Processing presuppositions: dynamic semantics *vs* pragmatic enrichment.[*]

Emmanuel Chemla (Institut Jean Nicod & LSCP, France)
Lewis A. Bott (Cardiff University, UK)

Abstract: *One defining and yet puzzling feature of linguistic presuppositions is the way they interact with linguistic operators. For instance, when a presupposition trigger (e.g., realise) occurs under negation (e.g., Zoologists do not realise that elephants are mammals), the sentence is most commonly interpreted with the same global presupposition (elephants are mammals) as if negation was not present. Alternatively, the presupposition may be locally accommodated, i.e., the presupposition may become part of what is negated. In this paper, we develop and test two processing accounts of presupposition projection, the global-first model and the local-first model, inspired by dynamic semantics and pragmatic theories respectively. We tested these predictions using a verification task similar to Bott and Noveck's (2004) test of default models of scalar implicature. Across two experiments, using different materials and instructions, participants were faster to derive the global interpretation than the local interpretation, in contrast to the local-first model. We discuss the results in terms of dynamic semantics vs. pragmatic models of presupposition projection (e.g., Heim, 1983b vs. Schlenker, 2008).*

Keywords: *presupposition; processing; scalar implicatures; pragmatics; dynamic semantics*

Linguists and philosophers have long recognised that natural languages offer the means to distinguish between the main point of an utterance, its *assertive content*, and information that should be considered as background information for the participants to the conversation, so-called *presuppositions*. The large set of linguistic expressions or constructions that are said to "trigger" a presupposition includes definite descriptions (*the*), change of state predicates (*stop, start, continue*), additive particles (*too*), pseudo-clefts (*it is X who…*), or, importantly for this paper, factive verbs such as *realise, discover, know, regret, ignore*. For example, each of the sentences below has an *assertive* component and a *presupposition* component.

(1)   a. The king of Moldavia is wise.
        b. John stopped smoking.
        c. Mary got married too.
        d. It is Helen who killed Bob.
        e. Michael discovered Jenna was having an affair.

(1a) *asserts* that the king of Moldavia is wise, but *presupposes* that Moldavia is a monarchy; (1b) asserts that John does not smoke, but presupposes that he used to; (1c) asserts that Mary got married, but presupposes that someone else also got married; (1d) presupposes that Bob was killed, and asserts that the killer was Helen; (1e) asserts that Michael became aware that Jenna was having an affair, but presupposes that she was. In this paper we derive and test processing predictions about how presuppositions interact with linguistic operators such as negation.

In the next section we describe the general phenomenon under investigation, known as the *projection problem*, followed by linguistic accounts of the problem. We then show that natural cognitive implementations of these theories lead to different processing predictions. Finally we present two experiments which offer the first behavioural results that address these issues and show how to adjudicate between these processing accounts.

Presupposition projection

Presuppositions behave differently from assertions in several ways. One such difference is how they behave when presupposition triggers are embedded in more complex sentences, such as under negation, or modals. This complex interaction between presupposition and linguistic operators has been studied as the *projection problem for presupposition*. Let us start with an example to illustrate. Given what may be assumed about negation, presuppositions interact with negation in a surprising way. Consider examples (2) and (3). Because (3) is the negation of (2), the assertive components of the two sentences are opposite. But they have the same presupposition, that elephants are mammals. The same would be true for each of the examples in (1): adding a negation would not affect their presupposition. More generally, linguistic operators – negation, conjunction (e.g., *and, but, because*), disjunction (e.g., *or*), quantifiers (e.g., *some, many, all*), etc. – do not act on presuppositions as it would be expected from their action on the assertive component of a given phrase.

> (2) Zoologists realise that elephants are mammals.
> - Presupposition: *elephants are mammals.*
> - Assertion: *Zoologists are aware that elephants are mammals.*
> (3) Zoologists do not realise that elephants are mammals.
> - Presupposition: *elephants are mammals.*
> - Assertion: *Zoologists are NOT aware that elephants are mammals.*

This interaction is studied as the *projection problem,* which is, as Heim (1983a) puts it, "the problem of predicting the presuppositions of complex sentences in a compositional fashion from the presuppositions of their parts". As illustrated above, presuppositions in complex sentences (e.g., with a negation) do not behave in a straightforward manner and the projection problem seeks to explain this behaviour. At issue is the compositionality of language, i.e. the standard way in which we construct sentence meaning from the meaning of individual parts of the sentence. Not surprisingly, such a fundamental problem has attracted a lot of attention since it was discovered but, more interestingly, there is still no accepted account of the phenomena (see Beaver and Geurts, to appear, for a recent overview). In the next section we discuss two types of influential linguistic accounts of presupposition projection and try to give a sense of the current theoretical considerations that animate the debates.

Linguistic accounts of presupposition projection

Dynamic semantics models (e.g., Heim, 1983b) claim that the meaning of any expression has two components: one for presupposition and one for assertion. A linguistic operator may thus act differently on the presupposition and on the assertion of the linguistic material X to which it applies. The projection problem is solved by setting appropriately the presuppositional component of the relevant operators. For example, negation is given a lexical entry as in (4), which recovers the pattern we described above: negation is active only at the level of the assertion.

> (4) negation $[X]$ = Presupposition of X

and
not [Assertion of X]

In this view, presupposition is a *semantic* phenomenon in the sense that presuppositions of sub-sentential linguistic constituents are computed along with other grammatical, recursive computations of meaning.

In recent years, this semantic approach has been criticised[1] and alternative *pragmatic* models of presupposition projection have been developed to offer new formal solutions (of particular importance for our purposes are Simons, 2004, Schlenker, 2008, Chemla, 2008a, 2010, Abusch, 2010, Romoli, 2011; see Schlenker, 2010 for an overview). In these views, negation is given a more traditional, non-dynamic lexical entry, i.e., a lexical entry which does not distinguish between presuppositional and non-presuppositional aspects of the meaning of the constituent it applies to:

(5)   negation $[X]$ = not [Meaning of X]

Presuppositions are thus silent at the semantic level at which, e.g., negation applies; they come into play when general rules of conversation are considered, and these rules become relevant only after the recursive computation of meaning has been terminated. The differing behaviour of assertions and presuppositions under operators is therefore not determined at a lexical level (as in (4)), but by conversational pragmatics.

There are many different versions of the semantic and pragmatic theories that we have outlined above and our experiments do not relate directly to any one of them in particular. Instead, our experiments seek to distinguish between two classes of theories: those theories that make presupposition projection obligatory and necessary (mostly *semantic* theories), and those theories that make it a process occurring after other more fundamental semantic computations (mostly *pragmatic* theories). We present below an overview of Schlenker (2008), to give a sense of the pragmatic considerations and tools that may be recruited in the second class of theories.

A recent example of a pragmatic theory
Schlenker (2008) seeks to explain presupposition projection as manner implicatures, in Grice's sense (Grice, 1967). The core idea is that the information conveyed by presuppositional phrases is dense, e.g., *realise* conveys both that the proposition expressed in its complement is true and that its subject argument holds it as true. Hence, phrases or sentences containing a presupposition trigger, e.g., (6)a, are claimed to evoke more 'articulated', albeit semantically equivalent alternative phrases, e.g., (6)b.[2] Now, if the

---

[1] While being highly influential, dynamic semantics has been criticised on conceptual grounds because the result it aims to achieve is stipulated in enriched lexical entries, thus failing to provide an *explanatory* solution to the projection problem and merely offering a framework well-suited for a systematic *description* of the facts (see e.g., Soames, 1989, for early criticism). In fact, not only would negation have to be revisited and receive a new lexical entry as above, but all other standard meanings of otherwise standard operators like conjunction, disjunction, quantifiers would have to be modified on a case by case basis to fit the presupposition data.

[2] For readers familiar with quantity or scalar implicatures, (6)b plays the role of what would be the stronger alternative. The claim here is that presuppositions are *manner* implicatures: the alternative is preferable although not because it is more informative.

speaker chose not to utter the more articulated form, e.g., 6(b), it can be assumed that the current conversational context is such that there is no advantage in using the more articulated form. In our example, this inference amounts to the fact that the truth of the information in the complement of the verb *realise* is agreed on. In short, an utterance of the condensed version (6)a will trigger the implicature that the additional explicitness of its competitor (6)b is unnecessary, which happens if participants to the conversation agreed on the truth of the alleged presupposition.

(6) a. Zoologists realise that elephants are mammals.
     b. Elephants are mammals and Zoologists realise that <elephants are mammals>.

A detailed formalisation of this competition between explicitness and briefness leads to a general pragmatic explanation of presupposition projection.[3] In general, the condensed version of the sentence will only be acceptable when the more explicit form would not make a different contribution than the shorter one. When applied to sentences involving negation, this principle dictates that presuppositions escape negation, without lexical specification about negation.

We have presented Schlenker's theory as an example of how pragmatic accounts explain presupposition projection. However, the details of the theory are not important for the present purposes. What matters is that under pragmatic views, presuppositions come out as a negotiation between different maxims, exactly as for other types of implicatures (e.g., scalar implicatures, e.g., Horn, 1972, are explained via maxims of quantity and relevance). This means that the explanations arise without any specific stipulations about the lexical properties of negation or any other linguistic operator (see the criticisms of dynamic semantics presented in Footnote 1), and that presuppositions only arise at a stage at which pragmatic considerations become relevant, that is, after the literal meaning of the sentence has been generated.

We next describe how the semantic and pragmatic theories derive multiple interpretations of presupposition sentences under negation, a phenomenon we use to test between processing instantiations of the different types of theories.

Presupposition under negation, another reading: local accommodation or *literal meaning?*
Both dynamic semantics and pragmatic accounts predict that a negative sentence like (3) has the same presupposition as its positive counterpart (2). However, under certain circumstances, the presupposition of a negative sentence seems to disappear. This is apparent in discourses like (7):

(7) Of course, zoologists do not realise that elephants are birds, because elephants are *not* birds!

---

[3] The interested reader is referred to the original theory but we provide some more details here. The theory predicts that a sentence such as $F(pp')$, where a phrase pp' with presupposition p and assertion p' is embedded in an environment F(…), will presuppose the following: $\forall x, F(p \text{ and } x) \Leftrightarrow F(x)$. The idea is that the condensed sentence $F(pp')$ is acceptable only if a more explicit, or articulated sentence of the form $F(p \text{ and } x)$ (the left-hand side of the equivalence) would not make a different contribution than a shorter one such as $F(x)$ (the right-hand side of the equivalence). (The quantification over xs can be understood as an abstraction away from the assertion or from the material following the occurrence of the presupposition). For instance, the presupposition of *negation (pp')* would be $\forall x, not\ (p\ and\ x) \Leftrightarrow not\ (x)$. In particular, with x instantiated as the tautology T, this leads to *not (p and* T) $\Leftrightarrow$ *not (*T), i.e. p is true. Hence, the result that a sentence and its negation have the same presupposition can be retrieved.

What appears to be happening in (7) is that both the assertion and the *elephants are birds* presupposition are being negated. The presuppositional clause is therefore interpreted as meaning something like *it is not the case that (elephants are birds and zoologists know this)*. This contrasts with the standard behaviour in which the presupposition escapes negation, as in (4). Negative sentences as in (3) or (7) can therefore be said to have a non-presuppositional interpretation in which, descriptively, the presupposition remains trapped under negation, as schematised in (8):

(8)   negation [Z realise p] = *not* [p *and* Z believe p]

In dynamic semantics, this non-presuppositional interpretation is explained by means of a process called *local accommodation*. According to this process, the presupposition of a constituent *X* may be treated as a genuine part of the assertion if something goes wrong with the application of the regular dynamic negation (4) to *X*. In (7), the application of the routine rule (4) leads to a contradiction in which the sentence would convey that *elephants are birds*, despite the fact that everyone knows that it is not true to begin with and that the opposite is actually asserted right after. To resolve this situation, the presupposition of *X* is cancelled, i.e., it becomes part of the normal asserted meaning. Descriptively, then, the presupposition is interpreted locally in (7) (i.e., under negation where it is triggered) and we will therefore refer to this interpretation as the *local reading*. The other interpretation, in which presupposition escapes negation and is interpreted at the level of the whole sentence, will be referred to as the *global reading*. This terminology is summarised in Table 1:

Table 1. Different readings for negative sentences with a presupposition, and schematic predictions of different approaches

	Global reading (Presupposition escapes negation and is interpreted at the *global* level)	Local reading (Presupposition is interpreted under negation, at the *local* level where it was triggered)
Negation of S	PRESUPPOSITION OF S AND NOT [ASSERTION OF S]	NOT [PRESUPPOSITION OF S AND ASSERTION OF S]
Zoologists do not realise that elephants are mammals	Elephants are mammals AND NOT[Zoologists believe that]	NOT [Elephants are mammals AND Zoologists believe that]
Semantic accounts	Core reading → Altered reading	
Pragmatic accounts	Pragmatic reading ← Literal meaning	

In the pragmatic approach, the non-presuppositional, or local, reading described in (8) is straightforwardly explained: it corresponds to the literal meaning of the sentence as predicted by (5), the standard lexical entry for negation, before pragmatic processes apply.

Hence, it is predicted to be an available reading for the sentence, although one that should be superseded when pragmatic processes come into play and enrich it into the global reading.

Processing models

Both dynamic semantics and the pragmatic approach can accommodate the different readings available for negative sentences but, crucially, they suggest different processing models. The dynamic semantics account suggests that the initial meaning of presupposition sentences with negation is the global interpretation, i.e., the presupposition escapes negation. This follows from the dynamic lexical entry for the negation operator (as in (4)). Only if this interpretation is judged unacceptable is the local interpretation (as in (8)) derived.[4] Consequently we refer to the processing instantiation of the semantic account as the global-first model.

To illustrate the predictions of the global-first model, consider sentence (3), *Zoologists do not realise that elephants are birds*. Under the global-first model, the presupposition trigger, *realise*, causes the negation operator to apply differently to the assertion and the presupposition, as in (4). This results in the global interpretation, something like *elephants are birds and Zoologists realise that elephants are birds*. Because the presupposition *elephants are birds* is false in this case, the listener may however, look for a more charitable interpretation of the sentence. One solution is to weaken the role of the faulty presupposition and to treat it as an asserted part of the sentence. This would lead to the local interpretation, *it is not the case that* [*elephants are birds and Zoologists believe so*].

The pragmatic account, on the other hand, assumes a standard lexical entry for negation, shown in (5), which means that the local interpretation is derived first: it is the literal meaning of the sentence in this view. The global interpretation thus comes out through pragmatic processes that take as input the local reading, and hence should unfold as a later (re)interpretation. The pragmatic processing account will therefore be referred to as the local-first model. Consider (3) again, *Zoologists do not realise that elephants are birds*. Under the local-first model, deriving the literal meaning of the sentence involves applying the negative operator to the sentence as a whole (see (5)). This results in the local interpretation, something like, *it is not the case that* [*elephants are birds and Zoologists believe so*]. However, pragmatic maxims may then be applied to the literal meaning of the sentence (i.e., the local interpretation), which results in the global interpretation.

Our main aim in this article is to derive and test processing models of presupposition projection. In doing so we make two standard assumptions about processing and pragmatics. These are (i) that the literal meaning of the sentence is the input to pragmatic procedures, and (ii) that the literal meaning is accessible prior to the application of pragmatic procedures. Given these assumptions, we claim that the pragmatic and dynamic semantics representational theories naturally lead to two distinct alternatives, the global-first and local-first hypotheses. It may be possible to defend different processing implementations of the dynamic and pragmatic theories, but we believe that the burden of the proof would then be on the independent arguments for such re-interpretations. We return to this issue in the General Discussion.

From a methodological perspective, we capitalise on the fact that a similar situation has been identified and studied in the domain of scalar implicatures (e.g., Horn, 1972), another

---

[4] This implementation is particularly warranted in so-called *cancellation* theories. Such theories solve the projection problems by means of late local accommodation processes, which wipe out problematic global presuppositions coming from a primary derivational step and help deliver contradiction-free presuppositional outputs, see e.g., Gazdar (1979).

phenomenon at the interface between semantics and pragmatics. Our experiments borrow from this literature and test the presupposition projection models using a paradigm with a logic similar to Bott and Noveck (2004), one of the pioneer studies in the domain of scalar implicatures which has now led to numerous refinements.

Experiment 1

Our approach was to compare processing times for global and local interpretations of sentences like (3). If people derive global interpretations and then potentially revert to local interpretations, as in the global-first account, processing times should be shorter for global interpretations than local interpretations. Conversely, the local-first account suggests that there will be an early derivation of the local, bare semantic interpretation of the sentence from which the global interpretation is constructed: local interpretation times should be shorter than global interpretation times. (Note that the amount of time needed to derive either of the interpretations individually is unimportant, as is the process by which either interpretation is derived; all that matters is the serial nature of the local- and global-first hypotheses.)

Participants completed a sentence verification task. The experimental sentences all involved a negated factive verb (*realise*) with a complement that generated a false presupposition (e.g., *elephants are birds*). The experimental sentences therefore had the following global and local interpretations (see also Table 1):

(9)   Zoologists do not realise that elephants are birds.
a.   Global: *[Elephants are birds] and not [zoologists believe so]*   (false)
b.   Local: *NOT [(Elephants are birds) AND (zoologists believe so)]* (true)

Global accommodation interpretations should therefore generate a false response whereas local accommodation interpretations should generate a true response. All things being equal, a local-first account consequently predicts shorter response times to true responses than false responses whereas a global-first account predicts the reverse pattern. We included experimental sentences, as described above, and four types of control sentences. These are shown in Table 2. Control sentences (b) and (c) were needed to ensure that participants were unable to predict the correct response prior to reading the final word (participants were given a cover story designed in particular to make the (c) condition unambiguously true) and control sentences (d) and (e) were needed to estimate any bias against false responding in general (see e.g., Clark & Chase, 1972).

Method

Participants

Thirty-three Cardiff University students participated for course credit. Three participants were removed for poor performance on the control sentences (see below).

Design

Each sentence was formed using *Geographers* or *Zoologists* as the subject, a factive verb with negation (*do not realise*) or a non-factive verb (*tell*), and a proposition about categories as the complement of the verb (e.g., *elephants are mammals* or *elephants are birds*), as in (10) below.

(10)   {Zoologists / Geographers} {do not realise / were told} that *subcategory* are *supercategory*.

We generated 60 place names and 60 animals as the subordinate category member of the category proposition. These exemplars formed the basis of each item in the design. The experimental sentences were generated using an exemplar and an incorrect supercategory, combined with the appropriate professional (zoologists or geographers), as in (9). Four control versions of each item were formed using the same exemplar but with a different subject or a different superordinate category (i.e., the correct superordinate) to obtain unambiguously true and false sentences, with the presuppositional phrase (*do not realise*) and without a presuppositional phrase (*were told*). Table 2 shows the five versions of the *elephant* item.

Table 2. Example sentences and observed accuracy.

Condition	Example sentences (Exp 1)	Expected answers (Exp 1)	Observed "true" proportions (Exp 1)	Observed "correct" proportions (Exp 2)
(a)	Zoologists do not realise that elephants are reptiles.	True or False	.38 (.32)	.36 (.31)
(b)	Zoologists do not realise that elephants are mammals.	False	.12 (.073)	.14 (.09)
(c)	Geographers do not realise that elephants are mammals.	True	.85 (.094)	.83 (.11)
(d)	Zoologists were told that elephants are mammals.	True	.93 (.052)	.91 (.07)
(e)	Zoologists were told that elephants are reptiles.	False	.11 (.079)	.11 (.10)

Note. Condition (a) provides the experimental sentences. Conditions (b), (c), (d) and (e) are control sentences. Standard deviations in parentheses.

No subordinate category member was used more than once but superordinate category members were used multiple times. There were four superordinate geographical categories (Africa, Asia, Europe, America) and six superordinate zoological categories (birds, dogs, fish, insects, mammals and reptiles). Each superordinate appeared equally often across conditions.

Items were assigned to five counterbalancing lists (distributed equally among participants) so that all items appeared equally often in each condition, but no participant saw the same exemplar twice. In all, each participant saw 24 items in each condition.

Procedure

Participants were given a cover story to remove ambiguity about the situation, controlling especially the knowledge of zoologists and geographers. The cover story described an alien invasion of Earth in which different groups of aliens were trained to have specialist knowledge of Earth geography but no knowledge of Earth zoology (the geographers) or *vice versa* (the zoologists). This scenario allowed us to construct unambiguously true and false control sentences, with and without presupposition triggers. Participants also went through a

training phase in which they judged 24 control sentences and received feedback on their responses. No experimental sentences were presented during the training phase and participants did not receive feedback during the main part of the experiment.

Sentences were presented one word at a time in the centre of the screen. Each word was presented for 200ms except for the last word, which remained on the screen until the participant made their response.

<div align="center">Results</div>

Data treatment

We removed three participants who scored less than .75 proportion correct on the control conditions. Responses with RTs greater than 10s were removed as outliers (18 out of 3600 data points) and the RT data were log-transformed to reduce positive skewness and inhomogenity of variance. All reported means and standard deviations correspond to the raw data.

Choice proportions

Table 2 shows response proportions for all five conditions. Responses to control sentences were accurate $M = .89$, $SD = .047$, and demonstrate that participants had no difficulty understanding the cover story. Consistent with the co-existence of global and local derivations, the experimental sentences prompted a high degree of bivocality. Whilst there was a slight bias away from the local interpretation, $M = .38$, there was significantly greater variation in the experimental sentences than any of the control sentences, F's > 40, p's < .005.

Response Times

Figure 1 shows the RTs as a function of the sentence type and the response type for the experimental sentences. Global responses (false) seem faster than local responses (true), consistent with a global-first account. We analysed this difference in two ways. First, we compared global and local RTs within each participant (and item). This analysis revealed that global interpretations to the experimental sentences were faster than local interpretations, $M = 2.47s$ ($SD = 1.37$) vs $M = 3.47s$ ($SD = 1.08$) $t_1(26) = 4.78$, $p < .005$. $\eta^2 = .47$ (three participants responded univocally and were therefore excluded from this analysis), $t_2(102) = 4.37$, $p < .005$, $\eta^2 = .16$ (17 items were excluded for the same reason). Second, we classified participants (and items) as local or global responders. Participants were ranked on the proportion of local responses they made to experimental sentences. The top half of participants were then classified as local responders and the bottom half as global responders. Consistent with the within-subject analysis, false responses from the global responders were faster than true responses from the local responders, $M = 2.75$ ($SD = .55$) vs $M = 3.23s$ ($SD = 1.11$) $t_1(28) = 3.86$, $p < .005$, $\eta^2 = .35$, $t_2(118) = 7.49$, $p < .005$, $\eta^2 = .32$. Overall then, the results are consistent with the global-first account.

Figure 1. Response time as a function of sentence type.
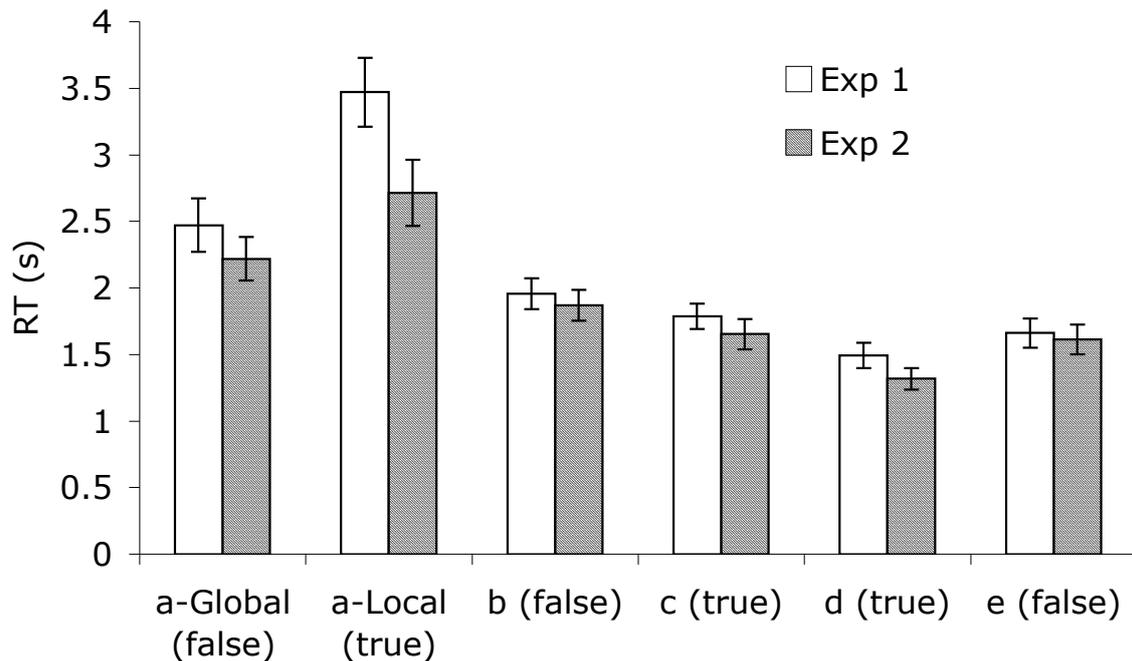
Note. Letters a, b, c, d, and e, correspond to the sentence types shown in Table 1. Condition (a) is broken down by response choice but all other conditions show RTs to correct responses only. Error bars refer to standard errors.

No comparable effect of answer (true *vs* false) was found on the RTs for the control conditions. Yet, one potential explanation for the slow local responses is that participants were slowed down by the *inconsistency* between the veracity of the embedded proposition (false) and the veracity of the local interpretation (true). We tested this by comparing local RTs to the experimental sentences against correct responses to (b) control sentences, in which the veracity of the embedded proposition was also contrary to the veracity of the correct response choice (false). Local interpretation RTs to (a) were significantly longer than correct RTs to (b), however, $M = 3.49s$ ($SD = 1.37$) *vs* $M = 1.95s$ ($SD = 0.63$), $t_1(27) = 7.24$, $p < .0005$, $\eta^2 = 0.66$, $t_2(102) = 9.15$, $p < .0005$, suggesting that long local RT cannot be due entirely to the inconsistency between embedded proposition and response choice.

A similar account of our results is that participants who responded globally to the experimental sentences may have done so because they were focusing entirely on the embedded proposition and not integrating material that occurred earlier in the sentence. This strategy falls short when confronted with the (b) control sentences, however, which require integrating the earlier part of the sentence with the proposition to generate the correct answer. While most participants scored very accurately on the (b) sentences, it is possible that the RT effect was driven by the few who didn't. If participants were using this strategy, we would expect low accuracy on the (b) sentences to be related to high global scores on the experimental sentences. In the event however, the correlation was very low, $r(30) = -.099$, and did not approach significance, $p = .60$. Participants were therefore integrating all parts of the sentence whether they responded globally or locally.

Discussion

Experiment 1 demonstrated that responding true to presupposition sentences such as (8) takes longer than responding false. This applies when responses are compared within participants and when participants are classified as either true responders or false responders and the correct response compared across participants. If a true response implies local interpretations, as in (8a), and a false response implies global interpretations, as in (8b), our results provide evidence against a local-first processing account of presupposition accommodation. Nonetheless, it is possible that when participants were responding to the experimental sentences they may have been accepting or rejecting the sentences on grounds other than the local or global interpretations. For example, participants could have responded "false" because they were unable to derive a meaningful interpretation of the sentence rather than because they derived the global interpretation and believed that it was a false sentence. Fast "false" responses could therefore have been a quick rejection of a sentence that did not make sense. We tested this explanation by conducting a similar RT experiment to Experiment 1 and collecting offline interpretation data about the experimental sentences.

Experiment 2

The logic of Experiment 2 was similar to that of Experiment 1. Participants saw the same five types of sentences that we used in Experiment 1 and we compared RTs to different interpretations of the experimental sentences. We introduced three changes however. First, because we wanted to improve the generality of our findings, we used the verb *know* rather than *realise* as the factive verb in conditions (a) to (c). An example of an experimental sentence might therefore be *Zoologists do not know that elephants are birds*. Second, we gave participants more instruction on the response choices to avoid any ambiguity in the terms "true" or "false." Specifically, participants were now instructed to say whether the sentences "correctly describe the world in the story or incorrectly described the world in the story."

The third and most important change was that we gave participants a post-experiment questionnaire on how they interpreted the experimental sentence. They were presented with the sentence *Zoologists do not know that elephants are birds* and instructed to say whether it correctly or incorrectly described the world in the story and why that was. They were given four possible choices: (1) "Incorrectly describes the world. The statement doesn't make sense." (2) "Incorrectly describes the world. Elephants are not birds so the statement can't be true. It doesn't matter whether zoologists know it or not." (3) "Correctly describes the world. Zoologists don't know that elephants are birds because elephants are not birds. Zoologists don't know things that aren't true." (4) "None of the above. Please write down whether you think the sentence is correct or incorrect and explain your reason." Our goal in introducing the questionnaire was to be able to remove participants who responded with options (1) or (4) from the analysis. We could then be sure that the remaining participants had access to the global (2) or local (3) interpretations when they responded "correct" or "incorrect" in the online task.

Method

Participants.

Forty-five Cardiff University students participated for course credit or payment. Six participants were removed for poor performance on the control sentences (see below).

Design, procedure and materials.

Participants first completed the online task and then the interpretation questionnaire. The online task was identical to Experiment 1 except that participants used a different response choice, as described in the Introduction, and conditions (a) to (c) were formed using *know* and not *realise*.

The questionnaire presented participants with an experimental sentence, "Zoologists do not know that elephants are birds," and asked, "Does this sentence correctly describe the world in the story, or does it incorrectly describe the world in the story? Please read each of the options below and tick the option that best describes your answer. Please tick only one box." The four response possibilities, as described above, were presented in a random order for each participant except for the "None of the above" option, which was always presented last.

Results

Interpretation Questionnaire

Participants ticked one of four boxes corresponding to their chosen interpretation of the experimental sentence. Seven participants chose the incorrect-nonsense interpretation (option 1, above), 22 participants the global interpretation (option 2), 14 participants chose the local interpretation (option 3), and 2 chose "none of the above" (option 4).

Online data treatment

We first removed all participants who chose either the nonsense option or the "none of the above" option during the offline interpretation questionnaire (9 participants). This meant that all of the analysed participants gave interpretation judgments consistent with either the local or the global interpretation. We then removed a further 6 participants who scored less than 0.75 proportion correct on the control conditions, as we did in Experiment 1. Overall 30 participants were included in the final analysis, 12 of whom responded with a local interpretation in the questionnaire and 18 with a global interpretation.

To test whether participants responded in a similar way to the questionnaire as they did in the online tasks we compared the choice proportions of the 12 local questionnaire participants against the choice proportions of the 18 global questionnaire participants. As expected, participants in the local group responded with "correct" interpretations to the experimental sentences more often than participants in the global group, $M = 0.62$ ($SD = 0.26$) *vs* $M = 0.18$ ($SD = 0.18$), $t(28) = 5.41$, $\eta^2 = 0.51$. There were no differences between groups in the other conditions, all $t(28)$'s < 1.9, p's > 0.05.

Responses with RTs greater than 10s were removed as outliers (5 out of 3600 data points) and the RT data were log-transformed to reduce skewness.

Choice Proportions

The far right column of Table 2 shows the mean accuracy and standard deviations for sentences used in Experiment 2. The cover story was understood correctly, as shown by high performance on the control sentences, $M = 0.87$ ($SD = 0.096$), and responses to the experimental sentences were more variable than those on the control sentences, $M = 0.36$, F's > 40, p's $< .005$. Overall the choice proportions were very similar to those of Experiment 1.

Response Times

The right hand bar of each of the conditions in Figure 1 shows the RTs for Experiment 2. The ordinal pattern of RTs across conditions is identical in both experiments but there is a general trend for shorter RTs in Experiment 2. Of principle concern was the comparison between local and global responses. As in Experiment 1, global interpretations were faster than local interpretations $t_1(26) = 2.23$, $p = .035$, $\eta^2 = 0.16$, $t_2(114) = 2.84$, $p < .005$, $\eta^2 = .066$, providing further evidence against a local-first hypothesis. Similarly, dividing the participants into local and global responders revealed that correct responses from local responders were slower than correct responses from global responders, $t(28) = 2.26$, $p = .032$, $\eta^2 = 0.15$, and equivalently for the global and local items, $t_2(118) = 4.3$, $p < .001$, $\eta^2 = 0.14$.

We also checked whether difficulty in using a response choice that was inconsistent with the truth of the embedded propositions could have explained the slow local responding. Just as in Experiment 1 however, local responses were slower than the (b) control condition (in which there was also an inconsistency between the embedded proposition and the response choice), $t_1(26) = 5.37$, $p < .001$, $\eta^2 = 0.53$, $t_2(114) = 4.10$, $p < .001$, $\eta^2 = .13$. Nor did we find evidence that participants were failing to integrate the embedded proposition with the rest of the sentence. The correlation between proportion local responding and accuracy in the (b) condition was larger than in the previous experiment but not significant, $r = 0.30$, $p = 0.11$, and, indeed, in the opposite direction to the predicted explanation.

The findings of Experiment 2 replicate those of Experiment 1. In both experiments we found that the local responses needed more time than global responses, contrary to a local-first account. This means our findings generalise to a different response type, "correctly describe the world" vs. "incorrectly describe the world," and, more importantly, our results cannot be due to participants who responded "false" because they thought the sentence did not make sense. In Experiment 2 we removed all participants who could not make a clear global or local interpretation and local responses remained slower than global responses.

Although the qualitative pattern of the results was replicated across experiments, the size of the difference between local and global responses was much smaller in Experiment 2 than Experiment 1. This can be seen in Figure 1 by comparing condition (a) across experiments, and by noting the effect size difference in the participant analysis between global and local responses, $\eta^2 = 0.47$ *vs* $\eta^2 = 0.16$. One potential explanation for this is that by removing participants who responded with nonsense interpretations in the questionnaire we removed a large proportion of the RT variance across interpretations in the experimental condition. However, when we reinstated the six participants who made nonallowable interpretation judgments but who still had accuracy on the control items above 0.75, the effect size decreased (to 0.10) rather than increased. The larger effect size is therefore unlikely to be due to the presence of unwanted participants in Experiment 1, whose answers would have been excluded from Experiment 2 thanks to the questionnaire.

Another difference between the two experiments is that Experiment 2 used *know* whereas Experiment 1 used *realise*. It is possible in principle that presuppositions triggered by different items have a different status, leading to differences in the relative difficulty to

compute presuppositional interpretations (see for instance the distinction between factives and semi-factives proposed at least since Karttunen 1971). However, notice that we obtained very similar rates of derivations of the respective interpretations in the two experiments, reducing the likelihood of an essential difference between the presuppositions triggered by *know* and *realise* which could explain our difference in RT effect size. Finally, given the overall trend to lower RTs in Experiment 2, the differences across experiments could be due to participant sampling differences. While we cannot say which of these possibilities is responsible for these differences, it is clear that local interpretations are more complicated whether *realise* or *know* is used as the presupposition trigger, neither of which is predicted by a local-first hypothesis.

Discussion

The goal of our study was to understand how presuppositions are computed and processed. To this end we distinguished local-first and global-first processing models. We relied on two specific current formal accounts of presupposition projection to show that these processing models are natural candidates. We argued that these processing models – and thus their corresponding theoretical approaches – make different predictions about false presupposition under negation, e.g., *Zoologists do not realise that elephants are reptiles*. The processing version corresponding to the semantic account predicts that people should initially apply the negation operator differently to the assertive and presupposition components of the sentence, as the lexical entry in (4) specifies, so that the global interpretation should be derived first. Only if this application fails would the local interpretation be derived. For this reason we associated the semantic view with the *global-first* processing model. In contrast, the processing version corresponding to the pragmatic account predicts the reverse order: the local interpretation should be derived first because the lexical entry of the negation operator makes no distinction between assertion and presupposition (unlike the semantic account) and it is only later, after pragmatic processes have been applied, that the global interpretation should be derived.

Using a sentence verification task with a similar logic to Bott and Noveck (2004), we found that participants were faster to derive a global, rather than a local, accommodation interpretation for negative sentences with a presupposition. This result applied across two experiments in which we varied the verb and the response choice. Our results therefore provide evidence against a two-step serial processing model of presupposition in which a local interpretation is computed first and then supplanted by a global interpretation.

Possible limitations due to the task

One potential objection to our design is that participants were required to use different response labels for the different interpretations, i.e., local interpretations were indicated with *true* responses and global interpretations with *false* responses. It is therefore possible that our results reflect processes involved in the mapping from a single interpretation to multiple judgment labels. One argument might be that people have a bias towards responding *false* and are therefore slow to respond *true*. For example, perhaps *true* responses were longer because participants were looking for ways to reject the statement but "timed out" on a number of occasions. Consideration of our control conditions suggests the reverse, however (see Figure 1): (c)(*true*) responses are faster than (b)(*false*) responses, and (d)(*true*) responses are faster than (e)(*false*) responses. Any bias is likely to be against responding false. Or perhaps participants were fast to respond *false* because the experimental sentences contained a negation and people can sometimes respond *false* more quickly to sentences with negation (Clark and Chase, 1972). This again appears to be counter to the evidence from the control

conditions, however. Participants responded faster to (c)(*true*), which contained a negation, than they did to (b)(*false*), which did not.

A similar argument is that only responses to ambiguous sentences, such as the experimental sentences, exhibit a *false* bias (this could also be phrased in terms of a facilitation for weaker interpretations, see below for discussion). While our data cannot directly address this point, the results from Bott and Noveck (2004), on which we based our current design, are consistent with a bias in the other direction: Participants were slower to respond *false* to similarly ambiguous sentences, such as *some elephants are mammals*, than they were to respond *true*, and there was a higher proportion of *true* responses than *false* responses. Additional explanation would therefore have to be generated as to why some types of ambiguous sentences are susceptible to the bias whereas others are not. Ultimately however, just as with Bott and Noveck, we hope that the present results will be replicated with other methodologies that will exclude task related alternative interpretations (e.g., Breheny, Katsos & Williams, 2005; Huang & Snedeker, 2009).

Another possibility is that longer RTs might be due to general preferences in offline interpretations and not processing *per se*. As is often claimed by theorists (based on introspective judgments), the global reading is the most salient. Perhaps the local reading takes longer because participants required more time to suppress their initial preferences. If so, our data should be seen as experimental evidence of a preference for the global reading. Under this interpretation of our results, how would the linguistic theories we discussed above account for such a preference? Although none of the linguistic theories explicitly make predictions, the most straightforward generalisation is that the preferred reading is the reading that is derived first. People might only go to the effort to derive the interpretation that is easiest, for example, and the easiest interpretation is the first sensible interpretation than can be derived. The global-first and local-first preference accounts are therefore consistent with the global-first and local-first processing accounts respectively, and our data would therefore support a global-first preference account. Nonetheless, we also suggest a stronger interpretation of our results than the preference account suggested above. In our experiments participants made interpretation judgments without any instructions guiding them towards a particular reading (c.f., Bott & Noveck, 2004, Experiment 1, in which participants were trained to respond with a particular interpretation). The response therefore reflected the preference experienced at the time of responding. When the participants responded with the local interpretation, for example, the local interpretation was the preferred response at that time. Thus, the 38% local responses were not biased towards global interpretations, by definition, and the longer local RTs were therefore unlikely to be caused by a simple dispreference for local interpretations.

Finally, the RT pattern could be explained by a greater difficulty in deriving local interpretations as opposed to the processing speed explanations that we have advanced. For example, if local interpretations required a greater memory search to verify than global interpretations, participants would be less able to derive an accurate local interpretation after a given amount of time (Grodner, 2009, argued for a similar explanation of the delay in responding to scalar implicatures). Participants could then have chosen to delay responding to local interpretations in our task so that they could retrieve an accurate judgement. While possible, there are several reasons to suggest otherwise. First, there is no intrinsic reason why the local interpretation should be more difficult to verify than the global interpretation. For example, the slower (local) interpretation does not entail the information contained in the faster (global) interpretation, unlike in the case of scalar implicatures, and more generally, there are no more propositions to verify in the local case than the global case. Secondly, Bott, Bailey and Grodner (submitted) found evidence that the RT delay in scalar implicatures was

not due solely to differences in the difficulty of retrieving the implicature interpretation. They implemented a speed-accuracy-tradeoff design (SAT; McElree, 1993) in which participants were prevented from delaying their responses by being forced to respond at fixed deadlines (50ms, 100ms, 150ms etc.) after the onset of the final word. Tomlinson, Bott, & Bailey (2011) came to the same conclusion using a mouse-tracking paradigm. To the extent that the our paradigm and that of Bott and Noveck (2004) share similarities, the results are likely to be explained by similar types of effects. Ultimately however, a generic difficulty explanation can only be eliminated by using an SAT design, and our results may reflect difficulties in derivation rather than processing speed *per se* (but note that how the linguistic theories under discussion predict "difficulty" may well turn out to be equivalent to the predictions above for "preference").

Implications for theories of presupposition projection

The global-first and local-first models were inspired by dynamic semantics and pragmatic accounts respectively, as we described above. Our results therefore argue against a class of theories which includes one processing implementation of the pragmatic account, while being consistent with the dynamic semantics account as presented above. The extent to which the pragmatic account is inconsistent with our results depends on three representational and processing assumptions. First, the pragmatic account we have described assumes a nondynamic lexical entry for linguistic operators. Clearly, a dynamic entry in conjunction with pragmatic machinery would make the same predictions as the semantic account. However, since the pragmatic component of the explanation would then be redundant, these accounts might be rejected on grounds of parsimony (and we know of no published pragmatic accounts that assume dynamic representations combined with pragmatic procedures like those above).

Second, we have claimed that pragmatic theories predict that local readings would be faster than global readings. The reason for this is that, in the pragmatic view, the local reading is necessary to obtain the global reading, because the local reading is the literal meaning that serves as the input of maxims and other pragmatic procedures. The assumption that this situation should give rise to RT advantages is in line with the recent psycholinguistic literature on scalar implicatures (Bott & Noveck, 2004; Breheny, Katsos & Williams, 2006; Huang & Snedeker, 2009), which involves the quantity maxim. However, it is conceivable that *some* maxims of conversation and pragmatic apparatus might be applied *automatically* to generate the global interpretation, and then, on some occasions, implicatures might be defeased to generate the local, literal interpretation (in the psycholinguistic literature about scalar implicatures, a similar default view of pragmatic processes is often associated with neo-Gricean accounts, e.g., Levinson, 2000, even though this association is highly disputable from a theoretical perspective). This would mean longer interpretation times for the global reading, as we observed here. While there is converging evidence that scalar implicatures based on the maxim of quantity are not derived in this way, there is no reason why other maxims may not be applied by default to generate implicatures (e.g., the manner maxims, as used in Schlenker, 2008). At the very least, our results firmly prove that there is a difference between the way scalar implicatures and presuppositions are obtained. This result goes immediately against recent attempts to reduce the differences between the two phenomena (Abusch 2010, Chemla 2008a, 2010, Romoli, 2011). To sum up, we relied on what we see as the default hypothesis that comparable pragmatic processes apply at the same stage in sentence interpretation. Yet, a complete assessment of the pragmatic theory we approached would require testing cases where the specifically relevant maxims are uncontroversially at work, to check whether these maxims are applied by default.

Finally, we have assumed that the different accommodation interpretations are generated serially in the pragmatic (and the semantic) models so that one interpretation must be rejected before the other interpretation. The reason that this assumption is important is that all of the linguistic models that we know of predict that derivation of the local (or global) interpretation is contingent on the rejection of the global (or local) interpretation. Any model that assumes this type of contingency predicts serial processing models by nature of the information needed to advance processing. Models that do no assume contingency are not restricted to serial processing, and could account for our findings as long as the local interpretation is more difficult to derive or to verify. However, we know of no accounts of presupposition projection that are implemented explicitly in this way and there is no reason *a priori* to assume that there could be such a facilitation for the global reading.

In addition to the processing and representational assumptions discussed above, our conclusions are necessarily restricted to the class of presupposition triggers tested in our experiments, namely factive verbs (*realise* and *know*). One may therefore worry about the validity of our conclusions for the whole class of presupposition triggers. For example, anaphoric triggers show a different behaviour when it comes to accommodation (see, e.g., Beaver and Zeevat 2007, van der Sandt and Geurts 2001, van der Sandt and Huitink 2003) and it may be that processing results would be different for those types of triggers. We thus believe that, indeed, these results should be replicated with a wider variety of presupposition triggers. Nonetheless, we claim that our results provide a solid first step. First, despite the differences described between different sub-classes of presupposition triggers, we are aware of no account that would distinguish them on processing grounds. Hence, there is no theoretical reason to guide investigations of possible processing differences for different sub-classes of triggers. Second, and more specifically, factive verbs may belong to a sub-class of presupposition triggers which, descriptively, easily allow local accommodation readings. We found that local accommodation was *not* the easy processing option for these factive verbs. Hence, there is little reason to doubt that this particular conclusion would not extend to anaphoric triggers, or other triggers that barely allow local accommodation.

Overall, our results require that a complete theory of presupposition projection include some process which explains why global interpretations arise faster than local interpretations. On the one hand, dynamic semantics and more generally so-called cancellation theories have a readily available explanation: the global interpretation is the root interpretation of the sentence. On the other hand, further efforts are needed to defend processing implementations of implicature-like, pragmatic theories of presupposition.

Conclusion

While there have been psycholinguistic studies of how presupposition sentences are interpreted (see, e.g., pioneering acquisition studies by Harris, 1975, and Hopmann & Maratsos, 1977; more recently Schwarz, 2007, processing results on *auch* (*too*); Klein et al., 2009, and Modyanova, 2009, and references therein on definite descriptions; Chemla, 2009, and Chemla & Schlenker, 2010, for offline data), there has been no systematic exploration of the processing aspects of presupposition *projection*. Our contribution to the debate is to describe a first plausible implementation of existing formal theories and to offer online results which should constrain these presupposition theories and their psycholinguistic implementations. We suggested in particular that our data distinguish between certain reconstructions of pragmatic and semantic theories, giving an advantage here to dynamic semantics over its implicature-like competitors. The advantage mostly comes from the

assumption that the pragmatic view naturally fits with a delay in the derivation of global presuppositions, while the cost we see is associated with *local* presuppositions.

More generally, the study presented here opens the way for a range of investigations, which should help understand the psycholinguistic aspects of presuppositions, and of other inferences. The target sentences of our experiments have been constructed so as to share the key property of the material used in the psycholinguistic literature about scalar implicatures (Bott & Noveck, 2004): their semantic value changes depending on whether their presuppositions are taken into account (see also Chemla 2009 for a larger comparison between presuppositions, scalar implicatures and inferences due to adverbial modifications which led to similar, although preliminary processing results). We hope that the methodologies developed recently in the growing literature in experimental pragmatics to study scalar implicatures (e.g., Breheny, Katsos & Williams, 2005; Huang & Snedeker, 2009; Noveck, 2001) will now be applied to confirm and extend the present results.

References

Abbott, B. (2000). Presuppositions as nonassertions. Journal of Pragmatics 32(10), 1419–1437.

Abusch, D. (2010). Presupposition Triggering from Alternatives. Journal of Semantics, 27(1), 37-80.

Beaver, D. & Geurts, B. (to appear). Presupposition. In E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*.

Beaver, D. I. and H. Zeevat (2007). Accommodation. In G. Ramchand and C. Reiss (Eds.), The Oxford Handbook of Linguistic Interfaces., pp. 503–538. Oxford University Press.

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

Bott, L., Bailey, T. M., & Grodner, D (submittted). Distinguishing speed from accuracy in scalar implicatures.

Breheny, R., Katsos, N., & Williams, J. (2006) Are generalised scalar implicatures generated by default? An on-line investigation into the role of context in generating pragmatic inferences. *Cognition*, *100*,434–463.

Chemla, E. (2008a) Présuppositions et implicatures scalaires: études formelles et expérimentales. EHESS Dissertation.

Chemla, E. (2008b). Projecting presuppositions with scalar implicatures. In Alte Grønn, editor, Proceedings of SuB 12, pages 81–91.

Chemla, E. (2009) Presuppositions of quantified sentences: experimental data. *Natural Language Semantics*, *17*,299–340.

Chemla, E. (2010) Similarity: Towards a unified account of scalar implicatures, free choice permission and presupposition projection. Under revision for *Semantics and Pragmatics*. Shorter version in Chemla (2008b).

Chemla, E., & Schlenker, P. (2010) Incremental vs. Symmetric Accounts of Presupposition Projection: An Experimental Approach. Ms. LSCP, IJN & NYU.

von Fintel, K. (2004) Would You Believe It? The King of France is Back! Presuppositions and Truth-Value Intuitions. In M. Reimer & A. Bezuidenhout (eds.) Descriptions and Beyond, Oxford University Press.

Gazdar, G. (1979a). Pragmatics: Implicature, Presupposition and Logical Form. New York: Academic Press.

Grice, H. P. (1967). *Logic and conversation*. the William James Lectures, delivered at Harvard University.

Grodner, D. (2009) Some Reasons to Doubt that 'Some' (and Probably All) Scalar Inferences are Delayed. Talk given in Moving Beyond Truth Conditions: The Computation of Meaning, University of Maryland.

Grodner, D., Klein, N., Carbary, K., & Tanenhaus, M.K. (submitted). Rapid interpretation of pragmatic some. *Cognition*.

Harris, R.J. (1975) Children's comprehension of complex sentences. *Journal of Experimental Child Psychology, 19*, 420-433.

Heim, I. (1983a) On the projection problem for presuppositions. In M. Barlow, D. Flickinger & M. Westcoat (eds.), Second Annual West Coast Conference on Formal Linguistics, Stanford University, 114–126.

Heim, I. (1983b) File Change Semantics and the Familiarity Theory of Definites. In R. Baüerle, C. Schwarze & A. von Stechow (eds.), Meaning, Use and Interpretation of Language, De Gruyter, Berlin, 164–189.

Hopmann, M. R. & Maratsos, M.P. (1977) A developmental study of factivity and negation in complex syntax. *Journal of Child Language, 5*, 295–309.

Huang, Y. & Snedeker, J. (2009) Online interpretation of scalar quantifiers: Insight into the semantics–pragmatics interface. *Cognitive Psychology, 58*, 376–415.

Karttunen, Lauri. (1971). Some Observations on Factivity. Papers in Linguistics 5, 55–69.

Klein, N.M., Gegg-Harrison, W., Sussman, R.S., Carlson, G. & Tanenhaus, M.K. (2009) Weak definites: rich, but not strong, special, but not unique. In U. Sauerland & K. Yatushiro (eds.) Studies in Pragmatics, Language, and Cognition, Palgrave, 264–275.

Modyanova, N. (2009) Semantic and Pragmatic Language Development in Typical Acquisition, Autism Spectrum Disorders, and Williams Syndrome with Reference to Developmental Neurogenetics of the latter. MIT Dissertation.

McElree, B. (1993). The locus of lexical preference effects in sentence comprehension: A time-course analysis. *Journal of Memory and Language, 32*, 536-571.

Noveck, I.A. (2001) When children are more logical than adults: experimental investigations of scalar implicature. *Cognition, 78*,165–188.

Romoli, J. (2011). The presuppositions of soft triggers aren't presuppositions. Proceedings of SALT XXI.

van der Sandt, R. and B. Geurts (2001). Too. In Proceedings of the 13th Amsterdam Colloquium. University of Amsterdam.

van der Sandt, R. and J. Huitink (2003). Again. In P. Dekker and R. van Rooij (Eds.), Proceedings of the 14th Amsterdam Colloquium, pp. 181–186. Amsterdam: ILLC.

Schwarz, F. (2007) Processing Presupposed Content. *Journal of Semantics 24*,373–416.

Schlenker, P. (2008) Be Articulate: A Pragmatic Theory of Presupposition Projection. *Theoretical Linguistics, 34*, 157–212.

Schlenker, P. (2010) Presupposition Projection: the New Debate. In Proceedings of Semantics and Linguistic Theory (SALT) 18.

Simons, M. (2001) On the Conversational Basis of some Presuppositions. In Hastings, R., Jackson,B. and Zvolenzky, S. (eds.) Proceedings of SALT 11. Cornell University: CLC Publications.

Simons, M. (2004) Presupposition and Relevance. In Z. Gendler Szabó (ed.) Semantics vs. Pragmatics, Oxford University Press.

Soames, Scott. (1989). Presupposition. In Gabbay & Guenthner (Eds.), Handbook of philosophical logic (Vol. IV, pp. 553–616). Dordrecht: Reidel.

Stalnaker, R. (1972) Pragmatics. In D. Davidson and G. Harman (eds.), Semantics of Natural Language, Dordrecht, Reidel, 389–408.

Stalnaker, R. (1973) Presuppositions. *Journal of Philosophical Logic 2*, 447–457.

Stalnaker, R. (1974) Pragmatic presuppositions. In M. Munitz & P. Unger (eds.), Semantics and Philosophy, New York University Press, 197–214.

Tomlinson, J., Bott, L. & Bailey, T.M. (2011). Understanding literal meanings before pragmatic inference: Mouse-trajectories of scalar implicatures. Paper presented at the 4th Biennial Conference of Experimental Pragmatics, June 2-4.