Non-monotonic reasoning in interpretation

Robert van Rooij and Katrin Schulz

1 Introduction

The original article of Thomason pursues two goals: first to outline the central logical issues of non-monotonic reasoning, and second to indicate possible applications of nonmonotonic reasoning techniques in linguistics. This appendix will follow up on Thomason's second goal and show that linguists have taken up his invitation to use nonmonotonic logic as a formal tool. Particularly in the new and very vivid area of formal pragmatics and at the intersection of semantics and cognitive psychology nonmonotonic logics are playing an important role.

The paper is structured as follows. The purpose of its first part is to present the recent progress made in formal pragmatics by using nonmonotonic logic to describe pragmatic meaning. We will show how minimal models can be used to describe and explain inferences of language use, in particular Gricean conversational implicatures. After this we will discuss how nonmonotonic logic can be used at the semantic-pragmatic interface to account for the *preferred interpretation* of a sentence. In the last part of the paper we will discuss the role of nonmonotonic logic for human reasoning in general. Here we will focus in particular on Stenning and van Lambalgen [2008].

2 Implicatures as nonmonotonic inferences

2.1 Grice's theory of conversational implicatures

We often make assumptions based on what is not said, and communicators make use of this. The standard example for this type of reasoning in artificial intelligence is the following: if a timetable does not mention any direct flight from A to B, we routinely assume that no such flight exists, which, in turn, allows the planner to represent this information simply by not listing such a flight. This type of reasoning is called in AI 'Negation as Failure'. Such reasoning is defeasible, however, because the underlying 'convention' can be explicitly abrogated or suspended. Many systems of non-monotonic reasoning developed within AI — e.g. McCarthy's Circumscription McCarthy [1980], McCarthy [1986] and various variants of Logic Programming — are meant to account for this type of defeasible reasoning.

In linguistics inferences of this type, going beyond what the sentence explicitly says, are called *conversational implicatures*. They are often illustrated with examples like the following: if the question whether Mr. X is a good linguist is answered by "Well, he speaks excellent English", you normally will infer that he is not a good linguist. The fact that this type of inference can be thought of in terms of negation as failure strongly suggests that such inferences can be modeled in terms of systems of non-monotonic reasoning. The purpose of the first part of this paper is to work out this suggestion.

Conversational implicatures are in linguistics essentially connected with the name of Paul

Grice (Grice [1989]), who proposed the name and claimed that these inferences play an important and systematic role in interpretation and need the theoretical attention of linguists. Grice proposed that conversational implicatures result from the assumption that speakers obey a number of maxims of conversation — the maxims of *quality*, *quantity*, *relevance*, and *manner*. These maxims — best thought of as rules of thumb of how speakers (ought to) behave — are stated in a very informal way. Over the years many phenomena have been explained in terms of the Gricean maxims of conversation and especially the first submaxims of *quantity*: *Make your contribution as informative as is required (for the current purposes of the exchange).*

The most famous class of inferences that have been explained using this maxim are the *scalar implicatures* (see Horn [1972], Gazdar [1979] and many others). Classical examples of scalar implicature are, for instance, the inferences from 'ϕ or ψ' to 'not (ϕ and ψ)' and from 'John has two children' to 'John doesn't have *more than* two children'. According to the standard analysis, scalar implicatures are thought of as *generalized* conversational implicatures (GCIs) triggered by specific lexical items. Specific items like 'some', 'or', 'two', and 'possible' are said to come with a conventionally given *scale* of alternative expressions which determine *what* the implicature is. If two lexical expressions S(trong) and W(eak) form a conventionally given linear scale, $\langle$S, W$\rangle$, a (non-complex) sentence in which the *weaker* expression W occurs will according to the GCI view on implicatures always trigger the implicature that the corresponding *stronger* sentence where S is substituted for W is not true. Horn, Gazdar, and others argued that we *require* such linguistically given scales, because an unlimited use of negation as failure in the message will overgenerate enormously.

2.2 Scalar implicatures in empirical perspective

Based on the idea that scalar implicatures are due to interpretation by *default* triggered by a *conventional* given scale prompted Levinson [2000] to make the strong psychological claim that scalar implicatures are *fast* and *automatic* inferences, inferences which don't take any more time to draw than inferences based on standard semantic/conventional meaning. Psycholinguistic evidence of at least three different types strongly suggests that scalar implicatures should not be accounted for as automatic, 'grammatical-like' inferences as suggested by Levinson. The first type of evidence comes from psychological research on human reasoning (e.g. Paris [1973]; Evans and Newstead [1980]).[1] In this type of research, participants are presented with simple sentences in which a scalar term like 'or' occurs, such that (i) the disjunctive sentences are completely unrelated to each other (as in 'The bird is in the nest, or the shoe is on the foot') and (ii) the contextual factors are factored out. It turns out that more than two-thirds of the participants prefer (in our terms) the reading without the scalar implicature to the one with the implicature. This suggests that scalar inferences are anything but 'automatic'.[2]

The second type of evidence pointing in the same direction is that especially young children are bad in inferring scalar implicatures. Noveck [2001], for instance finds that in contrast to adults, most children below age four treat 'Some elephants have trunks' as not being false or misleading if all (shown) elephants have trunks. Later experiments (e.g. Papafragou and Tantalou [2004], Pouscoulous [2006]) have confirmed these results. Other well-known experiments (such as 'False belief'-tasks) indicate that very young children cannot reason about other people's beliefs, goals, and intentions. This suggests that these

[1] Thanks to Bart Geurts for pointing us to this literature.

[2] A different experiment, but leading to the same conclusion, is conducted by Breheny [2006]. He looks at the difference between the time needed to read and comprehend a sentence with and without an implicature.

abilities are also needed to account for the derivation of scalar implicatures, which rules out a 'grammatical-like' view on them.

One reason such abilities might be relevant is to account for the fact that scalar inferences are *context dependent*. As observed by Levinson [2000], for instance, the scalar implicature (1-c) of (1-b) due to the ⟨all, some⟩ scale, is not available in the context of question (1-a) where all that matters is whether *at least some* of the documents are forgeries.

(1) a. Is there any evidence against them?
b. Some of their identity documents are forgeries.
c. Not all of their identity documents are forgeries.

Levinson [2000] accounts for this by proposing that the implicature is still triggered, but later *cancelled* for reasons of relevance. Alternatively, one might propose that it depends on what is taken to be relevant by speaker and hearer whether an implicature is even *triggered*.[3] Most naturally, only Levinson's proposal predicts that in the context of question (1-a) the *reading time* of (1-b) takes longer than when the latter was uttered in a more neutral context. Looking at the reading times of expressions containing scalar terms in different types of contexts, Noveck and Posada [2003], Bott and Noveck [2004], Breheny and Williams [2006], and Zondervan [in press] consistently find that this prediction of Levinson is wrong.[4] A natural conclusion is that whether a scalar implicature is triggered depends on what is relevant in the context. For instance, on which question the sentence in which the trigger occurs gives an answer to.

2.3 Scalar implicatures in minimal models

In the last section we have seen that there is good empirical evidence that scalar implicatures are pragmatic inferences that are dependent on the utterance context. In this section we focus on formally describing these implicatures with their characteristic properties.

Scalar implicatures can often be paraphrased in terms of 'only'. In a context where it is relevant how many students passed, for example, (2-a) gives rise to the scalar implicature that not all students passed. This inference can also be derived from (2-b) — but now it follows from the semantic meaning of the sentence.

(2) a. Some of the students passed.
b. Only $[\text{some}]_F$ of the students passed.[5]

Given this, it is natural to connect the analysis of scalar implicatures with the semantic analysis of 'only'. The standard analyses (especially Rooth [1985]; Krifka [1995]) of 'only' proposes that 'Only ϕ' should be interpreted as (3)[6] ($[\![\phi]\!]$ stands for the denotation of ϕ and $Alt(\phi)$ for the set of alternatives of ϕ.)

(3) $$[\![only(\phi, Alt(\phi))]\!] \stackrel{def}{=} \{w \in [\![\phi]\!] | \forall \psi \in Alt(\phi) : w \in [\![\psi]\!] \Rightarrow [\![\phi]\!] \subseteq [\![\psi]\!]\}.$$

To account for the scalar implicatures of 'ϕ', one could assume that a sentence simply has a silent 'only' in front of it, which we will call '*Prag*':[7]

[3]Proponents of this alternative view of scalar-implicatures include Hirschberg [1985], van Kuppevelt [1996], Carston [1998], and the present authors.

[4]It should be remarked, though, that a substantial ammount of research in this area is performed by proponents of Relevance theory. As an antidote, see Storto and Tannenhaus [2004] for evidence which might point in the opposite direction.

[5]The notation $[\cdot]_F$ means that the relevant item receives focal stress, i.e., an H^*L prosodic contour.

[6]In this representation the distinction between presupposition and assertion is ignored.

[7]Krifka [1995] introduces our *Prag* under the name '*Scal.Assert*. Chierchia et al. [to appear] explicitly

(4) $\quad Prag(\phi, Alt(\phi)) \quad \stackrel{def}{=} \quad \{w \in [\![\phi]\!] | \forall \psi \in Alt(\phi) : w \in [\![\psi]\!] \Rightarrow [\![\phi]\!] \subseteq [\![\psi]\!]\}.$

In case the alternative of (2-a) is '*All* of the students passed', the desired scalar implicature is indeed accounted for.

McCawley [1993] noticed that in case one scalar item is embedded under another one — as in (5),[8] — an interpretation rule like *Prag* does not give rise to the desired prediction that only one student passed.

(5) Alice passed, Bob passed, *or* Cindy passed

This observation can be straightforwardly accounted for if we adopt a different pragmatic interpretation rule together with a different way to determine alternatives. First, we will assume that the set of alternatives is not so much dependent on the particular sentence 'ϕ' that is uttered, but is just a set of sentences $\mathcal{L}$ whose truth or falsity is relevant in the context. What is relevant depends on what is at issue, and if the issue is who passed, it is natural to think of $\mathcal{L}$ as a set like {*Alice passed, Bob passed, Cindy passed, Delia passed*} (which should perhaps be closed under conjunction and disjunction). According to the new pragmatic interpretation rule — call it $Prag_{\mathcal{L}}$ —, w is compatible with the pragmatic interpretation of ϕ iff (i) ϕ is true in w, and (ii) there is no other world v in which ϕ is true where less alternatives in $\mathcal{L}$ are true than are true in w, see (6). Notice that by pragmatically interpreting a sentence in terms of $Prag_{\mathcal{L}}$, we do not assume, or are not required to assume, that certain expressions come with a *conventionally* triggered set of alternative expressions which are *linearly* ordered as was standardly assumed. The pragmatic interpretation rule $Prag_{\mathcal{L}}$ correctly predicts that from (5) we can pragmatically infer (i) that only one of Alice, Bob, and Cindy passed, and (ii) that Delia did not pass.

(6) $\quad Prag_{\mathcal{L}}(\phi) \quad \stackrel{def}{=} \quad \{w \in [\![\phi]\!] | \neg\exists v \in [\![\phi]\!] : \forall \psi \in \mathcal{L} : v \in [\![\psi]\!] \Rightarrow w \in [\![\psi]\!]\}.$

As it turns out, this new interpretation rule is closely related to McCarthy [1980]'s *predicate circumscription.* Circumscription was introduced in AI to account for the type of reasoning called 'Negation as Failure' as mentioned in section 2.1. To see the connection with $Prag_{\mathcal{L}}$, let us first observe that $Prag_{\mathcal{L}}$ describes a minimal interpretation of ϕ. Certain worlds among those where ϕ is true are selected, and these are worlds that are in some sense minimal. The relevant order is defined in terms of the language $\mathcal{L}$. A world is the more minimal, the fewer $\mathcal{L}$-sentences it makes true: $v \leq_{\mathcal{L}} w$ iff$_{def}$ $\forall \psi \in \mathcal{L} : v \in [\![\psi]\!] \Rightarrow w \in [\![\psi]\!]$. In case $\mathcal{L}$ is defined as above as describing the extension of some predicate, i.e. $\mathcal{L}$ consists of sentences $P(a), P(b), ...$ for some predicate P (probably closed under conjunction and disjunction) and if the language is rich enough, then the order is equivalent to one comparing the extension of P: $v \leq_P w$ iff$_{def}$ $P(v) \subseteq P(w)$. That means that *Prag* is equivalent to *circ* defined as given in (7). But this comes very close to the model theoretic idea of predicate circumscription, in fact it is the special case of McCarthy's *predicate circumscription* in which everything else in the formal language is allowed to vary.[9]

---

propose that sentences, as well as subsentences, might have such a silent 'only' in front of it in its logical form.

[8] Landman [2000] and Chierchia [2004] discuss structurally similar examples like 'Mary is either working at her paper *or* seeing *some* of her students.'

[9] There is a strongly related link from linguistics to predicate circumscription. As has been pointed out by Benthem [1989], predicate circumscription is also closely related to the exhaustivity operator *Exh* introduced by Groenendijk and Stokhof [1984]. This operator takes as arguments (i) the predicate P of a question, and (ii) the meaning of a term-answer, and turns it into a new formula, describing the exhaustive interpretation of the answer: $exh(F, P) =_{def} F(P) \wedge \neg\exists P' \subseteq D : F(P') \wedge P' \subset P$.

$$circ(\phi, P) \stackrel{def}{=} \{w \in [\![\phi]\!] \mid \neg\exists v \in [\![\phi]\!] : v <_P w\}. \quad (7)$$

2.4 Prospects and problems of the circumscription account

Quite a number of conversational implicatures (including *scalar* ones) can be accounted for in terms of a minimal model analysis.[10] Except for the obvious result that from the answer 'Alice passed' to the question 'Who passed?' we derive that Alice is the only one who passed, we also derive (i) the *exclusive* reading of a disjunctive answer like 'Alice passed or Bob passed' to the same question, and the reading that one and only one person passed from 'Alice passed, Bob passed, or Cindy passed'; (iii) the implicature that not everybody passed from the answer that *most* did; (iv) the so-called *conversion*-inference that only men passed, if the answer is 'Every man passed'. We also derive (v) for 'Alice ate three apples' that Alice ate *exactly* three apples[11], and (vi) the biconditional reading of 'Alice will pass if Bob will', if this sentence is given as answer to the polar question 'Will Alice pass?'. Another pleasing property of a circumscription analysis of implicatures is that it predicts that it depends on the context, or question-predicate, whether we observe these inferences.

Schulz and van Rooij [2006] suggest that some obvious problems of standard pragmatic interpretation rules (such as the rule given in Groenendijk and Stokhof [1984] of exhaustive interpretation) can be solved when we take the minimal models into account. They propose, for instance, that to account for the context-dependence of exhaustive interpretation, the beliefs and preferences of agents are relevant to determine the ordering relation between worlds required to define the minimal models. In this way they get a better grasp of the context (and relevance) dependence of implicatures, and can account for, among others, both mention-all and mention-some readings of answers (which Groenendijk & Stokhof could not). Thinking of pragmatic interpretation in terms of minimal models also might help to tackle the so-called *functionality problem*. The functionality problem follows from the fact that circumscription works immediately on the *semantic meaning* of an expression. It follows that if two sentences have the same semantic meaning, they are predicted to give the same implicatures as well. It is, for instance, standardly assumed that 'three men' has the same semantic meaning as '*at least* three men', and that 'Alice passed or Bob passed' has the same semantic meaning as 'Alice passed, Bob passed, *or both* passed'. But sentences in which the former examples occur give rise to a 'scalar' implicature, while the latter do not. This problem suggests that the notion of *meaning* adopted is *too coarse-grained*.[12] There might be several ways to solve this problem: Schulz and van Rooij [2006] proposed that instead of thinking of meanings just in terms of possible worlds, we should think of them in terms of world-assignment pairs, as is standard in *dynamic semantics*. Sevi [2006], instead, proposed to make use of recent insights of how to represent pluralities.

2.5 Grice and the theory of 'only knowing'

The interpretation function *circ* described above only provides a description of certain conversational implicatures, but does not explain where these inferences come from. When Grice discussed conversational implicatures, he wanted to derive them from his maxims

[10] Also Wainer (Wainer [1991], Wainer [2007]) makes use of circumscription to account for scalar implicatures. He does so, however, by making use of an explicit abnormality predicate in the logical representation of the sentence. Accounting for implicatures in this way is thus not really in the spirit of Grice, though more standard in AI and also used in section 3 of this paper.

[11] Given an *at least* -semantics for numerals and in the context of a question 'How many apples did Alice eat?'.

[12] The well-known problem of 'logical omniscience' is closely related, and adopting a more fine-grained notion of meaning has been suggested by various authors to solve this problem as well.

of conversation. The implicatures discussed above should intuitively be based on Grice's maxim of *quality*, according to which a speaker may only say what he knows to be true, and his combined maxim of *relevance* and (the first submaxim of) *quantity*. In order to make this derivation precise, we would need a formalization of the inferences that an interpreter can derive given that she assumes the speaker to obey these maxims. In a second step we can then evaluate this formalization by comparing its predictions with those made by *circ*. It turns out that such a formalization can be provided. It builds on autoepistemic logic (e.g. levesque [1990]), in particular, on theories of 'only knowing' (Halpern and Moses [1984], Hoek et al. [1999]).

The Gricean maxims concern the epistemic state of the speaker. To be able to speak about what the speaker knows one can use basic modal logic. Assume that $M = \langle W, R, V\rangle$ is a Kripke model with W a set of possible worlds, V an interpretation function, and R an accessibility relation connecting a world w_0 with all those worlds consistent with what the speaker knows in w_0. $\mathbf{K}\phi$ is true given a world w_0 and a model M if ϕ holds in all worlds w accessible from w_0 by R. What kind of information can an interpreter infer on hearing an utterance of ϕ? If the interpreter assumes that the speaker obeys the maxim of quality, she can conclude that $\mathbf{K}\phi$ holds.[13] If '$\mathbf{K}$' represents knowledge, it follows that ϕ is the case.

But how to formalize the maxims of quantity and relevance? How can we formalize that the speaker has provided with ϕ all the relevant knowledge she has? Following Halpern and Moses [1984], Shoham [1988], and Hoek et al. [1999] the idea is again to select for minimal models. But this time minimal means minimal with respect to the knowledge of the speaker. There are different ways to define orders comparing what the speaker knows. One way is to use a formal language $\mathcal{L}$. This language is supposed to express exactly the information relevant in the utterance context. Then one can define an order that compares pairs of a world and a model with respect to how many $\mathcal{L}$ sentences the speaker knows in them: $\langle M, w\rangle \leq^{\mathbf{K}}_{\mathcal{L}} \langle M', w'\rangle$ iff $\forall\varphi \in \mathcal{L} : M, w \models \mathbf{K}\varphi \Rightarrow M', w' \models \mathbf{K}\varphi$. Based on this order we can define a new Gricean interpretation function *Grice* that is supposed to capture the maxims of quality, quantity, and relevance by selecting among those worlds where the speaker knows ϕ to be true those that are minimal with respect to the newly defined order.[14]

(8) $$Grice_{\mathcal{L}}(\phi) \stackrel{def}{=} \{\langle M, w\rangle \in [\![\mathbf{K}\phi]\!] : \forall\langle M', w'\rangle \in [\![\mathbf{K}\phi]\!] : \langle M, w\rangle \leq^{\mathbf{K}}_{\mathcal{L}} \langle M', w'\rangle\}.$$

So far we still left open how exactly the language $\mathcal{L}$ describing what counts as relevant information has to be defined. This is not at all trivial. For instance, if $\mathcal{L}$ contains *all* epistemic formulas, then no two knowledge states are comparable.[15] A choice that works well is to select a set of relevant primitive expressions and close this set under conjunction and disjunction. It remains to be seen whether this kind of positive closure can be given a Gricean motivation.

Let us now assume that what is relevant in a particular context is the extension of some predicate P. In this case $\mathcal{L}$ is determined as the positive closure of primitive sentence

[13]In this way Gazdar [1979] explains why nobody can appropriately say things like 'The earth is round, but I don't know it'. A speaker of this sentence violates the maxim of quality, if R is transitive.

[14]Some formulas have more than one minimal state. The formula '$\mathbf{K}p \vee \mathbf{K}q$', for instance, has two incompatible minimal states: one in which the speaker only knows p, and one in which she only knows q. Halpern and Moses [1984] dubbed those formulas that don't have a unique minimal state '*dishonest*', because one cannot consistently claim that one only knows the information expressed by such a formula. A natural proposal would be that except for selecting minimal states, the maxim of quantity also demands that a sentence ϕ can only be uttered appropriately if it gives rise to a *unique* minimal state.

[15]The problem is that if R is serial, transtive, and euclidean, and $\langle W, R, V\rangle$ thus an **KD45** model, it holds for all formulas ϕ that $M, w \not\models \mathbf{K}\phi$ iff $M, w \models \neg\mathbf{K}\phi$. Switching to **S4** helps, but as explained by Hoek et al. [1999], perhaps so for accidental reasons. The system **S4.2**, for instance, popular to model knowledge, has similar problems.

$P(a), P(b), ...$ — let us call this language $\mathcal{L}^+(P)$. The approach just developed allows one then to derive from the assertion '$P(a)$' that the speaker doesn't know that '$P(a) \wedge P(b)$' is true, and thus that the speaker doesn't know that '$P(b)$' is true. Similarly, from the assertion '[Alice or Bob]$_{\mathcal{F}}$ passed' we conclude that the speaker does not know of anybody that he or she passed. This is a nice result, but in many cases we conclude something stronger: in the first example that Bob *did not* pass, and something similar for the second example. How do we account for this extra inference in terms of our modal-logical setting? This can be accounted for by assuming that speakers, in addition to obeying the Gricean maxims, are *maximally competent* (as far as this is consistent with obeying these maxims).[16] This can be described by selecting among the elements of $Grice_{\mathcal{L}^+(P)}(\phi)$, the ones where the competence of the speaker with respect to the extension of property P is maximal. This can be formalized using the same techniques as before. We simply use as order-defining language the full closure of the primitive sentences $P(a), P(b), ...$[17] — let us call this language $\mathcal{L}(P)$. On the result of the function *Grice* we define a new order $\preceq^{\mathbf{K}}_{P}$ comparing the speaker's knowledge on P relative to $\mathcal{L}(P)$. This order is then used in the definition of a new interpretation function *Comp* selecting model-world pairs where the speaker is maximally competent — see (9). *Comp* applied on top of *Grice* gives us then our new pragmatic interpretation function.

(9) $$Comp_{\mathcal{L}(P)}(S) \stackrel{def}{=} \{\langle M, w\rangle \in S : \forall \langle M', w'\rangle \in S : \langle M, w\rangle \geq^{\mathbf{K}}_{\mathcal{L}(P)} \langle M', w'\rangle\}.$$

By *Comp* we strengthen the Gricean inference that the speaker does not know, for instance, that Bob passed from the utterance of 'Alice passed' to the inference that the speaker knows that Bob did not pass.

At the beginning of this section we formulated the goal of providing a formalization of common sense conversational inferences based on the maxim of quality, the first submaxim of quantity, and the maxim of relevance, and then checking the predictions by comparing the output with the interpretation function *circ*. Now it can be shown that if we model a Gricean speaker in terms of $Grice(\phi, P)$, assume that she is maximally competent and limit ourselves to non-modal statements ϕ,[18] we derive exactly the same implicatures as we can derive using predicate circumscription.[19]

Fact 1 *For all modal-free* ϕ, ψ: $Comp_{\mathcal{L}(P)}(Grice_{\mathcal{L}^+(P)}(\phi)) \models \psi$ *iff* $circ(\phi, P) \models \psi$.

On the basis of this fact, we can claim that we have provided a Gricean motivation of an analysis of conversational implicatures in terms of predicate circumscription.

2.6 Implicatures in embedded contexts

The above analyses of implicatures assumed that implicatures work on the output of the grammar. This is in line with the Gricean view that the conversational implicatures of an utterance are generated *globally*, *after* the grammar assigned a meaning to it. More recently, a so-called *localist* view (e.g. Landman [2000], Chierchia [2004]) has become rather popular. According to this view the pragmatic interpretation function need not be applied only to the whole asserted sentence, but can be applied to subparts of what is asserted as well.

The arguments against the standard global picture of implicatures go back a long way. One serious type of problem was already discussed by Cohen [1971], who noted that if we use a scalar term within the scope of 'believe' as in (10)

[16] For a related notion, see Sauerland [2004].

[17] Closure under conjunction and negation.

[18] The Gricean interpretation make better predictions, though, for modal statements, especially those that refer to the speaker's own information state.

[19] For the proofs, see van Rooij and Schulz [2004] and Spector [2003].

(10) John believes that some students passed the examination.

we can interpret the sentence as saying that John believes that it is not the case that *all* students passed. The problem with the standard analysis is that it only predicts that John *does not* believe some students passed, which is much weaker.[20]

Geurts and Pouscoulous [2008] experimentally compared the interpretation of the scalar term 'some' occurring in a non-embedded position with occurrences of the same word embedded in the scope of 'think', 'want', deontic 'must', and the universal 'all'. They found that the rates in which scalar implicatures are derived drop dramatically under embedding, which is not what localist theories would lead us to expect. But they also found that scalar terms embedded under 'want' hardly ever give rise to local implicatures, while scalar terms embedded under 'think' do so much more often. This finding is hard to explain under a localist view. It might be explained, however, using a standard 'globalist' story. Suppose for a moment that the alternatives of a sentence of the form '$\Box_j \phi$' are all of the form '$\Box_j \psi$'. Following Schulz [2005], van Rooij and Schulz [2004] and Russell [2006] propose that implicatures under modal operators can be strengthened in case a competence assumption is made: If the speaker knows that for all (local) alternatives ψ it holds that John is competent about ψ — meaning that either $\Box_j \phi$ holds, or that $\Box_j \neg\psi$ holds —, we can infer from the assertion that $\Box_j \phi$ is true to the conclusion that (the speaker knows that) $\Box_j \neg\psi$ for all ψ stronger than ϕ. This required competence assumption might explain the drop of scalar implicatures under embedding. Geurts and Pouscoulous [2008] naturally assume furthermore that the competence assumption is much more natural for 'think' than for 'want', and propose to explain in this way the difference between 'local implicatures' under these two verbs.

A global analysis by itself doesn't immediately explain implicatures in embedded contexts. This is neither the case for embedded items under 'necessity' operators, nor for ones embedded under 'possibility' operators. Consider, for instance, the 'free choice permission' problem. The problem is how to account for the fact that from (11) one can infer that John may take an apple and that he may take a pear.

(11) You (John) may take an apple or a pear.

According to Kamp [1973], permissions are special speech acts and he proposes a solution of the problem in which this idea is essential. But (11) can be used as an assertion as well and still give rise to the free choice inference. Schulz [2005], therefore, proposes that the inference is a Gricean implicature, and accounts for it in a global way in terms of minimal models. This analysis works very well, but is based on the disputed assumption that 'John may not take an apple' is a relevant alternative.[21] To preserve a global explanation without assuming such alternatives, it seems most natural to assume that pragmatic interpretation is crucially *bidirectional*, taking also into account how the speaker would have expressed his alternative information states. We have nothing against such a move,[22] but it again does complicate the picture.

[20]Cohen [1971] proposed to account for this reading by assuming that 'or' is *ambiguous* between the inclusive and the exclusive reading. Unfortunately, neither reading of 'or' can account for the intuition that (5) is true iff only one of the three disjuncts is true.

[21]For a popular 'localist' analysis of (11), see Fox [2007].

[22]This move is made in bidirectional optimality theory (cf. Blutner [2000]), as well as in game theoretical analyses of pragmatic interpretation.

3 More on nonmonotonic reasoning and linguistics

3.1 Preferred interpretation

Conversational implicatures are not the only part of interpretation where nonmonotonic reasoning plays a role. It is well-known that for counterfactual conditionals, already the truth conditions show nonmonotonic behavior. But as shown already by Lewis [1973], other semantic phenomena show exactly the same pattern. For instance, contextually definite descriptions: 'The pig is grunting, but the pig with floppy ears is not grunting'. What is intuitively going on here is that the description 'the N' picks up the unique *most salient* individual with property N. But the most salient pig need not have floppy ears, which allows for the nonmonotonic behavior. What is a salient, or *normal*, exemplar also seems crucial for the interpretation of *generic* sentences. In fact, researchers in AI have produced many theories of non-monotonic reasoning that can be seen as attempting to give a semantics for genericity. If we know that x is a bird, our acceptance of the generic sentence 'Birds fly' allows us to expect that x flies, without guaranteeing that she can. Nonmonotonic theories of generics are attractive not only because they allow for exceptions, but also because they immediately account for the intuition that a generic sentence like *A lion has a mane* makes a claim about every possible lion, rather than a claim about the closed class of all existing lions. Some of the most detailed theories of generics (Delgrande [1998], Asher and Morreau [1995], and Veltman [1996]) are built on conditional logics like Lewis [1973]. Such theories predict that generic sentences have truth conditions, and thus can account for *nested* generic statements. Moreover, such theories can naturally account for conflicting rules, where some take priority over others. Being able to rank nonmonotonic rules seems very important in linguistic pragmatics as well.

Nonmonotonic logic is used in semantics, but its role is (potentially, at least) much more important in pragmatics. The reason is that these logics can be used especially to account for *preferred interpretation*. Take, for instance, the interpretation of plural reciprocals. It is well-known that sentences like 'The children followed each other' allow for many different interpretations. Still, such sentences are most of the time understood pretty well. Dalrymple et al. [1998] propose that this is due to a particular interpretation strategy. According to their "Strongest meaning hypothesis" a sentence should preferentially be interpreted in the strongest possible way consistent with its (underspecified) semantic meaning. This simple strategy predicts surprisingly well, and has become popular to account for other phenomena too. But it is important to note here that the hypothesis used is a nonmonotonic interpretation strategy because adding more information might make a stronger interpretation impossible. If we add 'into the church', for instance, our original sentence cannot be interpreted anymore as saying that all children followed another child, but can at most mean that any child followed, *or was followed*, by another child.

A systematic theory of preferred interpretation is also crucial for the resolution of pronouns, which clearly involves nonmonotonic inference patterns. This can already be illustrated by the following simple discourse (12).

(12) John met Bill at the station. *He* greeted *him*.

The pronouns *he* and *him* could refer to either John or Bill. For reasons of syntactic parallelism, however, there seems to be a preference for interpreting *he* as John and *him* as Bill. But this preference can be overruled if we add additional information. For instance, if we add 'John greeted him back', we have to reinterpret *he* as Bill, and *him* as John, due to the indefeasible semantics associated with the adverb *back*. In other cases, a non-preferred interpretation is suggested due to inconsistency with world-knowledge (e.g. in 'John hit

Bill. *He* got injured.'), or emphatic stress ('John called Bill a republican. Then *HE* insulted *HIM*.'). It has been argued that even (some of) Chomsky's binding rules in syntax can be violated. The well-formedness of the following discourse due to Evans [1980], for instance, suggests this for Chomsky's constraint saying that a non-reflexive pronoun like 'him' cannot be bound locally.

(13) Everyone despises poor John_1. Mary despises him_1. Sue despises him_1, and even John_1 despises him_1.

Centering theory (Grosz et al. [1995]) was specifically designed to account for pronoun resolution, and is stated in terms of a set of defeasible, or nonmonotonic, rules. Hendriks and de Hoop [2001], Beaver [2004], and others seek to account for anaphora and pronoun resolution in terms of Optimality Theory (OT), a more general theory of interpretation which makes use of violable constraints. It is clear, however, that any system of ranked violable rules can be seen as a nonmonotonic logic that can account for conflicting rules, where some take priority over others (compare the discussion on generics). Using such a logic is exactly what was proposed by Kameyama [1996] for the case of pronoun resolution.[23]

Nonmonotonic logic plays a dominant role in the work of Asher and collaborators. As mentioned already by Thomason, they were pioneers in their use of nonmonotonic logic to infer discourse relations and employ that for inferring temporal anaphora, lexical disambiguation, bridging, and many other things (e.g. Asher and Lascarides [1993], Asher and Lascarides [1995]). This work has been extended on significantly and elaborated over the last 15 years (see especially Asher and Lascarides [2003]). In their work on temporal anaphora, for instance, they observed that syntax sometimes only provides preferences which can be overridden by world-knowledge. Normally the event a first sentence in simple past is about is temporally located before the event the consecutive sentence with simple past is about. But world-knowledge somestimes forces us to interpret otherwise, as in the discourse 'John fell. Mary pushed him'. The theory of Asher and Lascarides [2003] deals with many other things required for interpreting texts, including inferring speech acts, and presupposition resolution. Most of these phenomena have been accounted for separately before (e.g., the treatment of speech acts in Appelt and Konolige [1988] and the analysis of bridging using abduction in Hobbs [1993]), but Asher and Lascarides brought them all under a very general theory of nonmonotonic discourse interpretation.

3.2 Bringing psychology into the picture

Why does nonmonotonic reasoning play a role in such diverse areas of interpretation? A very interesting answer to this question has been brought forward by van Lambalgen and Hamm [2005]. They propose that in order to interpret discourse, humans have excepted the cognitive processes and facilities involved in off-line planning. Furthermore, they claim that off-line planning involves mentally constructing minimal models, i.e. nonmonotonic reasoning.[24] In van Lambalgen and Hamm [2005] the focus lies on temporal discourse. In this case a relation to planning is very intuitive, especially to interpret progressives like 'building a barn'. Following Moens and Steedman [1988], they argue that our construction of time essentially involves our planning facilities and our conception of causal dependencies. But

[23]It is one thing to *state* violable rules and rank them, but quite another to *explain* these. Reinhart [1983], Levinson [1991], and others made interesting attempts to *derive* (some of) the binding constraints from some more general pragmatic principles, and much of the more recent interest in Bidirectional Optimality Theory (e.g. Blutner and Zeevat [2004]) and (evolutionary) Game Theory (e.g. Benz et al. [2005]) should be seen in this light as well.

[24]For a similar view on the semantics of perception reports see van der Does and van Lambalgen [2000].

van Lambalgen and Hamm [2005] propose that the cognitive planning device is used for discourse interpretation in general. Thus, when processing incoming sentences, an interpreter is constantly building and revising a minimal discourse model. This is consistent with the analyses of (temporal) pronoun resolution discussed above: 'linguistic' constraints can be overruled by additional information and world knowledge. This raises the question what the role of classical logic, or compositional semantics, is in such a picture. This question touches upon topics studied in psychology.

There exists a substantial literature on reasoning within psychology. This work is nearly independent from related work in AI and linguistics in this area. This independence might partly be an effect of the conviction among psychologists that logic does not provide the right framework to describe or explain human reasoning. Originally, things were quite different: psychologists started out with the idea that to reason is to be logical. But the picture of logic psychologists embraced was an absolute one: classical logic. Thus to test the thesis that people reason logical in this sense one distinguished a set of premises, translated them into classical logic, calculated the valid inferences — valid in the sense of classical logic — and then checked whether humans presented with this set of premisses count the same sentences as valid or invalid inferences. Unfortunately, it turned out that the classical notion of inference does not match human reasoning. This led to the growing conviction among psychologists that logic is not the proper framework to describe human reasoning behavior.

Of course, this is not the only way to deal with these results: one could, for instance, also hypothesize that humans do not reason classically but nonmonotonically. Though the predictions become better if one adopts this hypothesis, what still cannot be explained is the fact that there remains a lot of variation between the performance of humans in these reasoning tasks. A crucial observation in this respect is that the reasoning strategies people apply to particular tasks appear to vary in correlation with how they interpret the reasoning task they have to perform. Should the reasoning involve inference to the preferred interpretation, or not? If so (and only if so), the reasoner will try to integrate the given information in a representation of one single model and will add as much as possible additional assumptions available. The result will be a nonmonotonic notion of inference based on constructing minimal models.

3.3 An application: the suppression task

The suppression task is one of the classical experiments in the psychology of reasoning. It was originally designed by Byrne [1989]. The subjects are presented with the information given in (14-a) and (14-b). When asked whether from this information (14-c) can be infered 90% of the subjects agree. However, if (14-d) is added to the premisses the percentage drops to 60 % while in case (14-e) is added the percentage is approximately the same as in the first setting.

(14) a. If she has an essay to write she will study late in the libary.
b. She has an essay to write.
c. She will study late in the libary.
d. If the libary is open she will study late in the libary.
e. If she has a textbook to read she will study late in the libary.

These results suggested to many psychologists that formal logic should not be used to describe human reasoning. For if logic were fit to describe human reasoning, expressions with the same logical form should allow the same type of inferences. However, (14-d) and (14-e) have the same logical form, but while in the first case the inference to (14-c) is suppressed, in the second case the inference is still available.

Stenning and van Lambalgen [2008] argue, instead, that this conclusion was based on the mistaken assumption that (14-d) and (14-e) by necessity have the same logical form. In their preferred interpretation, (14-a)/(14-b) + (14-d) and (14-a)/(14-b) + (14-e) have to be assigned a different logical form. Conditional sentences are analyzed as law-like relations *If A and nothing abnormal is the case, then B*. Stenning and van Lambalgen [2008] propose to formalize law-like relations making use of propositional logic programming with negation as failure.[25] The logical form of the first conditional sentence is given by the program $essay.to.write \wedge \neg ab_1 \rightarrow stays.in.library$, together with the condition $\bot \rightarrow ab_1$. Obviously, ab_1 is an abnormality clause, expressing that something abnormal is happening. The extra condition demands that abnormality should be minimized. Notice that the $\rightarrow$ in these logical forms is simply a marker distinguishing body and head in a program clause, and should not be interpreted as a truth-conditional connective. The intended denotation of the program is a minimal model. To reach it, a notion of completion for logic programs is defined that makes sure that negation as failure is only applied to abnormality clauses, and not to arbitrary propositional variables. The authors show that the minimal models the completion describes can be constructed as the fixed points of a monotonic operator on three-valued models that can be reached in finitely many steps.

To account for the difference between (14-a)/(14-b) + (14-d) and (14-a)/(14-b) + (14-e), Stenning and van Lambalgen [2008] argue that whereas (14-d) is interpreted as interacting with the abnormality condition ab_1, this is not the case for (14-e). The conditional in (14-e) adds to the scenario built up by (14-a) and (14-b) a program clause $book.to.read \wedge ab_2 \rightarrow study.in.library$ with the abnormality minimization clause $\bot \rightarrow ab_2$. From these conditions one can infer with closed world reasoning that she will study late in the library. Thus the theory predicts correctly that in the scenario extended with (14-e) no suppression of the inference (14-c) is observed. Also the conditional (14-d) adds to the scenario a program clause $library.open \wedge ab_2 \rightarrow study.in.library$ with the abnormality minimization clause $\bot \rightarrow ab_2$. But due to inferring to preferred interpretation, two constraints on the interaction of the involved abnormality clauses are added: $\neg library.open \rightarrow ab_1$ and $\neg essay.to.write \rightarrow ab_2$. In this case, one cannot infer that she will study late in the library by applying closed world reasoning. This explains suppression in case (14-d).

How do we determine whether the antecedent of one conditional will affect the abnormality clause of the other? Stenning and van Lambalgen [2008] say that it is based on context and world knowledge. Thus, as far as ‘linguistic’ constraints are concerned, almost anything goes. To a certain extent this is also what Stenning and van Lambalgen want, because as they show in their reports on dialogue experiments, people exploit the most diverse interpretation strategies.

4 Conclusions

In this paper we have discussed some recent applications of nonmonotonic logic in linguistics. As already mentioned at the beginning, it is remarkable to see that some of the phenomena we discussed were already among the original motivations for the development of nonmonotonic logic. For instance, when McCarthy [1986] discusses possible applications of nonmonotonic logic in general and circumscription in particular he also mentions communicational conventions. The example he provides actually involves Grice’s maxims of

[25]In their use of logical programming they follow Suppes [1980] and others in taking a procedural stance towards meaning.

conversation.[26] Given this observation, the recent developments in the area appear to mark only a shift in perspective. In the 70s and particularly the 80s in AI one came across some interesting questions concerning reasoning and intelligent behavior in general. This led to the development of very sophisticated logical tools in order to answer these questions. However, the study of the involved phenomena (reasoning, communication) itself stayed at a relatively basic level, compared to the expertise available for instance in linguistics and psychology. At the same time — or even before that — similar problems also have been discussed in linguistics and psychology, but without the means to developing formal tool to deal with them. At a certain point in the 90s the interests of AI in nonmonotonic logic weakened, because the conviction emerged that nonmonotonic logic is not fit to deal with modeling the relevant phenomena in an efficient way. But now linguistics and psychology are taking over, with a more sophisticated understanding of communication. Nonmonotonic logic is rediscovered as a very useful tool to describe and explain certain aspects of interpretation and reasoning. What we might observe within the next years is that the more sophisticated applications of the original techniques will stimulate the development of the formalisms. In the end it might turn out that the situation for nonmonotonic logic is not as devastating as was thought before. For instance, it might well prove to be the case that those restrictions needed to obtain computationally tractable versions of nonmonotonic logics are empirically very natural. This is suggested by Asher and Lascarides [2003], and is certainly the case for the formalism of logic programming as employed in van Lambalgen and Hamm [2005] and Stenning and van Lambalgen [2008]. In the latter book it is shown that the employed solution concept — i.e., executing logical programe by means of monotonic operators — corresponds to computation in a suitable recurrent neural network. This implementation of the formalism into neural networks allows grounded speculation about the neural reality underlying reasoning.

References

D. Appelt and K. Konolige. A practical nonmonotonic theory for reasoning about speech acts. In *Proceedings of the 26th Meeting of the Association for Computational Linguistics*, pages 170–178. Association for Computational Linguistics, Association for Computational Linguistics, 1988.

N. Asher and A. Lascarides. Temporal interpretation, discourse relations, and commonsense entailment. *Linguistics and Philosophy*, 16:437–493, 1993.

N. Asher and A. Lascarides. *Logics of Conversation.* Cambridge University Press, Cambridge, 2003.

N. Asher and A. Lascarides. Lexical disambiguation in a discourse context. *Journal of Semantics*, 12:96–108, 1995.

N. Asher and M. Morreau. What some generic sentences mean. In G. Carlson and J. Pelletier, editors, *The Generic Book*, pages 300–338. Chicago University Press, Chicago, 1995.

D. Beaver. The optimization of discourse anaphora. *Linguistics and Philosophy*, 27:3–56, 2004.

J. van Benthem. Semantic parallels in natural language and computation. In H. D. Ebbinghaus et al., editors, *Logic Colloquium '87*, pages 331–375, Amsterdam, 1989. Elsevier Science Publishers.

[26] Though McCarthy [1986] doesn't seem to be aware of Grice's work.

A. Benz, G. Jaeger, and R. van Rooij. *Games and Pragmatics*. Palgrave Macmillan, Hampshire, 2005.

R. Blutner. Some aspects of optimality in natural language interpretation. *Natural Language Semantics*, 17:189–216, 2000.

R. Blutner and H. Zeevat. *Optimality Theory and Pragmatics*. Palgrave Macmillan, Hampshire, 2004.

L. Bott and I. Noveck. Some utterances are underinformative: the onset and time course of scalar implicatures. *Journal of Memory and Language*, 51:437–457, 2004.

N. Katsos Breheny, R. and J. Williams. Are genearised scalar implicatures generated by default. *Cognition*, 100:434–463, 2006.

R. Breheny. Communication and folk psychology. *Mind and Language*, 21:74–104, 2006.

R.M.J. Byrne. Suppressing valid inferences with conditionals. *Cognition*, 31:61–83, 1989.

R. Carston. Informativeness, relevance and scalar implicature. In R. Carston and S. Uchida, editors, *Relevance Theory: Applications and Implications*, pages 179–236. John Benjamins, Amsterdam, 1998.

G Chierchia. Scalar implicatures, polarity phenomena and the syntax/pragmatics interface. In A. Belletti, editor, *Structures and Beyond*. Oxford University Press, Oxford, 2004.

G. Chierchia, D. Fox, and B. Spector. The grammatical view of sclar implicatures and the semantics pragmatics interface. In P. Portner en anderen, editor, *Handbook of Semantics*. Mouton de Gruyter, Berlin, to appear.

L.J. Cohen. The logical particles of natural language. In Y. Bar-Hillel, editor, *Pragmatics of Natural Language*, pages 50–68. Reidel, Dordrecht, 1971.

M. Dalrymple, M. Kanazawa, Y. Kim, S. Mchombo, and S. Peters. Reciprocal expressions and the concept of reciprocity. *Linguistics and Philosophy*, 21:159–210, 1998.

J. Delgrande. An approach to default reasoning based on a first-order conditional logic. *Artificial Intelligence*, 36:63–90, 1998.

G. Evans. Pronouns. *Linguistic Inquiry*, 11:337–362, 1980.

J.S. Evans and S.E. Newstead. A study of disjunctive reasoning. *Psychological research*, 41: 373–388, 1980.

D. Fox. Free choice disjunction and the theory of scalar implicatures. Technical report, MIT, 2007.

G. Gazdar. *Pragmatics*. Academic Press, London, 1979.

B. Geurts and N. Pouscoulous. Local implicatures? Technical report, Radboud University of Nijmegen, 2008.

H. P. Grice. Logic and conversation. In *Studies in the Way of Words*. Harvard University Press, Cambridge, 1989. Typescript from the William James Lectures, Harvard University, 1967.

J. Groenendijk and M. Stokhof. *Studies in the Semantics of Questions and the Pragmatics of Answers*. PhD thesis, University of Amsterdam, 1984.

B. Grosz, A. Joshi, and S. Weinstein. Centering: a framework for modeling the local coherence of discourse. *Computational Linguistics*, 21:203–226, 1995.

J.Y. Halpern and Y. Moses. Towards a theory of knowledge and ignorance. In *Proceedings 1984 Non-monotonic reasoning workshop*, pages 165–193, New Paltz, NY, 1984. American Association for Artificial Intelligence.

P. Hendriks and H. de Hoop. Optimality theoretic semantics. *Linguistics and Philosophy*, 24:1–32, 2001.

J. Hirschberg. *A theory of scalar implicature*. PhD thesis, University of Pennsylvania, 1985.

J. et. al. Hobbs. Interpretation as abduction. *Artificial Intelligence*, 63:69–142, 1993.

W. van der Hoek et al. Persistence and minimality in epistemic logic. *Annals of Mathematics and Artificial Intelligence*, 27:25–47, 1999.

L. Horn. *The semantics of logical operators in English*. PhD thesis, Yale University, 1972.

M. Kameyama. Indefeasible semantics and defeasible pragmatics. In M. Kanazawa, C. Pinon, and H. de Swart, editors, *Quantifiers, Deduction, and Context*, pages 111–138. CSLI, Stanford, 1996.

H. Kamp. Free choice permission. *Proceedings of the Aristotelian Society, N.S.*, 74:57–74, 1973.

M Krifka. The semantics and pragmatics of polarity items. *linguistic Analysis*, 25:209–258, 1995.

F. Landman. *Events and Plurality*. Kluwer, Dordrecht, 2000.

H.J. levesque. All i know: a study in autoepistemic logic. *Artificial Intelligence*, 42:263–309, 1990.

S. C. Levinson. *Presumptive Meaning. The Theory of Generalized Conversational Implicatures*. MIT Press, Massachusetts, 2000.

S. C. Levinson. Pragmatic reduction of the biding conditions revisited. *Journal of Linguistics*, 27:107–161, 1991.

D. Lewis. *Counterfactuals*. Blackwell, Oxford, 1973.

J. McCarthy. Circumscription - a form of non-monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.

J. McCarthy. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence*, 28:89–116, 1986.

J. McCawley. *Everything that linguists always wanted to know about Logic**, volume 2. The University of Chicago Press, Chicago, 1993.

M. Moens and M. Steedman. Temporal ontology and temporal reference. *Computational Linguistics*, 14(2):15–27, 1988.

I. Noveck. When children are more logical than adults: experimental investigations of scalar implicatures. *Cognition*, 78:2165–188, 2001.

I. Noveck and A. Posada. Characterizing the time course of implicature: An evoked potential study. *Brain and Language*, 85:203–210, 2003.

A Papafragou and N. Tantalou. Children's computation of implicatures. *Language Acquisition*, 12:71–82, 2004.

S.G. Paris. Comprehension of language connectives and propositional relationships. *Journal of experimental child psychology*, 16:278–291, 1973.

N. Pouscoulous. *Processing scalar implicatures.* PhD thesis, EHESS, Paris, 2006.

T. Reinhart. *Anaphora and semantic interpretation.* Croom Helm, London, 1983.

M. Rooth. *Association with focus.* PhD thesis, University of Massachusetts, Amherst, 1985.

B. Russell. Against grammatical computation of scalar implicatures. *Journal of Semantics*, 23:361–382, 2006.

U. Sauerland. Scalar implicatures of complex sentences. *Linguistics and Philosophy*, 27: 367–391, 2004.

K. Schulz. A pragmatic solution for the paradox of free choice permission. *Synthese: Knowledge, Rationality and Action*, 147(2):343–377, 2005. see also Chapter **??**.

K. Schulz and R. van Rooij. Pragmatic meaning and non-monotonic reasoning: The case of exhaustive interpretation. *Linguistics and Philosophy*, 29(2):205–250, 2006.

A. Sevi. *Exhaustivity. A Semantic Account of 'Quantity' Implicatures.* PhD thesis, Tel Aviv University, Tel Aviv, 2006.

Y. Shoham. *Reasoning About Change.* MIT Press, Boston MA, 1988.

B. Spector. Scalar implicatures: Exhaustivity and Gricean reasoning? In B. ten Cate, editor, *Proceedings of the ESSLLI 2003 Student session*, Vienna, 2003.

K. Stenning and M. van Lambalgen. *Human Reasoning and Cognitive Science.* MIT Press, Cambridge, Massachusetts, 2008.

G. Storto and M.K. Tannenhaus. Are scalar implicatures computed online? In *Proceedings of WECOL*, 2004.

P. Suppes. Procedural semantics. In R. Haller and W. Grassl, editors, *Language, Logic, Philosopy*, pages 27–35. Hölder-Pichler-Tempsky, Vienna, 1980.

J. van der Does and M. van Lambalgen. A logic of vision. *Linguistics and Philosophy*, 23: 1–92, 2000.

J. van Kuppevelt. Inferring from topics. Scalar implicatures as topic-dependent inferences. *Linguistics and Philosophy*, 19:393–443, 1996.

M. van Lambalgen and F. Hamm. *The Proper Treatment of Events.* Blackwell Publishing, Oxford, 2005.

R. van Rooij and K. Schulz. Exhaustive interpretation of complex sentences. *Journal of Logic, Language, and Information*, 13:491–519, 2004.

F. Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–261, 1996.

J. Wainer. Modeling generallized implicatures using non-monotonic logic. *Journal of Logic, Language, and Information*, 16(2):195–216, 2007.

J. Wainer. *Uses of nonmonotonic logic in natural language understanding: generalized implicatures.* Ph.D. dissertation, Pennsylvania State University, 1991.

A. Zondervan. Experiments on qud and focus as a contextual constraint on scalar implicatures calculation. In U. Sauerland and K. Yatsushiro, editors, *Semantics and Pragmatics: From Experiment to Theory.* Palgrave MacMillan, in press.