. This can be modeled in terms of the structure of the scales involved: relative adjectives are associated with open scales, whereas absolute adjectives are associated with bounded scales.¹⁰

¹⁰See Kennedy (2007) for an explicit formalization of *expensive* as denoting an open scale—excluding null costs—and Lassiter (2010), Lassiter & Goodman (2013) who take issue with it.

Importantly, there are no clear general conceptual grounds that would require the association of *cheap* with an open scale and that of *dry* with a bounded scale. To see this, imagine the adjective *freap*, defined as follows: (i) what is free is *completely freap*, and (ii) if not free, then the cheaper, the *freaper*. Formally put, *freap* is associated with a bounded scale: the degrees on the scale are ordered following the costs they formalize, and null costs are mapped to the endpoint degree of the scale. There is no *a priori* conceptual reason why *cheap* should not mean what *freap* means, for the scale of *freap* has the exact same structure as that of the attested adjective *dry* and it actually maps rather directly to the (non-linguistic) measure of cost. The apparent absence of a general pragmatic principle that could rule out the possibility of an absolute counterpart of *cheap* (i.e. *freap*) suggests that the relative nature of *cheap* is merely a conventionally fixed property, which accordingly has to be part of its lexical entry.

4.2.2 Relative adjectives and context

The relative vs. absolute split is relevant to the focus of the present chapter in yet another way, in that relative adjectives seem to exhibit a lexically encoded dependence on context that absolute ones do not. Compare contexts where we are talking about either a one-room apartments in a small town or about houses in Manhattan. The sentence *A rent of \$2,000 is cheap* will likely be false in the former but true in the latter case. Note that speakers can also offer a class of comparison explicitly, using

for-phrases (23a). However, as Siegel (1979) notes, this does not generally seem to be available for absolute adjectives (23b), suggesting that their semantic contribution is not context-dependent in the same way.

- (23) a. That's a cheap rent for a full house in Manhattan!
b. ?? That's a dry towel for a bath towel!

While such data suggest that some, and only some, gradable adjectives should encode context-dependence in their lexical entry (at least insofar as classes of comparison modulate their threshold) other data once again question the sharpness of the division Sassoon & Toledo 2011, McNally 2011.

- (24) a. The soil is dry for this usually very green land
b. The towel is very dry for a towel stored in a sauna

Two approaches to this offer themselves: (i) speakers store semantic representations for relative and absolute adjectives that formally differ in whether they require contextual parameters. (23b) would then be degraded because the semantic entry for *dry* is incompatible with contextual modification, though an additional coercion mechanism can override this. Alternatively, (ii) speakers store semantic representations of the same formal type for relative and absolute adjectives (thus explaining their co-occurrence in degree constructions) and context-sensitivity emerges as a split that is not lexically stipulated. The second solution has been fa-

vored both by supporters and critics of degree semantics. Defending a degree semantics approach, Kennedy (2007) assumes the open vs bounded scale structures we mentioned earlier and derives context sensitivity from there: all gradable adjectives convey the surpassing of a degree point on their scale, but only bounded scales define a non-arbitrary degree of reference (their endpoint) whereas speakers need to recruit contextual factors to establish a degree of reference on open scales. From another perspective, Burnett (2014), who assumes that gradable and non-gradable adjectives alike define binary properties at their core, proposes to derive gradability from a context-sensitive semantics for both relative and absolute adjectives, but argues that the two types of adjectives are dependent on context in crucially different ways. Lassiter & Goodman propose a more radical approach where all gradable adjectives define the exact same type of semantics at their core and where “the relative/absolute distinction is not a binary distinction, but a matter of degree.” (Lassiter & Goodman 2013, pp.601-602)

At the end of the day, the question remains relatively open as to what principles underlie the categorization of different adjectives as relative vs. absolute in the mental lexicon, specifically with respect to how much of this categorization is derived with every single interpretation and how much has to be learned and stored in the lexicon. The semantics-pragmatics division we posit makes this categorization question pervasive, for usage often does not perfectly align with ontologies posited by semanticists. As we

have seen with gradable adjectives in particular, one standard solution is an analysis in terms of coercion. Results from recent experimental work focusing on the relative-absolute distinction have elicited signals of additional processing associated with cross-category interpretations (Frazier et al. 2008, Bogal-Allbritten 2012, Aparicio et al. 2015). While those results are consistent with coercion-based views that assume a sharp lexical distinction between relative and absolute adjectives, more experimental work is needed before one can discard analyses that attribute more uniform semantics to categories of gradable, or even non-gradable, adjectives.

5 Presuppositions

The notion of presupposition constitutes the first instance in the history of the study of natural language meaning where a distinction between different layers of meaning was invoked. Frege (1892) not only laid the foundations for the compositional approach of meaning alluded to in Section 2, but also discussed issues that arise because speakers' truth-value judgments with regards to sentences like (25) tend to be unclear or inconsistent.

(25) The Queen of France is not bald.

Following the Frege-Strawson analysis Strawson (1964), the source of confusion is the meaning of *the*, which contributes a *presupposition* that there is a unique referent fitting the noun phrase description. If this presupposition (here that there is a unique Queen of France) fails to be sat-

isfied, then semantic composition cannot proceed and the sentence fails to receive a truth value. Expressions that exhibit properties parallel to those of *the* here are standardly referred to as *presupposition triggers*. Given the association of certain lexical items with presuppositions, one may posit an additional layer of presuppositional information as part of (certain) lexical entries. Unsurprisingly given the overall thread of this chapter, there are ongoing debates whether this is indeed warranted, or whether a more explanatory account, where presuppositional properties are derived in a more general way, is called for. In light of the multi-faceted set of issues in accounting for presuppositions theoretically, variations of this question arise concretely for separate aspects of the phenomenon. First, there is the question of where presuppositions come from, and how they attain their status—what is referred to as the ‘triggering problem.’ Second, the apparent variation in more fine-grained properties across (classes of) presupposition triggers has given rise to the suggestion that triggers differ in how and to what extent the presuppositional nature of certain bits of information is lexically encoded. Finally, the special behavior of presupposition triggers in embedded environments—what is known as the ‘projection problem’ (introduced in more detail in the following section)—raises questions about the extent to which embedding operators and connectives have to lexically encode specific aspects of how they pass on presuppositions to the larger structure they appear in. We turn to these respective issues in turn in the following sections.

5.1 The Triggering Problem

A full account of presuppositional phenomena requires answering the following questions: i) Which expressions are presupposition triggers? ii) What do they presuppose (potentially in contrast to other, non-presupposed content)? iii) What makes the relevant content have presuppositional status? The standard take on the first two questions relies on how presuppositions interact with the compositional process, in particular with regards to embedding under certain types of operators: presuppositions ‘escape’ (i.e., remain unaffected by) operators such as negation, modals, and *if*, all of which are ‘entailment-canceling’ operators (cf. the family of sentences tests in Chierchia & McConnell-Ginet 1990).¹¹ This phenomenon is known as ‘presupposition projection.’ To illustrate, start by considering the meaning of (26), which can be seen as involving the components below.

- (26) Lola stopped running.
- a. First Lola was running, . . .
 - b. . . and then she was not running.

These two components of the meaning contributed by *stop* to the sentence behave differently under embedding, e.g., only the second part appears to be affected by sentential negation (paraphrased below in bold):

¹¹Two caveats: first, there are other expressions with similar, though arguably distinguishable, behavior under embedding, e.g., conventional implicatures in the sense of Potts (2005); second, presuppositions are also taken to have a second defining property, of being backgrounded or taken for granted, which we’ll come back to momentarily.

- (27) Lola didn't stop running.
- a. First Lola was running, . . .
 - b. . . and then **it-was-not-the-case-that** she was not running.
(i.e., she was (still) running)

In other words, the first part, which is unaffected by negation, is presupposed and projects, and it is diagnostics like these that allow us to identify presupposition triggers and the content they presuppose. An additional aspect of presuppositions is that they tend to be pragmatically backgrounded, and thus not constitute the main point of the sentence.

The third question above constitutes the triggering problem, namely how it comes to be that the content in (26a) has presuppositional status. One answer, long dominant in the literature at least since Heim (1983, building on Karttunen 1973 and Stalnaker 1973, but in a new formal framework), is that presuppositional information is separately encoded as such in lexical entries of the relevant expressions, as a condition on which contexts a presuppositional sentence can be uttered in.¹²

However, much recent work, reviving a perspective first raised in early work by Stalnaker (1974), has raised what amounts to an explanatory challenge questioning why it should be lexically stipulated that certain parts of the meaning of an expression like *stop*, but not others, have presuppositional status. After all, it doesn't seem like a coincidence that the split

¹²In Heim's context change semantics, this is technically implemented by utilizing partial functions.

in (26) is as it is, as it can be found systematically across languages (and there don't seem to be any attested variants of *stop* that do this the other way around). The alternative pursued in this line of work is to derive the presuppositional status from more general principles. The key assumption that is generally shared across variants of such accounts is that the content that winds up being presupposed is lexically encoded as part of the simply entailed content, i.e., that (26a) and (26b) are entirely on a par as far as the lexical semantics is concerned. The special status of (26a) and its behavior under embedding operators is then derived pragmatically in one way or another. To sketch the general gist of the central idea, it's useful to start by considering what would happen if sentential negation applied to the content of the sentence in (27) wholesale:

- (28) **It-is-not-the-case-that** ...
- a. ... first, Lola was running, ...
 - b. ... and then she was not running.

Note, first, that the negation of a conjunction as a whole ($\neg(p\&q)$) is strictly weaker than the projection reading with only one conjunct negated ($p\&\neg q$), meaning that this is compatible with situations where the projection reading is false, for example, if Lola was not running to begin with. The negation of *stop* actually seems to permit such interpretations when called for in context:

- (29) A: I think Lola may have stopped running.

B: No - She didn't stop running, since she wasn't running in the first place.

Thus, in contrast to the general case where we find projection, the content in (26a) does not necessarily seem to escape negation, and any theoretical approach will have to account for this option. On lexicalist views, an additional operator such as that of 'Local Accommodation' in Heim (1983) has to be invoked.¹³ On the other approaches we're now considering, the availability of this interpretation falls out automatically, since this content is lexically encoded as regular content. But what do these accounts say about the interpretation involving projection above?

One family of approaches invokes parallels to implicatures, by positing alternatives for presupposition triggers that form the basis for reasoning leading to the projection interpretation. For example, Abusch (2002) assumes that *continue* is an alternative to *stop*, along with a principle that one of the alternatives of a sentence has to be true in a given context. The fact that *continue* shares the ingredient in (26a) then is used to account for the presuppositional status of this content (though the account has to say more to explain projection; see Abusch 2010 for more details). In a similar vein, Romoli (2015) accounts for presupposition projection by alluding

¹³The phrase '(Global) Accommodation' was initially reserved for situations where interlocutors accept some information as backgrounded that is yet new to them, for questioning the presupposition would otherwise disrupt the flow of the conversation (Stalnaker 1973). *Local* Accommodation also avoids such a presupposition crash, but it does so in linguistic environments that end up canceling the information in some form at the global level, for example negation as in (29).

to reasoning that is even more closely parallel to that involved in implicatures. He proposes that *stop* is associated with *used to* as an alternative, i.e., the lexical encoding of just the first ingredient in (26) above. Under negation, this yields a strictly stronger reading (i.e., *John didn't use to run* entails *John didn't stop running*), and the idea is that just as in the case of implicatures, the projection interpretation with the 'ran before' content not affected by negation comes about via strengthening through the negation of the stronger alternative: *It-is-not-the-case that Lola didn't used to run* amounts to *Lola used to run* - precisely the content of the projected presupposition. While we have to refer the reader to the original proposals for further details, an obvious issue that arises for these accounts is just why the relevant presupposition triggers are associated with the alternatives posited on the respective accounts (and not others) - this is entirely parallel to the issue of the lexical status of alternatives for implicatures discussed in Section 3. Another interesting aspect of Romoli's proposal in particular is its empirical prediction that speakers derive so-called presuppositions, as in (26), in the same way that they derive scalar implicatures. A number of experimental studies (Romoli & Schwarz 2015, Bill et al. 2015, Kennedy et al. 2015, Bill et al. 2018) have tested precisely this prediction, and while there are some parallels, the overall evidence of different behavioral results for presuppositions and implicatures is more in line with lexicalist accounts of presuppositions (or at any rate, accounts that predict different behavioral patterns for the two types of inferences).

Another influential line of work, starting with Simons (2001), has tied presuppositional status to the discourse structure, modeled explicitly in later work (Beaver et al. 2017) in terms of the Question Under Discussion (QUD Roberts 1996), the main idea being that whatever follows from the QUD will be not *at-issue*, and thus presupposed (in contrast to *at-issue* information that contributes to resolving the QUD). A key property, and potential problem, for this type of approach is that it attributes a central role to context in determining whether a piece of information ends up being presupposed, with no hard-wired lexical split between presuppositional vs. non-presuppositional lexical units. While proponents of this perspective have put forth experimental results suggesting a correspondingly gradient empirical picture of projection (Tonhauser et al. 2018), other recent experimental work (Djäv & Bacovcin 2017, Djäv 2019) finds evidence for clustering of presuppositional vs. non-presuppositional lexical items, with Mandelkern et al. (2019) adding further methodological variants of the relevant paradigms that strengthen the evidence for a more categorical split in what expressions introduce projecting content.

A final line of attack to tackle the explanatory challenge under consideration ties the backgroundedness of presuppositions to whether or not they are at the attentional focus in terms of the temporal unfolding of events (Abrusán 2011, 2016).¹⁴ For *stop*, for example, the idea is that the main reference time introduced by tense concerns the information in (26b)

¹⁴Also see Qing et al. (2016) for a broadly similar approach.

(of not running), whereas the part in (26a) is about a preceding time, which is not part of the main events attended to relative to the sentence. The backgroundedness of this information is then seen as the basis for it projecting, though projection itself is delegated to other mechanisms operating on backgrounded content.

In sum, the question of how presuppositional content winds up having its special status is subject to a lively and ongoing debate, with some relevant experimental work and potential for much more. One aspect that is shared by most of the approaches considered here is that they intend to only deal with certain types of presuppositions, but not others. We turn to some related questions about differences in the lexical nature across presupposition triggers in the next section.

5.2 Differentiating classes of presupposition triggers

One central topic in the recent literature has been the apparent need to distinguish different sub-classes of presupposition triggers (and many of the pragmatic approaches to the triggering problem above only are intended to account for one of these). A particularly influential line of argument from Abusch (2002) draws on contrasts such the following (adapted from Abusch):

- (30) John will either attend the first meeting, or miss it.
- a. ?? And he will either attend the second meeting too, or miss the

second meeting too.

- b. And he will either stop attending meetings, or stop missing them.

The version with the presupposition trigger *too* seems infelicitous, apparently due to the conflict between the two presuppositions (of attending/missing the first meeting), both with one another and the global context. In contrast, the very much parallel presupposition of *stop* does not give rise to the same issue, because it allows for a non-projecting interpretation (where the notion that he previously attended/missed meetings is interpreted within the respective disjuncts). Abusch terms triggers of the latter sort, which allow for non-projecting readings relatively easily, ‘soft’ triggers, and the former, which resist non-projecting readings, ‘hard’ triggers.

These and similar observations have led various authors to attempt to construe differences in how their presuppositional content is encoded in the lexical structure of the different types of triggers. While some authors (Kripke 2009, Zeevat 1992) distinguish between triggers that have an anaphoric element (such as *too*) from those that don’t, others - such as the various pragmatic approaches discussed in the previous section - are based on the idea that the presuppositional status of the relevant information contributed by soft triggers is pragmatically derived, and not lexically encoded, whereas the lexical entries of hard triggers may encode it more directly.

In another line of work, Klinedinst (2012, 2016) and Sudo (2012) in-

independently present an account of the soft-hard split that assumes separate layers of encoding for entailed and presupposed content. They propose that while hard triggers lexically encode totally independent information on their presuppositional and conventional layers, soft triggers redundantly encode the information from their presuppositional layer on their conventional layer as well (e.g., the information of having done X before presupposed by *stop* is both presupposed and entailed). This redundant encoding then gets called upon to explain the easy availability of non-projecting readings for soft triggers, which by assumption already contribute the presupposition to their local context, even if one ignores the presuppositional layer. In the case of hard triggers, ignoring their presuppositional layer would result in a net loss of information, unless additional steps were taken to add it to the entailed layer (e.g., via local accommodation). Notably, this approach predicts that even though soft triggers normally make a contribution at the presuppositional level, they also necessarily contribute the same piece of information at the conventional level as well.¹⁵

5.3 Connectives and Projection

The question of how presupposition-related information ought to be represented in the lexicon, and whether it should be represented at the lex-

¹⁵Corresponding predictions based on examples by Sudo (2012) for environments involving non-monotonic quantifiers have been experimentally examined by Zehr & Schwarz (2018), with some support for a split between triggers, though the results leave open a number of currently unsolved questions and issues.

ical level at all, is not limited to presupposition *triggers*, but also arises when considering how presuppositions interact with embedding operators to yield (or not yield) projection. In addition to the projection data considered above, a key aspect of presuppositions is that they can be ‘filtered’ (in the terminology of Karttunen 1973) by other parts of the sentence:

- (31) a. If Lola is taking night classes now, then she stopped running.
b. If Lola used to run Marathons, then she stopped running.
- (32) a. Lola is taking night classes now and she stopped running.
b. Lola used to run Marathons and she stopped running.

The respective first parts of these conditionals and conjunctions affect whether the sentences as a whole presuppose that Lola used to run (based on the trigger *stop*) - as the (a)-versions do but the (b)-versions do not. Intuitively this is due to the fact that at the point the trigger is considered, the notion of her having been a runner in the past is already introduced by the other clause. While we cannot go into the details of the formal attempts to capture these patterns - which have concerned a substantial literature for several decades -, a key tenet present from the start is that at least for certain connectives, linear order crucially interacts with the projection pattern, e.g., the reverse order of conjuncts yields a relatively odd sentence:¹⁶

¹⁶Changing the order in conditionals gives rise to even more complications; see Mandelkern & Romoli (2017), Romoli & Mandelkern (2017) for discussion.

(33) ??Lola stopped running and she used to run Marathons.

Since *and* is usually modeled in terms of conjunction of statement logic, which is symmetric, this asymmetry has to be accounted for in some way. The dynamic context change semantics proposed by Heim (1983) builds the asymmetry directly into the lexical entry for *and*, by stipulating that the presuppositions introduced in the second conjunct get evaluated relative to the first conjunct (in combination with the preceding discourse context), whereas those in the first conjunct get evaluated relative to the discourse context only. This has been criticized as lacking in explanatory adequacy, and an influential line of recent work (most notably, Schlenker 2009) has revived a notion first introduced by Stalnaker (1973), namely that the asymmetry results from the linear precedence of the first conjunct. It generalizes this approach further to other relevant connectives by providing a classical logical, and symmetric, account that is combined with a pragmatic notion of local context that unfolds linearly and thus can account for the observed asymmetry.

An important difference between these lexical and pragmatic accounts of the projection properties of connectives such as conjunction is that while the former take the asymmetry to be a hard-wired property, it is natural to treat them as an overridable default on the latter. In other words, pragmatic accounts in principle allow for the content of a linearly second conjunct to be taken into account when evaluating the presupposition of the linearly first conjunct. Several experimental projects have aimed to address

related questions, with Chemla & Schlenker (2012) arguing their data to support the pragmatic view when looking at parallel issues for disjunction, conditionals, and *unless*-clauses. However, more recent experimental evidence homing in on the key case of conjunction (Mandelkern et al. 2019) suggests that speakers generally treat projection from conjunction asymmetrically, and do not resort to considering linearly later information, even when that is the only way to rescue the felicity of a sentence. In particular, they judge the continuation (34a) where *stop* appears after *and* more natural than the minimally different continuation (34b) where *stop* instead appears before *and*. (Embedding the conjunction in a conditional preceded by an explicit ignorance context ensures that the presuppositional properties of *stop* are tested here.)

- (34) I never see Lola run. I don't know if she ever was a runner, but...
- a. if she used to run the Marathon and she stopped running, she needs to find another sports activity.
 - b. if she stopped running and she used to run the Marathon, she needs to find another sports activity.

While further work is needed to reconcile the seemingly conflicting results in the literature by extending this paradigm to other connectives, there clearly is a tension between the drive for explanatory adequacy on the theoretical side and the apparent rigidity of projection asymmetry found here empirically for conjunction.

6 Conclusion and Outlook

We have aimed to provide a broad overview of issues relating pragmatic phenomena arising in language use to questions about what related information has to be construed as being encoded in the mental lexicon. We focused on three concrete areas for more detailed illustration, with particular emphasis on issues pertaining to potential lexical encoding. In the realm of scalar implicatures, we reviewed the debates about whether the alternatives that crucially feature in implicature reasoning have to be represented at the level of the lexicon, in the form of Horn scales. We furthermore discussed the role of scales in the interpretation of gradable adjectives and their lexical representations, also with regard to different types of scales and the potential need for encoding this lexically. Finally, we reviewed presuppositional phenomena, with questions arising about whether the presuppositional status of certain information encoded by presupposition triggers has to be marked as such lexically, or whether it can be derived pragmatically. Additionally, the question of differences between triggers and corresponding differentiation in lexical encoding came up, as well as the issue of whether projection properties of connectives such as conjunction have to be lexically represented. There is, of course, a clear common thread throughout, which is at the core of the semantics-pragmatics interface, and the language-cognition interface more generally: For any given interpretive effect of a linguistic expression, we can ask whether it comes

about through conventional encoding, or whether it can be explained independently, typically in domain-general (i.e., not language-specific) terms. The latter route makes for a more minimal theory, but often leads to challenges in capturing empirical nuances. The former route tends to be more powerful, but risks being overly stipulative theoretically, and faces challenges in accounting for how the relevant conventional encoding is learned. There are, needless to say, many other phenomena in this realm that could have been discussed as well (e.g., implicated presuppositions, to name just one), but we hope to have given the reader a taste of the types of debates in the linguistic study of meaning, broadly construed, and how they relate to highly general questions about lexical representations and general reasoning and cognition; and to have conveyed that this is a rich area where new empirical directions combined with ever more refined theoretical considerations promise many fruitful lines of future work.

References

- Abrusán, M. (2011), ‘Predicting the presuppositions of soft triggers’, *Linguistics and Philosophy* **34**(6), 491–535.
URL: <http://link.springer.com/article/10.1007/s10988-012-9108-y>
- Abrusán, M. (2016), ‘Presupposition cancellation: explaining the ‘soft-hard’ trigger distinction’, *Natural Language Semantics* **24**, 165–202.

- Abusch, D. (2002), Lexical alternatives as a source of pragmatic presuppositions, *in* B. Jackson, ed., ‘Semantics and Linguistic Theory (SALT) 12’, Cornell University, Ithaca, NY, pp. 1–19.
URL: <http://elanguage.net/journals/salt/article/view/12.1>
- Abusch, D. (2010), ‘Presupposition triggering from alternatives’, *Journal of Semantics* **27**(1), 37–80.
URL: <http://jos.oxfordjournals.org/content/27/1/37.abstract>
- Aparicio, H., Xiang, M. & Kennedy, C. (2015), Processing gradable adjectives in context: A visual world study, *in* ‘Proceedings of SALT 25’, LSA and CLC Publications, pp. 413–432.
- Barner, D., Brooks, N. & Bale, A. (2011), ‘Accessing the unsaid: The role of scalar alternatives in children’s pragmatic inference’, *Cognition* **118**(1), 84–93.
- Beaver, D. I., Roberts, C., Simons, M. & Tonhauser, J. (2017), ‘Questions under discussion: Where information structure meets projective content’, *Annual Review of Linguistics* **3**(1), 265–284.
URL: <https://doi.org/10.1146/annurev-linguistics-011516-033952>
- Beltrama, A. (2018), Metalinguistic *just* and *simply*: exploring emphatic exclusives, *in* ‘Proceedings of SALT 28’, pp. 307 – 324.
- Bill, C., Romoli, J. & Schwarz, F. (2018), ‘Processing presuppositions and

implicatures: Similarities and differences’, *Frontiers in Communication*

3.

Bill, C., Romoli, J., Schwarz, F. & Crain, S. (2015), ‘Scalar implicatures versus presuppositions: The view from acquisition’, *TOPOI, Special Issue on “Presuppositions: Philosophy, Linguistics, and Psychology”* pp. 1–15.

Bogal-Allbritten, E. (2012), Slightly coerced: Processing evidence for adjectival coercion by minimizers, in ‘Proceedings of the 48th Meeting of the Chicago Linguistic Society’, Chicago Linguist. Soc. Chicago, pp. 77–92.

Bott, L. & Noveck, I. A. (2004), ‘Some utterances are underinformative: The onset and time course of scalar inferences’, *Journal of Memory and Language* **51**(3), 437–457.

Breheny, R., Klinedinst, N., Romoli, J. & Sudo, Y. (2018), ‘The symmetry problem: current theories and prospects’, *Natural Language Semantics* **26**(2), 85–110.

URL: <https://doi.org/10.1007/s11050-017-9141-z>

Burnett, H. (2014), ‘A delineation solution to the puzzles of absolute adjectives’, *Linguistics and Philosophy* **37**(1), 1–39.

URL: <https://doi.org/10.1007/s10988-014-9145-9>

Chanchaochai, N. & Zehr, J. (2019), On the ambidirectionality of thai mid-

scale predicates: How to get more ‘warm’ by getting less ‘hot’, in ‘Proceedings of SALT 29’.

Chemla, E. & Schlenker, P. (2012), ‘Incremental vs. symmetric accounts of presupposition projection: an experimental approach’, *Natural Language Semantics* **20**(2), 177–226.

URL: <http://www.springerlink.com/content/x1830j0872140612/abstract/>

Chemla, E. & Singh, R. (2014a), ‘Remarks on the experimental turn in the study of scalar implicature, part i’, *Language and Linguistics Compass* **8**(9), 373–386.

URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/lnc3.12081>

Chemla, E. & Singh, R. (2014b), ‘Remarks on the experimental turn in the study of scalar implicature, part ii’, *Language and Linguistics Compass* **8**(9), 387–399.

URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/lnc3.12080>

Chierchia, G. & McConnell-Ginet, S. (1990), *Meaning and Grammar. An Introduction to Semantics*, MIT Press, Cambridge, MA.

Cresswell, M. J. (1976), The semantics of degree, in B. H. Partee, ed., ‘Montague Grammar’, Academic Press, New York, pp. 261–292.

URL: <http://www.sciencedirect.com/science/article/pii/B9780125458504500157>

Djävrv, K. (2019), Factive & assertive attitude reports, PhD thesis, University of Pennsylvania.

- Djävrv, K. & Bacovcin, H. (2017), ‘Prosodic effects on factive presupposition projection’, *Proceedings of SALT 27* pp. 116–133.
- Fox, D. & Katzir, R. (2011), ‘On the characterization of alternatives’, *Natural Language Semantics* **19**(1), 87–107.
URL: <https://doi.org/10.1007/s11050-010-9065-3>
- Frazier, L., Clifton, C. J. & Stolterfoht, B. (2008), ‘Scale structure: processing minimum standard and maximum standard scalar adjectives.’, *Cognition* **106**(1), 299–324.
- Frege, G. (1892), On sense and reference, in P. Geach & M. Black, eds, ‘Translations from the Philosophical Writings of Gottlob Frege’, Blackwell, Oxford, pp. 56–78.
- Gazdar, G. (1979), *Pragmatics: implicature, presupposition and logical form*, New York : Academic Press.
- Grice, H. (1975), Logic and conversation, in P. Cole & J. L. Morgan, eds, ‘Syntax and Semantics’, Vol. 3, Academic Press, New York, pp. 41–58.
- Grodner, D. J., Klein, N. M., Carbary, K. M. & Tanenhaus, M. K. (2010), ‘“Some,” and possibly all, scalar inferences are not delayed: Evidence for immediate pragmatic enrichment’, *Cognition* **116**(1), 42–55.
URL: <http://linkinghub.elsevier.com/retrieve/pii/S0010027710000788>
- Heim, I. (1983), On the projection problem for presuppositions, in M. Bar-

- low, D. Flickinger & N. Wiegand, eds, 'Proceedings of WCCFL 2', Stanford University, pp. 114–125.
- Heim, I. & Kratzer, A. (1998), *Semantics in Generative Grammar*, Blackwell, Malden and Oxford.
- Hirschberg, J. L. B. (1991), *A theory of scalar implicature*, New York: Garland Pub.
- Horn, L. (1972), On the Semantic Properties of Logical Operators in English, PhD thesis, UCLA.
- Horn, L. (1989), *A Natural History of Negation*, University of Chicago Press.
- Karttunen, L. (1973), 'Presuppositions of compound sentences', *Linguistic Inquiry* 4(2), 169–193.
- Katsos, N. & Bishop, D. V. (2011), 'Pragmatic tolerance: Implications for the acquisition of informativeness and implicature', *Cognition* 120(1), 67–81.
URL: <http://linkinghub.elsevier.com/retrieve/pii/S0010027711000758>
- Katzir, R. (2007a), 'Structurally-defined alternatives', *Linguistics and Philosophy* 30(6), 669–690.
URL: <https://doi.org/10.1007/s10988-008-9029-y>

- Katzir, R. (2007*b*), ‘Structurally-defined alternatives’, *Linguistics and Philosophy* **30**(6), 669–690.
URL: <https://doi.org/10.1007/s10988-008-9029-y>
- Kennedy, C. (2007), ‘Vagueness and grammar: the semantics of relative and absolute gradable adjectives’, *Linguistics and Philosophy* **30**, 1–45.
- Kennedy, C. & McNally, L. (2005), ‘Scale structure, degree modification, and the semantics of gradable predicates’, *Language* **81**(2), 345–381.
- Kennedy, L., Bill, C., Schwarz, F., Crain, S., Folli, R. & Romoli, J. (2015), Scalar implicatures vs presuppositions: The view from broca’s aphasia, *in* ‘Proceedings of NELS 40’, GLSA, Amherst, MA.
- Kim, M. (2016), It’s not what you said, it’s what you didn’t say: How comprehenders use context to model alternative speaker utterances in pragmatic inferencing, B.A. thesis, Swarthmore College.
- Klinedinst, N. (2012), Totally hardcore semantic presuppositions. Ms., UCL.
- Klinedinst, N. (2016), Two types of semantic presuppositions, *in* K. Allan, A. Capone & I. Kecskes, eds, ‘Pragmemes and Theories of Language Use’, Springer International Publishing, Cham, pp. 601–624.
URL: https://doi.org/10.1007/978-3-319-43491-9_31
- Kripke, S. A. (2009), ‘Presupposition and anaphora: Remarks on the formulation of the projection problem’, *Linguistic Inquiry* **40**(3), 367–386.

URL: <http://dx.doi.org/10.1162/ling.2009.40.3.367>

- Ladusaw, W. A. (1980), *Polarity Sensitivity as Inherent Scope Relations*, Garland, New York, NY.
- Lassiter, D. (2010), Gradable epistemic modals, probability, and scale structure, in ‘Proceedings of SALT’, Vol. 20, pp. 197–215.
- Lassiter, D. & Goodman, N. (2013), ‘Context, scale structure, and statistics in the interpretation of positive-form adjectives’, *Semantics and Linguistic Theory* **23**, 587–610.
- Lauwers, P. & Willems, D. (2011), ‘Coercion: Definition and challenges, current approaches, and new trends’, *Linguistics* **49**(6), 1219 – 1235.
URL: <https://www.degruyter.com/view/journals/ling/49/6/article-p1219.xml>
- Levinson, S. C. (2000), *Presumptive meanings: the theory of generalized conversational implicature*, Language, speech, and communication, MIT Press, Cambridge, Mass.
- Mandelkern, M. & Romoli, J. (2017), ‘Parsing and presupposition in the calculation of local context’, *Semantics & Pragmatics* **10**(17).
- Mandelkern, M., Zehr, J., Romoli, J. & Schwarz, F. (2019), ‘We’ve discovered that projection across conjunction is asymmetric (and it is!)’, *Linguistics and Philosophy*.

URL: <https://doi.org/10.1007/s10988-019-09276-5>

Matsumoto, Y. (1995), ‘The conversational condition on horn scales’, *Linguistics and Philosophy* **18**(1), 21–60.

URL: <https://doi.org/10.1007/BF00984960>

McNally, L. (2011), The relative role of property type and scale structure in explaining the behavior of gradable adjectives, *in* R. Nouwen, R. van Rooij, U. Sauerland & H.-C. Schmitz, eds, ‘Vagueness in Communication’, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 151–168.

Papafragou, A., Friedberg, C. & Cohen, M. L. (2018), ‘The Role of Speaker Knowledge in Children’s Pragmatic Inferences’, *Child Development* **89**(5), 1642–1656.

URL: <http://doi.wiley.com/10.1111/cdev.12841>

Partee, B. H. (1973), ‘Some structural analogies between tenses and pronouns in English’, *The Journal of Philosophy* **70**(18), 601–609.

Potts, C. (2005), *The Logic of Conventional Implicatures*, Oxford Studies in Theoretical Linguistics, Oxford University Press, Oxford.

Qing, C., Goodman, N. D. & Daniel, L. (2016), A rational speech-act model of projective content, *in* A. Papafragou, D. Grodner, D. Mirman & J. C. Trueswell, eds, ‘Proceedings of the 38th annual meeting of the Cognitive Science Society (CogSci-2016)’, Cognitive Science Society, Austin, TX, pp. 1110 – 1115.

- Roberts, C. (1996), Information structure in discourse: Towards an integrated formal theory of pragmatics, in J. H. Yoon & A. Kathol, eds, ‘Papers in Semantics’, number 49 in ‘OSU Working Papers in Linguistics’, The Ohio State University, Columbus, OH, pp. 91–136.
- Romoli, J. (2015), ‘The presuppositions of soft triggers are obligatory scalar implicatures’, *Journal of Semantics* **32**, 173–219.
- Romoli, J. & Mandelkern, M. (2017), Hierarchical structure and local contexts. To appear in *Sinn und Bedeutung 21*.
- Romoli, J. & Schwarz, F. (2015), An experimental comparison between presuppositions and indirect scalar implicatures, in F. Schwarz, ed., ‘Experimental Perspectives on Presuppositions’, Springer International Publishing, Cham, pp. 215–240.
- Sapir, E. (1944), ‘Grading, a study in semantics’, *Philosophy of Science* **11**(2), 93–116.
- Sassoon, G. & Toledo, A. (2011), ‘Absolute and relative adjectives and their comparison classes’, *Unpublished manuscript*.
- Sauerland, U. (2012), ‘The computation of scalar implicatures: Pragmatic, lexical or grammatical?’, *Language and Linguistics Compass* **6**(1), 36–49.
- Schlenker, P. (2009), ‘Local contexts’, *Semantics and Pragmatics* **2**.

- Sevi, A. (2005), A Semantic Account of 'quantity' Implicatures, PhD thesis, Tel-Aviv University.
- Siegel, M. (1979), Measure adjectives in montague grammar, *in* S. Davis & M. Mithun, eds, 'Linguistics, Philosophy and Montague Grammar', University of Texas Press, Austin, pp. 223–262.
- Simons, M. (2001), On the conversational basis of some presuppositions, *in* R. Hastings, B. Jackson & Z. Zvolenszky, eds, 'Proceedings of SALT 11', CLC Publications, Ithaca, NY, pp. 431–448.
- Skordos, D. & Papafragou, A. (2016), 'Children's derivation of scalar implicatures: Alternatives and relevance', *Cognition* **153**, 6–18.
URL: <https://linkinghub.elsevier.com/retrieve/pii/S0010027716300944>
- Stalnaker, R. (1973), 'Presuppositions', *Journal of Philosophical Logic* **2**(4), 447–457.
- Stalnaker, R. (1974), Pragmatic presuppositions, *in* M. K. Munitz & P. K. Unger, eds, 'Semantics and Philosophy', New York University Press, New York, pp. 197–213.
- Strawson, P. F. (1964), 'Intention and convention in speech acts', *The Philosophical Review* **73**(4), 439–460.
- Sudo, Y. (2012), On the Semantics of Phi Features on Pronouns, PhD thesis, Massachusetts Institute of Technology, Cambridge, MA.

- Swanson, E. (2010), ‘Structurally defined alternatives and lexicalizations of xor’, *Linguistics and Philosophy* **33**(1), 31–36.
URL: <https://doi.org/10.1007/s10988-010-9074-1>
- Tonhauser, J., Beaver, D. I. & Degen, J. (2018), ‘How Projective is Projective Content? Gradience in Projectivity and At-issueness’, *Journal of Semantics* **35**(3), 495–542.
URL: <https://doi.org/10.1093/jos/ffy007>
- Trinh, T. & Haida, A. (2015), ‘Constraining the derivation of alternatives’, *Natural Language Semantics* **23**(4), 249–270.
URL: <https://doi.org/10.1007/s11050-015-9115-y>
- Unger, P. (1975), ‘Ignorance: A case for scepticism’, *Mind* **86**(343), 456–458.
- Uppili, H. (2018), Alternatives Across Multiple Scales, B.A. thesis, Swarthmore College.
- van Tiel, B., Van Mittenburg, E., Zevakhina, N. & Geurts, B. (2014), ‘Scalar diversity’, *Journal of Semantics* **33**(1), 137–175.
- von Fintel, K. (1994), Restrictions on quantifier domains, PhD thesis, University of Massachusetts Amherst, Amherst, MA.
- Westerstahl, D. (1984), Determiners and context sets, *in* J. van Benthem & A. ter Meulen, eds, ‘Generalized Quantifiers in Natural Language’, Foris, Dordrecht, pp. 45–71.

Zeevat, H. (1992), ‘Presupposition and accommodation in update semantics’, *Journal of Semantics* **9**(4), 379–412.

URL: <http://jos.oxfordjournals.org/content/9/4/379.short>

Zehr, J. & Schwarz, F. (2018), Returning to non-entailed presuppositions again, in U. Sauerland & S. Solt, eds, ‘Proceedings of Sinn und Bedeutung 22’.