THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Pragmatics: From Theory to Experiment and Back Again

Citation for published version:

Katsos, N & Cummins, C 2010, 'Pragmatics: From Theory to Experiment and Back Again', *Language and Linguistics Compass*, vol. 4, no. 5, pp. 282-295. <https://doi.org/10.1111/lnco.2010.4.issue-5>

Digital Object Identifier (DOI):

[10.1111/lnco.2010.4.issue-5](https://doi.org/10.1111/lnco.2010.4.issue-5)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Language and Linguistics Compass

Publisher Rights Statement:

© Katsos, N., & Cummins, C. (2010). Pragmatics: From Theory to Experiment and Back Again: Pragmatics. *Language and Linguistics Compass*, 4(5), 282-295. 10.1111/lnco.2010.4.issue-5

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Pragmatics: From theory to experiment and back again

Napoleon Katsos and Chris Cummins

Research Centre for English and Applied Linguistics, University of Cambridge

Abstract

This paper reviews some cases in which the collaboration of theoretical pragmaticians and psychologists of language has been most fruitful for all parties. Linguists have benefited from experimental data confirming the psychological validity of their observations and providing theory-critical evidence in cases beyond the reach of reflective intuition, while psychologists have benefited from having a wealth of linguistic phenomena to study as well as multiple theories furnished by semantics and pragmatics. Focusing on the much-discussed question of whether scalar implicature is a default inference, we aim to showcase the benefits of the experimental approach to pragmatics.

Keywords: experimental pragmatics, scalar implicature, psycholinguistics

0.1 Introduction

Theoretically-motivated empirical investigations of pragmatic and semantic aspects of meaning are not novel. However, in recent years, interest in experimental pragmatics has intensified, leading to a flourishing literature in both linguistic and psychological journals, comprising overview and position articles (Novick & Reboul, in press; Sedivy, 2007; i.a), a series of meetings and conferences since 1998 and a dedicated website (www.experimentalpragmatics.org, Clark & Noveck, 2005), and a European Science Foundation network grant (Euro-Xprag, see www.euro-xprag.org, Noveck, Geurts & Sauerland, 2009), among much else.

This article does not offer a historical review of previous and current investigations, although such a review might provide a useful sidelight on the evolution of psychological and linguistic approaches to the study of meaning (see Sperber & Noveck, 2004). Rather, we present cases where interdisciplinary investigations between linguistics and psychology have proved mutually beneficial, whilst also highlighting problems to be overcome in this approach.

The discussion of several topics could serve this aim. Studies on reference, speech act, metaphor, irony, and so-called figurative language comprehension have produced a wealth of psycholinguistic findings and informed models of human cognition. Many of these investigations were directly motivated by, and have fed back into, the linguistic theory of the phenomena.

For instance, studies by Clark and colleagues on reference (see Clark & Bangerter, 2004 for a review) indicate that the coordination condition proposed by Grice (1989) and Lewis (1969) is insufficient, and that reference assignment requires more collaboration between speaker and hearer. This work has significance not only in a linguistic context but also to the psychology of language processing and Theory of Mind, the ability to take into account interlocutors' epistemic state (see Barr & Keysar, 2007; Hanna, Tanenhaus & Trueswell, 2003; i.a.). On another front, research on figurative language by Gibbs, Glucksberg and colleagues, among others, has produced evidence against the pervasive notion in psycholinguistics that pragmatics is secondary to syntax and semantics (see chapters 2 and 3 in Noveck & Sperber, 2004 for an overview). Specifically, studies in metaphor comprehension have shown that apparently pragmatic processes of the type suggested by Grice (1989) and Searle (1979)

are not costly, nor effortful, but exhibit the automaticity customarily associated with modular processes (see also Giora, 2003).

However, in this paper we shall focus on *scalar implicature*, a topic that has generated much of the most recent literature. This phenomenon is particularly amenable to empirical investigation for several reasons. First, there is widespread agreement between theoretical accounts of scalar implicature about the interpretation of scalar terms (at least in the most straightforward cases). The predictions that could differentiate between competing accounts are often too fine-grained to be tested with the traditional tools of the theoretical linguist, reflective introspection and intuition. Empirical investigation is therefore a more suitable testing procedure, as argued by Chierchia (2004), Levinson (2000) and Sperber & Wilson (1986/1995), among others. Furthermore, the competing theoretical accounts of implicature have achieved a considerable level of conceptual clarity and precision, making it possible to generate clear predictions about processing and acquisition from the linguistic claims about the nature and availability of scalar implicature. Finally, the proliferation of research on scalar implicatures reflects the importance of this topic to determining the nature of the semantics-pragmatics interface, as well as to key questions in psycholinguistics. Issues such as the locus of implicature, the role of context, and the pragmatic intrusion on truth-conditions have repercussions for the organisation of the whole semantic and pragmatic system, while issues such as the time-course of integrating contextual information and its interaction with grammatically encoded meaning are broadly relevant to psycholinguistics.

This paper focuses in particular on the issue of whether scalar implicatures are default inferences. Despite this comparatively narrow remit, we hope to draw some wider conclusions about how linguistic claims can be associated with experimental findings, and which methodologies can be used to address such claims. However, we first turn our attention to the linguistic phenomena.

1. Scalar implicature and the generalised-particularised distinction

Consider the following conversation:

- (1) The Bride: Have you seen Tommy?
 Bill: Big guy in the tux?
 The Bride: Yes.
 Bill: Then I saw him. I like his hair.
 The Bride: You promised you'd be nice!

(Kill Bill, Volume 2; script by Quentin Tarantino;
Uma Thurman as the Bride; David Carradine as Bill)

Let us further assume that the following inferences can be drawn from Bill's last utterance¹:

- (2) a. Bill doesn't love, adore or worship Tommy's hair.
 b. Bill doesn't like anything else about Tommy besides his hair.
 c. Bill doesn't approve of Tommy as a groom for the Bride.

¹ These inferences are licensed on the grounds that Bill's last utterance is not ironic, i.e. that he does mean he likes Tommy's hair. This is clear in the film, in that there is nothing wrong with Tommy's hair, while Bill has other reasons to disapprove of Tommy.

According to Grice (1989), (2a), (2b) and (2c) are cases of conversational implicature. These pragmatic inferences arise because of certain conventions of rational behaviour: for example, because the interlocutors assume one another to be cooperative and to obey certain rules regulating the quantity and relevance of information that ought to be exchanged.

In all versions of current pragmatic theory, the derivation of (2c) would be (more or less) straightforwardly accounted for by some version of a maxim of relation or relevance. However, the precise derivation of (2a) and (2b) has been the subject of much controversy. In the early neo-Gricean tradition, and particularly by Horn (1972, 1984 i.a.), it was noted that these inferences rely upon the existence of a scale of informativeness. For (2a), the scale of informativeness consists of verbs of liking, <like, love, adore, worship...>. Terms of the right of this scale are informationally stronger than terms on the left, since the former entail the latter, but not vice versa. For (2b), the situation is more complex. An entailment scale could consist of things that one could like about a prospective groom: <{hair}, {personality}, ..., {hair, personality}, {hair, personality, ...}, ...>. Alternatively, the scale for (2b) could involve a person's characteristics, ranked in terms of noteworthiness, e.g. <hair, ..., physique, ..., personality, ...>. In this case there is no entailment relation between the terms of the scale, but we can see how the terms could be ranked according to this (subjective) criterion.

However the scale is formed, the reasoning that derives (2a) and (2b) goes approximately as follows. The interlocutors are interested in establishing what Bill thinks of Tommy. Hence, any information on that point is relevant to the conversation. Now, the speaker said 'I like [Tommy's] hair'. If the speaker wanted to say something more informative, he would have done so. Assuming that he is cooperative and knows what he is talking about (valid as he is expressing his own views), then his failure to say anything more informative implies that he did not want to. That is, the fact that Bill only said he *liked* Tommy's hair implies that he did not feel any more strongly about it, and the fact that Bill only said he liked Tommy's *hair* implies that he did not like anything else about Tommy. Which arguably justifies the Bride's reproach.

Since these inferences are conversational in nature (rather than grammatically encoded), and since the derivation relies upon the existence of informativeness scales, they are traditionally known as *scalar implicatures*. However, within this category, there may be important differences, for instance between (2a) and (2b). Consider the utterance 'I like Tommy's hair' occurring out of the blue, without any conversational context. Arguably, inference (2b) never occurs in this case, while (2a) still seems to be available. This is the difference that Grice (1989: 37) himself had in mind when he introduced the distinction between Particularised and Generalised Conversational Implicatures (henceforth *PCIs* and *GCI*s):

“[PCIs] are cases in which an implicature is carried by saying that *p* on a particular occasion in virtue of special features of the context, cases in which there is no room for the idea that an implicature is normally carried by saying that *p*. But there are cases of generalized conversational implicature. Sometimes one can say that the use of a certain form of words in an utterance would normally (in the absence of special circumstances) carry such-and-such an implicature or type of implicature.”

While the PCI-GCI distinction was explicitly introduced by Grice himself, it should be emphasised that, for Grice, calculating an implicature of either type “is to calculate what has to be supposed in order to preserve the supposition that the Cooperative Principle is being observed” (1989: 39-40).

Thus, both PCIs and GCIs satisfy the criteria for conversational implicature status; for example, they are defeasible, calculable, indeterminate, and non-detachable (see Horn, 1984; Levinson, 1983; i.a.).

It has generally been agreed that conversational implicatures ought to have at least some of these properties (Sadock, 1978), and there is by now ample evidence that scalar implicatures (henceforth *SI*s) do. For example, with regard to defeasibility, it is expected that *SI*s can be cancelled without giving rise to strong contradiction. Katsos (2007; chapter 3) reports several studies where participants are presented with short discourses with an utterance containing a scalar expression followed by an utterance that contradicts either the *SI* or a semantic entailment of that expression. It was consistently found that the continuation contradicting the *SI* was rated as more coherent than that which contradicted the entailment (although not perfectly coherent).

Moreover, there is evidence that *SI*s are indeed context-dependent, and are not generated in conversational contexts where they do not contribute to the participants achieving the discourse goal. As Horn (1984: 13) notes, “[A scalar]...implicatum...may be...cancelled implicitly (by establishing the appropriate contexts, in which all that is relevant, or can be known, is the lower-bound).” (See also Levinson, 2000: 51ff; Sperber and Wilson, 1986/1995: 276ff for similar comments.) Let us call contexts in which “all that is relevant...is the lower bound” *lower-bound contexts*, and contexts in which pragmatic inference is required to achieve the discourse goal *upper-bound contexts*. Then, in the following example (adapted from Levinson, 2000), (3a) is an upper-bound context and (3b) a lower-bound context.

- (3) a. Were all their identity documents forgeries?
 b. Is there any evidence against them?
 c. Some of their identity documents are forgeries.

In answer to (3a), the speaker of (3c) implicates that not all of their identity documents were forgeries. This is a classical scalar implicature: ‘all’ is more informative than ‘some’ in this proposition, and if it were true to say ‘all’, the speaker would have said so, because it would have been relevant (given (3a)). Hence, assuming that the speaker has the knowledge to make the stronger statement, it follows that (3c) gives rise to the above-mentioned *SI*. However, the same utterance (3c), when used as a response to (3b), could receive a different interpretation. The discourse goal set by (3b) – to establish whether there is any incriminating evidence – is satisfactorily addressed by the semantic meaning of (3c), which answers (3b) in the affirmative. In such a context, it is irrelevant whether or not the strengthened alternative proposition containing ‘all’ is true, and hence the *SI* need not be derived. Katsos, Breheny & Williams (2005) report studies where participants who are asked to interpret scalar expressions (in pen-and-paper tasks) are indeed sensitive to the difference between implicature-generating upper-bound contexts and implicature-free lower-bound contexts.

There is also evidence that, in accordance with the Gricean account, interlocutors modify their expectations of informativeness according to each other’s epistemic state. Chierchia, Crain, Guasti & Thornton (1998), Musolino & Lidz (2006) and Katsos & Bishop (2008) have demonstrated that comprehenders only generate *SI*s in conditions where their speaker is known to have all the relevant information. For instance, given a scenario in which ‘all’ is applicable, a speaker’s description of that scenario using the quantifier ‘some’ is consistently rejected on the grounds of underinformativeness. However, if the speaker is making a guess about an action that is yet to take place, or a bet about an action that has happened unseen, then participants are significantly more willing to accept ‘some’ in

the scenarios where ‘all’ applies, demonstrating sensitivity to whether or not the speaker is in a position to make a more informative statement (Katsos & Bishop, 2008).

From this brief review, we can see that the predictions common to all accounts of SI – that such inferences are defeasible, persist only when relevant to the discourse goal, and take into account epistemic state – are empirically supported. This allows us to proceed to the central issue of this paper, namely the default or contextual nature of scalar implicature.

2. Positions on defaultness

In the derivation of generalised implicatures, the concept of a lexical scale of informativeness is crucial. Strong constraints have been proposed on what constitutes a lexical scale: terms must be equally lexicalised, of the same form class, of a similar register, and about the same semantic field (Horn, 1984, 2004; see also Atlas & Levinson, 1981; Levinson, 1983, chapter 3; 2000: 79ff). Lexical scales meeting these criteria are formed by quantifiers, connectives and modals, as in (4a) to (4c), as well as gradable adjectives (4d/e), articles (4f; see Hawkins, 1991), and numerals (4g), among others.

- (4)
- a. <some, many, most, all>, <sometimes, always>, <something, everything>
 - b. <or, and>
 - c. <possible, certain>, <may, must>
 - d. <warm, hot>, <cool, cold>
 - e. <like, love, adore>, <dislike, hate>
 - f. <a, the>
 - g. <one, two, ...>

Alongside this constrained set of lexical scales, there is a practically unlimited range of scales that can be formed in an *ad hoc* fashion for specific conversational settings. This contrast is exhibited by (2a) and (2b) in section 1: the former is context-independent because it relies on a lexical scale, while the latter is context-dependent, because it relies on an *ad hoc* scale that has only been introduced in a specific context. Removing or changing the context stops the generation of the implicature. This is captured by Horn (2004: 3, emphasis in the original), who observes that “it is apparent that some implicatures are induced **only** in a special context [=PCIs]...while others go through **unless** a special context is present [=GCIs].”

Emphasising this distinction, Levinson (2000) and Chierchia (2004) explicitly argue that context is not at all necessary for deriving GCIs. They argue that GCIs are derived by the application of default principles that generate preferred, albeit potentially defeasible, inferences. Contextual factors may then cause these to be cancelled (for example, if the speaker is identified as not fully cooperative or not in a position to make a stronger statement). In Levinson’s (2000) account, GCIs are derived by specialised pragmatic rules that have evolved to maximise the speed of communication. Chierchia (2004: 57) bases his radical proposal “on the idea of adding [SIs] in as soon as possible, as part of the computational system of grammar, rather than as part of the interaction of such a system with other extra-grammatical modules”. In either account, PCIs like (2b) and (2c) would be derived by a distinct context-sensitive Gricean pragmatic system.

By contrast to these *default*² approaches, other accounts collapse the PCI-GCI distinction and maintain that (a) the same pragmatic principles are involved in generating PCIs and GCIs, and (b) certain contextual conditions are required for the derivation of both kinds of implicature. For example, Hirschberg (1991) provides an elaborate survey of lexical scales and ‘posets’, partially-ordered sets of terms including a wide range of relations (whole/part, type/subtype, etc.) and scales based on contextually-salient contrastive relations in a non-entailment ordering such as the one in (2b). In this account, lexical and logical scales are perhaps the most salient kind of scale, but nevertheless unexceptional.

While radically different in the specifics of how implicatures are derived, Relevance theory (Sperber & Wilson, 1986/1995; see also Carston, 1998, 2002; i.a.) also undermines the GCI-PCI distinction. According to this theory, all implicatures are derived by the same strategy, and contextual assumptions are assumed to be equally integral to deriving any type of pragmatic inference. Thus, if an SI is not contextually relevant, it will not be generated. As the Relevance theory account, like Hirschberg’s account and other accounts such as Geurts (1998), assumes a single context-dependent pragmatic system to be responsible for generating all SIs, we will refer to these as *context-driven* accounts.

3. Theory-critical predictions

Let us outline the issues where default and context-driven accounts make contrasting predictions. First, the accounts envisage different roles for discourse context. Context-driven accounts predict that all implicatures are generated only when relevant to the discourse goal, while default accounts predict that this only holds for PCIs, with GCIs also being generated in ‘neutral’ contexts where it cannot be established whether or not the implicature is relevant. In contexts where the implicature is discourse-irrelevant, default accounts – unlike context-driven accounts – further predict that the implicature is first generated and then cancelled.

Secondly, default accounts predict that GCIs are more easily generated than PCIs, for reasons of processing economy. According to Levinson’s account, default heuristics have developed to maximise the speed of communication, hence generating a GCI should be rapid and extra effort could be required when cancelling them. By contrast, Relevance theory-inspired accounts tend to predict that generating any pragmatic inference is costly in terms of processing resources (but do not normally need to be cancelled, hence not giving rise to additional costs – see Bott & Noveck, 2004). However, these predictions are not necessarily shared by all accounts in their respective categories. For example, there is no stipulation in Chierchia’s default account that implicature generation or cancellation should be easy or slow, respectively, merely that SI generation is context-independent and therefore occasionally leads to cancellation and backtracking.

Theories of SI also differ in whether they assume local or global generation of implicatures, and whether these implicatures intrude upon truth-conditions. For now, we leave aside these dimensions of variation and consider the experimental evidence available to address the question of default versus context-driven generation of SIs in terms of the role of context and the time-course of generating SIs.

4. Experimental investigations

² It should be emphasised that by ‘default’ accounts we refer to approaches that propose defaultness at a *pragmatic* level, not on a conceptual level (e.g. Jaszczolt, 2005).

Carston, Levinson, Chierchia and Sperber and Wilson are among those who consider that their accounts should enjoy psycholinguistic as well as linguistic validity, and that their predictions about the availability and generation of SIs should be empirically (dis)confirmable (cf. Levinson, 2000: 5, 81, 162ff, 370; Chierchia 2004: 51, 68, 93; Carston & Powell, 2006 for the Relevance theory perspective). Relevance theory and Levinson's account are actually motivated by cognitive considerations (principles arising from cognitive evolution and the psychology of sentence processing, respectively).

Furthermore, as we shall see in the following section, experimental investigation is necessary because several theory-critical predictions concerning SIs lie outside of the reach of the traditional tools of the linguist – conceptual argumentation, introspection and intuition. At a time when it is widely acknowledged that syntax has run out of 'cheap data' and that a rapprochement of cognitive psychology with theoretical linguistics is highly desirable for both disciplines (see Ferreira, 2005; Marantz, 2005; i.a.), Chierchia (2004: 46) makes precisely this point for SIs.

In the following subsections, we will review theoretically-motivated empirical investigations of SIs, paying particular attention to their contribution to the theoretical debate.

4.1 The role of context

Context-driven and default accounts agree that SIs are ultimately available when relevant to the discourse goal and not available otherwise; however, they differ as to how this situation arises. Default accounts predict that the SI is first generated, regardless of context, then cancelled if it is contextually irrelevant, whereas context-driven accounts predict that the SI is not generated at all if it would not be relevant. Clearly intuition alone cannot settle this question, as we are concerned not with the ultimate interpretation but with the process through which it is derived. However, psycholinguistic investigations of the time-course of scalar implicature can shed light on this.

More generally, a central debate in the area of sentence processing concerns whether there is an initial, encapsulated stage of structure assignment, or whether different categories of information (syntax, semantics, context, co-occurrence frequencies, etc.) interact and coordinate from the earliest possible stage (Spivey-Knowlton & Sedivy, 1995; Tanenhaus, Spivey-Knowlton, Eberhard & Sedivy, 1995; i.a.). The question of how SIs are generated could be seen as part of this debate.

According to the default approach, more processing time should be required when participants are processing scalar terms in lower-bound contexts (where the SI is irrelevant) than in upper-bound contexts (where it is relevant), because in the latter the inference must be generated and then cancelled. This additional processing should manifest itself as a delay in reading time, in self-paced reading studies. Context-driven accounts do not share this prediction, as no cancellation is required; in fact, accounts based on Relevance theory (see Bott & Noveck, 2004; i.a.) predict that reading time will be slower in upper-bound contexts, because an additional inference is generated in these contexts. Thus, by investigating the reading times of scalar expressions in upper- and lower-bound contexts, we can gather evidence about the role of context as well as the automaticity of the inference.

In a series of self-paced reading experiments, Breheny, Katsos & Williams (2006; see also Katsos, 2008; Katsos, Breheny & Williams, 2005) investigated the reading time of scalar expressions in such contexts. For example, for the existential quantifier 'some of the Fs', Breheny et al. (2006) constructed upper-bound (UB) and lower-bound (LB) conditions, by manipulating whether 'all of the

Fs' was relevant to the discourse (8a/c). They included a control condition with 'only' (8b) that explicitly encoded the SI³.

- (8) a. (UB) Mary asked John whether he intended to host all his relatives in his tiny apartment. John replied that he intended to host some of his relatives. The rest would stay in a nearby hotel.
- b. (Control) Mary asked John whether he intended to host all his relatives in his tiny apartment. John replied that he intended to host only some of his relatives. The rest would stay in a nearby hotel.
- c. (LB) Mary was surprised to see John cleaning his apartment and asked the reason why. John told her that he intended to host some of his relatives. The rest would stay in a nearby hotel.

The reading time was measured on the scalar term itself ('some of his relatives') and on a dependent phrase ('the rest') whose interpretation is facilitated by the SI. If the participants interpret 'some' as 'some but not all', then the reference set for 'the rest' is already available when they read the dependent phrase; if they do not, then extra processing time is required to interpret who 'the rest' refers to.

As predicted by all accounts, the dependent phrase was read faster in the UB and control conditions than in the LB. This indicates that the SI was ultimately generated in UB and not in LB, indicating that the conditions were appropriately constructed. To answer the theory-critical question – whether the SI is generated and then cancelled in the LB condition – we have to look to the scalar expression itself. In this study (and in Katsos et al., 2005 for English), the trigger phrase 'some of the Fs' was read more slowly in UB than in LB contexts. This is hard to reconcile with the default account, since no reading time delay was evident in the LB condition. Thus, these findings show that contextual factors have a primary role in the process that generates SIs, not merely in cancelling them at a later stage. (For experimental studies that reach different conclusions, see Bezuidenhout & Cutting, 2002; but see also the discussion of that paper in Breheny et al., 2006).

The second prediction on the role of context concerns 'neutral' situations in which the discourse goal cannot be clearly established. In these cases, default accounts predict generation of the SI, while context-driven accounts do not. Note that single sentence utterances are not necessarily neutral, as they can create their own context through a variety of presupposition triggers and information-structure triggers. For example, in (9), the NP 'Mary' is linguistically marked as focus, i.e. the answer to an implied contextual question.

- (9) It was Mary who ate shark-soup.

If the question 'who ate shark-soup?' is not already salient for the hearer, an utterance of (9) would trigger its accommodation into the context. Similarly, (10a) is more likely to relate to an issue like 'what happened to the consultants?' or 'who did the consultants meet with?' than is (10b).

- (10) a. Some of the consultants had a meeting with the director.
- b. The director had a meeting with some of the consultants.

³ The original experiment was conducted in Greek.

This was controlled for in a further experiment by Breheny et al., conducted in Greek. Greek has a strong tendency for grammatical subjects in sentence-initial position to be construed as old information and post-verbal objects to be construed as unrelated to a contextual issue (see Philippaki-Warbuton, 1987; i.a.) Breheny et al. created two-sentence discourses in Greek in which ‘some of the Fs’ was either a pre-verbal subject (11a) or post-verbal object (11b).

- (11) a. Some of the consultants had a meeting with the director.
The rest did not manage to attend.
- b. The director had a meeting with some of the consultants.
The rest did not manage to attend.

The target phrase (‘the rest’ in the above examples) can be interpreted immediately but requires making a bridging inference that the collection referred to by ‘the rest’ is the rest of the consultants. If the SI ‘not all the consultants’ has been generated on-line in the first sentence, this inference will be facilitated. Given the assumptions about the effect of the word order on accommodating contextual issues, context-driven accounts predict that this will affect the reading time for ‘the rest’. If there is no accommodation of an appropriate contextual issue, such as ‘what happened with the consultants?’, there is no reason for the SI to become available in the neutral context. Thus, reading time should only be facilitated by the trigger being in sentence-initial position. By contrast, default accounts predict that the implicature should be generated on-line regardless of where the implicature trigger is, and hence that sentence position should not affect reading time.

Additionally, control conditions were created in which the inference was explicitly marked with ‘only’, as in (11c) and (11d), which was expected to result in maximally short reading times.

- (11) c. Only some of the consultants had a meeting with the director.
The rest did not manage to attend.
- d. The director had a meeting with only some of the consultants.
The rest did not manage to attend.

The results showed that ‘the rest’ was read equally quickly in the two cases where some was sentence-initial (11a/c): this suggests that the SI was generated in (11a) and the referent set was as readily available as when introduced by ‘only’. However, the reading time for ‘the rest’ was slower in (11b) than in (11d), suggesting that when ‘some of the consultants’ was sentence-final, the contrast set was not as readily available as a referent for ‘the rest’. Since, as we have argued, (11b) is a neutral context for a scalar expression, these findings suggest that the predictions of default accounts are not upheld.

4.2 The time-course of scalar implicature

In a series of studies in French, Noveck and colleagues (Noveck, 2001; Noveck & Posada, 2003; Bott & Noveck, 2004) investigated the time-course of SI generation. They presented participants with utterances containing scalar terms which are logically true but underinformative given readily available world knowledge, such as in (12).

- (12) Some elephants have trunks.

Participants were asked to give a truth-value judgement on such sentences. These would be ‘false’ if, as Noveck argues, they generated the SI (‘some but not all’), or ‘true’ if they accessed the plain meaning of the quantifier ‘some’. In experiment 3 of Noveck (2001), it was found that about 60% of

such utterances are rejected. Noveck & Posada (2003) measured response times to utterances containing a scalar term after the whole utterance had been processed, and found that the time taken by the participants who were instructed to answer ‘false’ was significantly longer than those who were instructed to answer ‘true’. Bott & Noveck (2004) replicated these findings both with respect to the percentage of no responses and the time-course of the two responses, even when substituting the response ‘true’/‘false’ (which might bias for logical interpretations) for ‘I agree’/‘I disagree’ (which might be more relevant to felicity than to logical truth-value judgements).

Bott & Noveck also excluded the possibility that the time difference between the two responses was due to the difference in response type (positive/negative). They showed that no time difference was evident in the corresponding control conditions (accepting or rejecting utterances that are patently true and informative, or patently false). They also showed that the same pattern of results obtained in a more complex experiment in which the utterances were narrated by a character (thus allowing answers of both types to have the same response type). This pattern persisted when all instructions were removed. In a fourth experiment, they found that the number of responses based on the SI increased as the participants were permitted more time to respond.

The studies by Noveck and colleagues yielded evidence that is relevant to two predictions. First, they showed that adult native speakers do not always generate implicatures in simple declarative sentences, as predicted by default accounts. Second, they show that generating an SI is not an automatic process, but rather an effortful one. If, as according to default accounts, generating an SI were automatic and cancelling it required extra resources, then the acceptance of underinformative utterances ought to take longer than the responses where they reject them; moreover, when participants are given less time to respond, then the response pattern that should increase is the acceptance, rather than the rejection, of underinformative utterances. The findings of Noveck and colleagues have been replicated in English using the same methodology (De Neys & Schaeken, 2007, but see also Feeney et al., 2004), and there is also evidence from studies using the visual-world eye-tracking paradigm that corroborates the claim that generating SIs is indeed a time-consuming process (Huang & Snedeker, 2009).

5. Overview and outlook

In the previous section, we reviewed empirical investigations that have been motivated by the debate between default and context-driven accounts of SI. We saw that there is mounting evidence that SIs are only generated in contexts where the inference is relevant to the discourse goal; in other contexts, it is not the case that they have first been generated and then cancelled. These findings are predicted by context-driven accounts and cannot be accommodated by default accounts. Moreover, generating SIs was found to be a time-consuming rather than a fast process, which also tends to cohere with context-driven rather than default proposals, although claims about the time-course of the inference need not be an integral part of default, or context-driven, accounts. The mounting evidence in favour of context-driven accounts that consider pragmatic inferences as extra-grammatical and effortful lends support to accounts similar to the original Gricean proposal, but may also accept other interpretations in future work.

With regard to methodology, we have discussed several paradigms for investigating different aspects of the SI generation process. Pen-and-paper studies are perfectly adequate when we are interested in the eventual interpretation of the expression being studied (e.g. defeasibility of SIs, relevance of epistemic state). However, self-paced reading, eye-tracking and reaction time studies are called for when we are investigating the process through which the eventual interpretation is generated.

In all these cases, the relation between conceptual clarity and successful experimentation can be seen as a two-way street; clear predictions at the theoretical level render a topic straightforwardly amenable to empirical study, while the drive for empirical investigations can underscore the need for clarity at the theoretical level.

We conclude this paper by highlighting further outstanding questions in the debate on SI generation, and discuss how the collaboration of linguists and psychologists could be fruitful in resolving these.

5.1 Why should scalar implicature be effortful?

In section 4.2, we saw that generating an SI is an effortful, costly process. But why should SIs be costly? There is by now ample evidence that pragmatic inferences *per se* are not necessarily costly in processing terms. For example, work on metaphor comprehension and indirect speech acts (Gibbs, 1986; Gibbs & Tendahl, 2006; Glucksberg, 2003; McElree & Nordlie, 1999; i.a.) has shown that drawing the pragmatically appropriate interpretation can be automatic and not involve extra cost when the appropriate interpretation is clear-cut.

So why should SIs *specifically* be costly? Multiple suggestions are available: SIs are metalinguistic inferences requiring the addressee to consider the possible alternative utterances available to the speaker; the negative component of the SI ('p or q but *not* both' for disjunction, 'some but *not* all' for the existential quantifier) adds to processing cost; or, as proposed by Noveck and colleagues, inferences that contribute to making an utterance more informative are generally costly, a cost which is balanced by the gain in informativeness.

Still another proposal is that the SI is not intrinsically costly, but gives rise to costly operations on the mental discourse model. When participants interpret 'Some of the Fs G-ed' without an SI, they introduce the set of Fs who G-ed into their discourse model. When they interpret it with an SI, they have to introduce the set of Fs and also partition it into two non-empty subsets (those who G-ed and those who did not G). The suggestion that the introduction of contrast sets, rather than the SI, induces a processing cost is compatible with two pieces of evidence. Grodner, Gibson & Watson (2005) showed that introducing contrast sets through restrictive relative clauses gives rise to processing costs. It has also been shown that, once a contrast set is established, inferential processes dependent on the second maxim of quantity (requesting speakers not to be over-informative) are initiated at the earliest possible stage, without manifesting any processing cost (Sedivy, 2001; Sedivy, Tanenhaus, Chambers & Carlson, 1999; i.a.). Further research along these lines could shed light on the time-course of SI processing, and whether this is something inherent to scalar reasoning.

5.2 Are there no defaults?

The on-line studies reported in section 4.2 concerns the role of relevance of the SI to the discourse context. A distinct question concerns whether global conversational assumptions about cooperativity, epistemic state, etc. are constructed for each communicative exchange, or whether they are introduced across the board by general principles of communicative behaviour. In the experiments where participants generated SIs, they must have assumed that the speaker was cooperative and had adequate knowledge of the topic under discussion, despite there being no explicit information to that effect.

One possibility is that people, at least in certain communicative settings with fixed social roles, use stereotypical assumptions on whether their interlocutors are cooperative and well-informed on the issues they discuss. This could indicate that there are 'defaults' at the social level, or merely indicate

that the assumptions are context-dependent and constructed from scratch for each communicative exchange.

Testing these assumptions would be another straightforward extension of the research presented in this overview, and would further broaden the scope of the debate by introducing wider sociological considerations on human interaction.

References

- Atlas, J. & Levinson, S. (1981). It-Clefts, Informativeness and Logical Form. In Cole, P. (ed.), *Radical Pragmatics*. New York: Academic Press. 1-61.
- Barr, D. J. & Keysar, B. (2007). Perspective taking and the coordination of meaning in language use. In Traxler, M. J. & Gernsbacher, M. A. (eds.), *Handbook of Psycholinguistics*, 2nd edition. New York: Academic Press. 901-38.
- Bezuidenhout, A. & Cutting, J.C. (2002). Literal meaning, minimal propositions, and pragmatic processing. *Journal of Pragmatics*, 34: 433-456.
- 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.
- 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(3): 434-463.
- Carston, R. (1998). Informativeness, relevance and scalar implicature. In: Carston, R. & Uchida, S. (eds.), *Relevance Theory: Applications and Implications* (pp 179-236). Amsterdam: John Benjamins.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Carston, R. & Powell, G. (2006). Relevance theory – new directions and developments. In Lepore, E. and Smith, B. (eds.), *Oxford Handbook of Philosophy of Language*. Oxford: Oxford University Press.
- Chierchia, G. (2004). Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In Belletti, A. (ed.), *Structures and Beyond*. Oxford: Oxford University Press. 39-103.
- Chierchia, G., Crain, S., Guasti, M. T. & Thornton, R. (1998). "Some" and "or": A study on the emergence of Logical Form. In *Proceedings from the 22nd Annual BUCLD*. Somerville, MA: Cascadilla Press. 97-108.
- Clark, B. & Noveck, I.A. (2005). www.experimentalpragmatics.org/ visited on 11 September, 2009.
- Clark, H. H., & Bangerter, A. (2004). Changing conceptions of reference. In Noveck, I. & Sperber, D. (eds.), *Experimental pragmatics*. Basingstoke, England: Palgrave Macmillan. 25-49.
- De Neys, W., & Schaeken, W. (2007). When people are more logical under cognitive load: Dual task impact on scalar implicature. *Experimental Psychology*, 54, 128-133.
- Feeney, A., Scafton, S., Duckworth, A. & Handley, S.J. (2004). The story of *some*: Everyday pragmatic inference by children and adults. *Canadian Journal of Experimental Psychology* 58:2, 121-132.

- Ferreira, F. (2005). Psycholinguistics, formal grammars, and cognitive science. *Linguistic Review*, 22: 365–380.
- Geurts, B. (1998). Scalars. In Ludewig, P. & Geurts, B. (eds.), *Lexikalische Semantik aus kognitiver Sicht*. Tübingen: Gunter Narr. 95-118.
- Gibbs, R. (1986). *The poetics of mind*. Cambridge: Cambridge University Press.
- Gibbs, R & Tendahl, M. (2006). Cognitive effort and effects in metaphor comprehension: Relevance Theory and psycholinguistics. *Mind & Language*, 21 (3): 379-403.
- Giora, R. (2003). *On our mind: Salience, context, and figurative language*. New York: Oxford University Press.
- Glucksberg, S. (2003). The psycholinguistics of metaphor. *Trends in Cognitive Science*, 7, 92-96.
- Grice, H. P. (1989). *Studies in the Way of Words*. Cambridge MA: Harvard University Press
- Grodner, D., Gibson, E., & Watson, D. (2005). The influence of contextual contrast on syntactic processing: evidence for strong-interaction in sentence comprehension. *Cognition*, 95: 275–296.
- Hanna, J.E., Tanenhaus, M.K., & Trueswell, J.C. (2003). The effects of common ground and perspective on domains of referential interpretation. *Journal of Memory and Language* 49: 43–61.
- Hawkins, J. (1991). On (in)definite articles: implicatures and (un)grammaticality prediction. *Journal of Linguistics*, 27 (2): 405-42.
- Hirschberg, J. (1991). *A Theory of Scalar Implicature*. New York: Garland.
- Horn, L. (1972). *On the Semantic Properties of Logical Operators in English*. UCLA dissertation, distributed by Indiana University Linguistics Club, 1976.
- Horn, L. (1984). Toward a new taxonomy for pragmatic inference. In Schiffrin, D. (ed.) *Meaning, Form and Use in Context: Linguistic Applications, Proceedings of GURT '84*. Washington D.C.: Georgetown University Press. 11-42.
- Horn, L. (2004). Implicature. In Horn, L. and Ward, G. (eds), *Handbook of Pragmatics*. Malden, MA: Blackwell Publishers. 3-28.
- Huang, Y. & Snedeker, J. (2009). Online interpretation of scalar quantifiers: Insight into the semantics–pragmatics interface. *Cognitive Psychology*, 58(3), 376-415.
- Jaszczolt, K. (2005). *Default Semantics: Foundations of a Compositional Theory of Acts of Communication*. Oxford: Oxford University Press.
- Katsos, N. (2007). *Experimental investigations on the effects of structure and context on the generation of scalar implicatures*. PhD thesis, University of Cambridge.
- Katsos, N. (2008). The semantics/pragmatics interface from an experimental perspective: the case of scalar implicature. *Synthese*, 165: 358-401.
- Katsos, N. (2009). Evaluating under-informative utterances with context-dependent and context-independent scales: experimental and theoretical implications. In Sauerland U. & Yatsushiro K. (eds.) *Experimental Semantics and Pragmatics*. Basingstoke: Palgrave Studies in Pragmatics, Language & Cognition. 51-73.
- Katsos, N. & Bishop, D. (2008). The development of informativeness from a speaker's and comprehender's perspective. Paper presented in *IASCL 2008*, Edimburgh, UK.

- Katsos, N., Breheny, R. & Williams, J. (2005) The interaction of structural and contextual constraints during the on-line generation of scalar inferences. In Bara, B., Barsalou, L. and Bucciarelli, M. (eds.), *Proceedings of the 27th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum. 1108-13.
- Levinson, S. (1983). *Pragmatics*. Cambridge: Cambridge University Press
- Levinson, S. (2000). *Presumptive Meanings*. Cambridge, MA: MIT Press.
- Lewis, D. K. (1969). *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.
- Marantz, A. (2005). Generative linguistics within the cognitive neuroscience of language. *Linguistic Review*, 22: 429-445.
- McElree, B., & Nordlie, J. (1999). Literal and figurative interpretations are computed in equal time. *Psychonomic Bulletin & Review*, 6, 486-494.
- Musolino, J. & Lidz, J. (2006). Why children aren't universally successful with quantification. *Linguistics*, 44(4), 817-852.
- Noveck, I. A. (2001). When children are more logical than adults. *Cognition*, 86: 253-282.
- Noveck, I.A., Geurts, B. & Sauerland, U. (2009). <http://www.euro-xprag.org/> visited on 11 September 2009.
- Noveck, I. A. & Posada, A. (2003). Characterizing the time course of an implicature: An evoked potentials study. *Brain and Language*, 85: 203-210.
- Noveck, I. A. & Reboul, A. (in press). Experimental pragmatics: a Gricean turn in the study of language. *Trends in Cognitive Sciences*.
- Noveck, I. A. & Sperber, D. (2004). *Experimental pragmatics*. Basingstoke: Palgrave Macmillan.
- Philippaki-Warbuton, I. (1987). The theory of empty categories and the pro-drop parameter in Modern Greek. *Journal of Linguistics*, 23: 289-318.
- Sadock, J. (1978). On testing for conversational implicature. In P. Cole (ed.), *Syntax and Semantics 9: Pragmatics*. New York: Academic Press: 281-297.
- Searle, J. (1979). *Expression and Meaning*. Cambridge: Cambridge University Press.
- Sedivy, J. (2001). Evidence for Gricean mechanisms in on-line language processing. Paper presented at the *14th Annual CUNY Conference on Human Sentence Processing*. Philadelphia, PA.
- Sedivy, J. (2007). Implicatures in real-time conversation: A view from language processing research. *Philosophy Compass* 2/3, 475-496.
- Sedivy, J., Tanenhaus, M., Chambers, C., & Carlson, G. (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition*, 71, 109-147.
- Sperber, D. & Noveck, I. A. (2004). Introduction to Experimental Pragmatics. In Noveck, I. A. & Sperber, D. (eds.), *Experimental Pragmatics*. New York: Palgrave Macmillan. 1-23.
- Sperber, D. & Wilson, D. (1986/1995). *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Spivey-Knowlton, M. & Sedivy, J. (1995). Resolving attachment ambiguities with multiple constraints. *Cognition*, 55, 227-267.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268, 1632–1634.