A Note on the Absence of XOR in Natural Language[*]

Roni Katzir
Raj Singh
MIT

1 Introduction

Some connectives and quantifiers are systematically absent from the lexical inventories of natural language. This is an old observation, which was brought back to prominence within a modern setting by Horn (1972, 1989). The fundamental observation is that out of the four corners of the Square of Opposition, the two that correspond to $\exists$ and $\forall$ receive a priveleged status: the one that corresponds to $\neg\exists$ is sometimes found, though not always, and it tends to be morphologically marked (*n*-words); and the fourth corner, the one corresponding to $\neg\forall$ is never lexicalized. So, for example, we find *or* and *and*, as well as the *n*-construction *neither. . . nor*, but there is no **nand* (nor do we find *xor*). Similarly, we find *some* and *all*, and also the *n*-word *none*, but we do not find **nall* (*not all*; we also have no word for *some but not all*). In the domain of individuals, we find *a(n)* and *the* but neither **nthe* nor a lexical item that would mean *a but not the*. And the same holds for modals (*may*,*must*,**nmust*), adverbs (*sometimes*, *always*, *never*, **nalways*; *possibly*, *certainly*, **ncertainly*).

It should be noted that there does not seem to be a cognitive restriction that prevents us from dealing with the missing corner. The absent entries can readily be constructed and used analytically. Moreover, they are routinely arrived at via conversational reasoning (Grice, 1989 and much subsequent work). For example, *some but not all*, which entails *not all*, is often the preferred reading for *some*.

A proposal for deriving the typological pattern is presented by Horn (1989), who makes direct use of Grice's conversational principles, along with a general economy condition on inventories and a specific pressure against negation. It is precisely the ability of conversational reasoning to derive the missing corner, Horn suggests, that makes that corner redundant. By strengthening *some* to implicate *not all* (using *all* as a stronger scalar alternative) we can achieve almost the same effects as by lexicalizing *not all*. An economy condition on inventories now applies, penalizing systems that have a lexical *not all* in addition to the strengthenable *some*. But why these particular three corners? After all, the same reasoning could let us eliminate *some* using the corners *none*, *all*,

[*]We thank Asaf Bachrach, Danny Fox, Irene Heim, and Ivona Kučerová.

and *not all*.[1] Horn's answer is that negation is avoided whenever possible. The attested 3-corner system has $\neg$ as part of the *n*-word *none*, while the unattested one has $\neg$ both as part of *none* and as part of *not all*. The former system has fewer occurrences of negation, and consequently it is the winner.

We think Horn's intuition is correct, and we would like to use it to derive the typological pattern. As it stands, however, the idea faces certain challenges, which we think are inherent in the standard framework of quantification and negation in which Horn's proposal is couched. For example, in certain cases, the economy setting would lead us to expect a much more compact inventory than the one we actually find. What we have in mind here is the case of binary connectives, where the whole system can be defined in terms of just $\neg\wedge$. That is, we should expect to find *nand* and none of the other corners. Strangely, this is the complement of what we do find.

Another puzzling aspect of Horn's proposal is that it forces us to ignore the duality of $\exists$ and $\forall$, treating them instead as primitives. But the duality of these quantifiers is important. We only need one primitive quantifier. The other one can and should be defined using duality (either define $\exists := \neg\forall\neg$ or define $\forall := \neg\exists\neg$). Once we do that, however, we observe two new instances of negation in our inventory. How many negations we have in total now depends on whether we take $\exists$ to be the primitive quantifier or whether we choose $\forall$. If $\exists$ is primitive, the attested system of *some,all,none* has three instances of negation: two for $\forall$ and one for *none*. The unattested *all,none,not all* has four negations: two in *none* and one each of the other two. If we choose $\forall$ as our primitive quantifier, the attested trio will again have three negations. The unattested one, on the other hand, will now have only two negations. For Horn's proposal to go through, then, we would need to either ignore the duality of $\exists$ and $\forall$ or to stipulate that $\exists$ is primitive.

2 The proposal

We propose that Horn's idea can be made to work if instead of seeing the inventory of natural language operators through the quantificational prism of $\exists$ and $\forall$ we switch to talking about those operators in terms of min and max. The main difference between the two perspectives is that the quantificational approach relies on the notion of *truth*, while our min / max perspective relies on the notion of *ordering*. The domain of truth values is naturally ordered (with $0 < 1$), allowing for some overlap between the two perspectives. However, from the min / max perspective the domain of truth values is simply a special case, and the same inventory of operators can be applied to other domains, as long as an ordering is defined. Moreover, even within the domain of truth values (and more generally, domains of types that end in *t*) the min / max perspective will sometimes make different predictions than those of the quantificational one.

Before spelling out the details of our proposal, let us see how we could talk about a few familiar operators using min and max. Take *and*, for example. As a binary sentential connective, *and* takes two sentences and returns 1 if and only if both of its arguments have the truth value 1. In $\{\min, \max\}$ talk, this can be straightforwardly rephrased by saying that *and* returns the minimum truth value of its two arguments. Extending beyond two sentential arguments is not much more

[1] The other two options may or may not be possible. Much depends on what gets negated and what gets strengthened. Our presentation assumes, with Horn (1972), that only stronger alternatives are negated. This has been standard in the literature, but see Groenendijk and Stokhof (1984), van Rooij and Schulz (2004), Spector (2006), and Fox (2007a) for arguments that this assumption might have to be revised.

challenging: *and* still returns the minimum truth value of the set of truth values of its arguments. A lexical entry for *and* would be the following (where D_{t} is the domain of propositions):

(1) $[\![\text{ and }]\!] = \lambda P \in {D_{\mathrm{t}}}^{\mathrm{n}}.\min P$

Syntactically, *and* can take non-sentential arguments, as in *John and Mary are students*. A standard treatment of non-sentential coordination reduces these cases to the sentential case, either in syntax (by means of so-called *Conjunction Reduction*) or in semantics. Here is an entry based on the latter option, in analogy with the quantificational approach (cf. Partee and Rooth, 1983):

(2) $[\![\text{ and }]\!] = \lambda A \in {D_{\alpha}}^{\mathrm{n}}.\lambda f_{\alpha\mathrm{t}}.\min\{f(A_{\mathrm{i}}) : i = 1,\ldots n\}$

The treatment of *or* is entirely parallel. We only need to substitute max for min in all of its occurrences above. Generalizing to *every* and *some* is likewise straightforward. The thing to observe is that the entry for *and* involved two somewhat artificial constraints. First, the cardinality of A was limited (exactly n). Second, A was ordered (an n-tuple). For what we are trying to do, though, all we need is for A to be composed of elements, and for the application of the predicate f to those elements to yield something of type t. Once we notice this, we can simplify the way our entries look, and we can account for *every* (and *some*) along the way:

(3) $[\![\text{ every }]\!] = \lambda A \subseteq D_{\alpha}.\lambda f_{\alpha\mathrm{t}}.\min\{f(x) : x \in A\}$

We are almost ready to state our general constraint on operators. The last thing to get rid of in our entries is the restriction to types that end in t. For entries of the kind we just saw, any set that ends in a type for which min and max are defined would work just as well. D_{t} is one such type, and the natural ordering, $0 < 1$, yields the results we needed for things like *and, or, every,* and *some*, but we see no particular reason to stipulate that other types are prohibited. Here, then, is the form of a natural language operator (stated here in terms of $\mu \in \{\min,\max\}$):

(4) $[\![\ Op_{\mu}\ ]\!] = \lambda A \subseteq D_{\alpha}.\lambda f_{\alpha\beta} : \beta \in Dom(\mu).\mu\{f(x) : x \in A\}$

Using (4) we can state a preliminary constraint on lexicalization that handles morphologically simplex lexicalizations.[2]

(5) CONSTRAINT ON LEXICALIZATION (the basic case): Morphologically simplex operators in natural language are of the form Op_{μ}, as defined in (4), where $\mu \in \{\min,\max\}$

What counts as a legal input to μ (=min / max)? First, the input must be a set. Then, μ must be able to compare elements within the set, so it is good if the set is at least partially ordered. Finally, the set cannot be empty (what would it mean to return the maximal element of the empty set?). These are all near-trivial points, but they have some interesting consequences, often quite different from those of the quantificational approach. Let us see some examples.

3 Some immediate consequences

3.1 Other domains

The quantificational approach relies on the notion of truth: the contribution of a quantifier Q is defined in terms of the truth of substitutions to sentences in which Q appears. For example,

[2] In (9) below the constraint will be extended to accommodate *n*-words.

Someone smokes is true *iff* there is an x such that the sentence *x smokes* is true. Our proposal, on the other hand, makes no reference to truth. The primitive notion our entries rely on is ordering, and their application to propositions is done derivatively, using the ordering within the domain of truth values. We should expect to find operators similar to *every* and *some* over domains that do not end in t, as long as an ordering is defined. Individuals, for example, have often been taken to be collections ordered by the `part-of` relation, but they quite clearly do not end in t. A maximality operator along the lines of (4) but applying to the partially-ordered domain denoted by a noun will give us something very similar to Link (1983)'s notion of the definite article, one that can apply to both plural and singular individuals. An entry that uses (4) as a template could look like this:

(6) $[\![\max_{\mathrm{e}}]\!] = \lambda A \subseteq \{\emptyset\}.\lambda f_{<\{\emptyset\},\beta>} : \beta \in D_{\mathrm{e}}.\max\{f(x) : x \in A\}$

It is not obvious how the quantificational approach to *every* etc. could generalize to this domain.

It is interesting to note that while our proposal derives the entry for a definite article, it has nothing to say about indefiniteness. Replacing max with min in (6) would try to return the minimum element in a collection of individuals (ordered, as before, by `part-of`). When there is more than one atomic individual, min would be undefined, resulting in presupposition failure. When there is only one atomic individual, min would return exactly the same result that max would, namely that individual. Either way, there is no need to lexicalize min. We seem to predict, then, that natural language should have at most a definite article.

How bad is this prediction? We are not sure. Some languages do have what many people refer to as an indefinite article. What this indefinite article does, however, is less clear. Heim (1982), for example, argues that indefinites are not themselves operators (instead, they introduce variables). If this analysis is correct, it is probably a good idea to avoid predicting an indefinite counterpart to (6).[3] Typological evidence also seems to distinguish between the two articles. Many languages have a definite article but not an indefinite one. Even in languages like English, German, and French, where both articles exist in the singular, only the definite article makes it to the plural part of the paradigm. And in many languages (e.g., Danish, Greek) the indefinite article looks more like an adjective, often similar in form to the numeral *one*, than like a full fledged determiner. We are therefore happy to be able to derive the definite article, and not overly worried at this point about our inability to generate an indefinite article as its mirror image.[4]

3.2 Existential import

We mentioned above the rather unremarkable fact that min and max need a non-empty set as an argument. An immediate consequence of this is that whenever a natural language operator happens to receive an empty argument, the result is undefined. The arguments for *and* and *or* are provided explicitly, so we cannot use them to test this prediction. On the other hand, *every* and *some*, as well as *the*, take an intensional definition of their domain, in the form of their *NP* sister. We predict, then, that these elements will introduce an existential presupposition over the domain of their sister.

[3] Heim (1982) also argues against the view that definites are operators. Instead, she treats them as things that select a referent that has already been introduced. This may seem at odds with our operator-based definition in (6), but we think the two accounts are consistent.

[4] The indefinite article might still end up being definable in terms of min / max. If this is the case, though, we suspect that its domain will be different from that of the definite article and closer to the domains of adjectives and numerals, elements about which we have nothing to say at this point.

Adding an existence presupposition to the restrictor of a quantifier has been argued by Strawson (1952) to be necessary for capturing intuitively valid inferences in natural language, specifically the inference from $\forall$ to $\exists$. More recently, von Fintel (1999) has argued that this notion of entailment (his *Strawson Entailment*) is central to the proper treatment of NPIs. While the facts are complex, standard projection tests suggest that our prediction might be right:[5]

(7) Do you happen to know whether every/some/the delegate from France will need a translator?

If (7) is asserted, the hearer would be well within her rights to complain by saying 'Hey, wait one minute, I didn't know that there would be delegates from France.'[6]

To the extent that this argument is correct, the quantificational perspective fails to provide an explanation. A definition of $\exists$ and of $\forall$ that presupposes a non-empty domain is certainly possible, but nothing about the notions of existential and universal quantification requires that.[7] Within the $\forall, \exists$ framework, introducing Strawson entailment required a significant addition to the theoretical machinery; the min / max framework derives it.

4 Ordering and negation

4.1 Internal negation

The min / max approach relies on ordering, and in orderings, directionality matters: $< a,b > \in \leq$ is not the same as $< b,a > \in \leq$.[8] So far we assumed certain natural orderings over domains, but the directionality of those choices was largely a matter of convention. For any $\leq$ we can just as easily define an ordering that goes in the opposite direction: $\leq_{\text{new}} = \{< a,b >:< b,a > \in \leq\}$. Clearly, $\min_{\leq} \equiv \max_{\leq_{\text{new}}}$ (and similarly, $\max_{\leq} \equiv \min_{\leq_{\text{new}}}$). This suggests a role for a directionality-reversing operator, as well as a simple way to state the intuitive duality of min and max. Define an operator $\circ$ (pronounced *suhrk*) that reverses ordering relations in the way just described (that is, $\circ(R) \equiv \{< a,b >:< b,a > \in R\}$), and use the following entry for one of the operators, say for min:

(8) $[[\ \min_{\leq}\ ]] = \lambda X \in Dom(\max).\max_{\circ(\leq)}(X)$

Of course, we could also use $\circ$ to define max in terms of min. The important point is that our two operators are mutually interdefinable in terms of $\circ$ (which serves as some kind of internal negation), and that the relationship between them is captured by means of a simple ordering-reversal operation. The choice of what counts as min and what counts as max, then, depends on what we take to be the ordering relation $\leq$. Given the standard definitions of $<_{\text{t}} = \{< 0,1 >\}$ and $<_{\text{e}} = \texttt{part-of}$, max is used to define existential-like operators (as well as

[5] While much of the literature on the existential import of quantifiers has focused on truth-value intuitions, we believe that von Fintel (2004) is right in putting more faith in projection as a diagnostic. For discussion of the complexities involved in attributing existential presuppositions to quantifiers see de Jong and Verkuyl (1985), Lappin and Reinhart (1988), Diesing (1992), Heim and Kratzer (1998), Abusch and Rooth (2004), and Geurts (2007), among others.

[6] We follow von Fintel (2004) in using the *Hey, wait a minute!* test (HWAMT, modeled after Shanon, 1976) as a diagnostic for presupposition. See Singh (2008a) for a discussion of cases where the HWAMT seems to target additional entailments of the sentence. As far as we can tell, that discussion does not affect our current point.

[7] In fact, it is standard textbook practice to say that $\exists$ is false and $\forall$ is true when the domain is empty.

[8] In fact, by weak antisymmetry, the one would preclude the other, unless $a = b$.

the definite article, confusingly enough), while min is used to define universal-like operators (and perhaps the indefinite article). In terms of which operator should be our primitive one, nothing so far helps us choose, but we will soon see that it makes more sense to take max as our primitive operator, and to define min in terms of max and $\circ$, as outlined above.

Note that stating the duality of the traditional quantificational operators $\forall$ and $\exists$ requires two reversal operations, one on either side of the quantifier that is taken to be primitive. For example, $\forall \equiv \neg\exists\neg$. The possibility of inserting $\neg$ on either side of a quantificational operator suggests a four-member inventory of operators along the lines of the Square of Opposition: $\{Op, \neg Op, Op\neg, \neg Op\neg\}$. The intuitively priviliged status of the pair $\{Op, \neg Op\neg\}$ remains something of a mystery. From the min / max perspective, the privileged pair is $\{Op, Op\circ\}$, which results from the single choice of the direction in which $\leq$ is read.

A subtle point, closely related to the previous one is that our notion of duality depends not on the domain of the operators but on the symmetry inherent in the notion of ordering. In contrast, the duality of $\forall$ and $\exists$ is based on the notion of negation within the domain of truth values. As mentioned above, it is not obvious how the quantificational perspective could generalize to domains that do not end in t. Now we can say something stronger. Even if we were to define a variant of one of the propositional operators to a different domain, we wouldn't have an obvious way to make duality carry over. What does the negation of an individual mean, for example? The fact that duality does seem to hold seems to be an advantage of our system.

4.2 External negation

We mentioned that the $\leq$-reversing operator $\circ$ can be thought of as an internal negation, applying to the ordering argument of μ. It is important to notice that whenever μ is defined, $\circ$ will be defined as well. Differently from the traditional perspective, where $\exists$ and $\neg$ are primitives, on our account something like $\circ\mu\circ$ will not always be well-defined. For example, when applied to D_{e}, max takes a collection of individuals and their sums and returns its maximal element, if such an element exists. Using internal negation amounts to min, which may correspond to something like the indefinite article.[9] In any case, it seems rather pointless to try to externally negate the maximal element. max was defined on a collection, where the relation of `part-of` holds. Once we find the maximal element, we arrive at a singleton domain. No non-trivial ordering is definable on this element. It makes sense, then, to treat max and min as providing us with pairs of operators rather than with squares.[10] We can also see why languages lexicalize *the* but not **nthe*. The operator $\circ$ is an argument of μ, and may also be present syntactically in the right place for all we know. More importantly, regardless of the domain, the two operators resulting from using μ with and without $\circ$ are always well-defined. External negation, on the other hand, is not always well-defined, and in any case the external reversor will not be an argument of μ, and will probably not even form a constituent with it. There is no reason to lexicalize such a thing.

So we have pairs like *some* and *all*, and we don't have **nall*. What about *none*? The standard story predicted four operators per domain and had to explain why we find only three out of those four in natural language. Our system avoided that problem, predicting only the basic duality. This avoided the overgeneration problem of the standard story, but now we have to explain how to get

[9] Though, as mentioned earlier, this is probably incorrect.

[10] Contrasting with the $\neg, \exists$ perspective.

three operators out of two.[11] We said earlier that external negation is not always well-defined. But sometimes it is. Even then it would not be quite as organic as internal negation (the external operator is not a semantic argument of μ and does not form a constituent with it), but at least we can figure out what it does. In the domain of propositions, for example, μ returns a proposition. This can be reversed, assuming that a higher operator is sensitive to the natural ordering of truth values. External negation, in such cases, will amount to standard truth-conditional negation. If our basic μ is max, external negation will mean what the English *none* seems to mean. If our basic μ is min, external negation will mean what **nall* would have meant. It seems, then, that we should take max as primitive. As far as we can tell, all of the so-called *n*-words (*none*, *no*, *neither... nor*, *never*, etc.) go in the same direction.

(9) CONSTRAINT ON LEXICALIZATION (final version):

a. Basic case (repeated from (5)): Morphologically simplex operators in natural language are of the form Op_μ, as defined in (4), where $\mu \in \{\min, \max\}$

b. *N*-words: It is possible to lexicalize $Op_{\circ\max}$, and the result is an *n*-word

4.3 Back to the missing corner

We moved from a $\exists, \neg$ perspective to a $\max, \circ$ one so as to avoid overgeneration. We played with different ways to state our new inventory, but we just saw that using max as primitive might make more sense. The evidence, though, came from those cases where external negation is possible. By allowing external negation to be used in defining a lexical operator we have opened the door again to the Square of Opposition. If *external-not* can combine with max, why couldn't it combine with $\min(\equiv \max\circ)$? As suggested in the introduction, we believe that an answer can be found in Horn (1989). Recall that Horn attributes the missing corner of the Square to (the diachronic effects of) Gricean reasoning combined with language's abhorrence of negation: using $\exists, \forall$, and $\neg$ we get the attested three corners, $\{\exists, \forall, \neg\exists\}$; the fourth corner, $\exists\neg$, is derived by strengthening $\exists$ to imply $\neg\forall$. Within our framework, $\{\max, \circ\max, \max\circ\}$ is sufficient, while $\circ\max\circ$ ($\approx$ **nall*) can be derived by Gricean reasoning.

Horn notes that Gricean reasoning could also support a different system, where the lexicalized operators are $\{\forall, \neg\forall, \forall\neg\}$, and where the missing corner is $\exists$. Systems of this kind, however, are empirically unattested, and Horn suggests that the reason is a markedness constraint that penalizes negation in lexical inventories: the unattested system has two negations, compared to one in the attested system. We mentioned, however, that while Horn's intuition seems compelling, using it within his framework would force us to ignore the inherent duality of $\exists$ and $\forall$. If we want to be able to express the duality of $\exists$ and $\forall$, maintaining Horn's idea of negation avoidance requires the stipulation that $\exists$ is primitive.

Within our system, on the other hand, Horn's idea of avoiding negation yields the correct results without any such stipulation. The attested system, as mentioned, is $\{\max, \circ\max, \max\circ\}$

[11] Maybe we don't have to. Penka (2007) has argued that negative quantifiers are never part of natural language. She treats *n*-words as uniformly indefinite, attributing their apparent negative behavior to agreement with a higher negative operator. If she is right, our $\{\min, \max\}$ would be the full inventory of lexicalizable operators. However, this would still not account for the systematic absence of universal-like *n*-words. All things considered, then, it would be good to find a more principled explanation for the distribution. As mentioned above, we will do this by following Horn (1989)'s proposal.

(with $\circ\max\circ \approx$ **nall* derived Griceanly). The other 3-way system based on max is $\{\max\circ, \circ\max\circ, \circ\max\}$ (with max $\approx$ *some* derived Griceanly), but this system has four instances of $\circ$, and hence loses to the attested system, which has only two. In this respect, the min / max setting is similar to the quantificational one. In contrast with the quantificational setting, however, switching to a system where min is primitive (rather than max) no longer helps. The two contenders are $\{\min\circ, \circ\min\circ, \min\}$ and $\{\min, \circ\min, \circ\min\circ\}$, each of which has three instances of $\circ$, hence losing to the attested, max-based system.

5 Non-classical logics

In talking about the propositional domain, we restricted ourselves to classical, two-valued logic. Even there, as is obvious from even a cursory glance at the long and controversial history of the Square, there was ample space for logically possible but empirically unattested operators. Moving to non-classical systems adds more possible states for each proposition, an increase that gets polynomially amplified when several propositions are combined by a connective or a quantifier. The space of logically possible connectives increases exponentially with the size of the table. This is not just a conceptual issue. Non-classical logics have often been argued to be the proper way to account for presupposition. It has been supposed that presupposition corresponds to partiality in the domain of a function, a new truth value (say, 0.5), or uncertainty about which of the two classical truth values hold. These perspectives are interestingly different in many ways, but for our purposes, to which we will return shortly, a crucial feature they share in common is that this new state, which we will designate neutrally as $\circledast$, can be directly compared with the classical values, and that $0 \leq \circledast \leq 1$. The potential exacerbation of the overgeneration problem has been a serious concern, as discussed in detail by Soames (1982, 1989) and Heim (1983, 1990), and more recently by Schlenker (2008a), Fox (2007b), and George (2008). To take a simple example, the classical unary operator $\neg$ is one of $2^2 = 4$ possible operators of this kind. Adding $\circledast$ increases the number of such operators to $3^3 = 27$, of which 3 overlap with classical negation with respect to 0 and 1. More dramatically, the classical binary propositional connective *and* is one of $2^{(2\times2)} = 16$ possible connectives of this kind. $\circledast$ expands the space to $3^{(3\times3)} = 19,683$ possible connectives, of which $3^{(9-4)} = 243$ are consistent with classical *and*. Even if one can think of an explanation for why we find the operators that we do, why don't we also see some of the other possible ones, at least occasionally?

Once we switch from $\{\exists, \forall\}$ to $\{\min, \max\}$, the problem with adding $\circledast$ disappears. The only unary operator is reversal, $\circ$, which gives rise to the ordering $1 \leq {}_\circ\circledast \leq {}_\circ 0$. From 19,683 possible tables for a binary propositional connective we are down to exactly two:

(10) a. *and*

min	0	$\circledast$	1
0	0	0	0
$\circledast$	0	$\circledast$	$\circledast$
1	0	$\circledast$	1

b. *or*

max	0	⊛	1
0	0	⊛	1
⊛	⊛	⊛	1
1	1	1	1

No other binary propositional operator is definable using min and max. Under our earlier assumption that these are the only possible entries for morphologically simplex operators in natural language, we have nothing to worry about. We mentioned above that external negation can sometimes be used, and we suggested that when this happens it is by combining with max (rather than with min $\equiv$ max $\circ$) and that the result is a morphologically complex *n*-word. For this, too, overgeneration doesn't arise. The only possibility is this:

(11) *nor*

∘ max	0	⊛	1
0	1	⊛	0
⊛	⊛	⊛	0
1	0	0	0

Interestingly, our system not only avoids the overgeneration problem, it coincides exactly with the Strong Kleene system, a historically significant contestant for a descriptively adequate framework for presupposition. Moreover, it extends straightforwardly to any number of arguments (including an unbounded set, as in quantifiers) as well as to any set of truth values (finite or not) that has the relevant kind of ordering. Unfortunately, it also goes against a long history of work on presupposition that demonstrates that the empirical patterns of projection involve asymmetries between the arguments of connectives. Here are some early examples.

(12) (Horn, 1972 ex. 2.6)

a. []John is a man and (he is) a bachelor

b. [#]John is a bachelor and (he is) a man

(13) (Karttunen, 1973 ex. 16)

a. []Jack has children and all of Jack's children are bald

b. [#]All of Jack's children are bald, and Jack has children

Such examples have motivated the incorporation of asymmetry into the basic mechanism of presupposition and its projection. Within frameworks that rely on non-classical logics, from Peters (1979) to the more recent proposals of Beaver and Krahmer (2001), Fox (2007b), and George (2008), this has motivated asymmetric truth tables, often following an intuition that the left-to-right order in which arguments are presented has truth-functional consequences. As should be clear by now, our system has no way to produce an asymmetric table.

This is hardly the place for a detailed discussion of the intricacies of presupposition projection, but we would like to mention some of our reasons to think that our symmetric system is on the right track.[12] First, to the extent that our diagnostics can serve to probe non-classical truth values, there

[12]For more detailed proposals involving symmetric projection see Chemla (2008) and Schlenker (2008b). Both follow Horn (1972) in relegating the above asymmetries to a presupposition-independent constraint against incremental redundancy. We will do the same in (18) below.

is no difference between the two propositional arguments in (12) and (13). For example, while the two sentences in (13) differ in felicity, uttering 'Hey, wait a minute! I didn't know that Jack had children,' would be just as inappropriate as an objection to (b) as it would be for (a). Could it be that the inapplicability of the *HWAM!* test for (b) has to do with the oddness of the original sentence? We think not. It is sometimes possible to get rid of the oddness of Karttunen-like sentences by embedding them, as can be seen in (14) below. This is in itself a surprising fact, but more importantly, the *HWAM!* test remains inapplicable.

(14) a. There are exactly five workers in this factory [who LOVE their wife] and [who HAVE a wife]
b. There are exactly five workers in this factory [who HAVE a wife] and [who LOVE their wife]

(15) [#]Hey, wait a minute! I didn't know that someone in this factory has a wife.[13]

So it looks like the left-to-right asymmetry in Karttunen's sentences, while real, has little to do with the truth-functional component. What is the asymmetry about, then? We are not sure, but left-to-right asymmetries have been observed in many other domains that could be loosely grouped together under the heading of discourse appropriateness.[14] Grice (1989), to cite an early example, noted that the order in which conjuncts are presented is taken to imply a temporal or causal order between their denotations. Thus the two sentences in (16) both entail the same two events but suggest different orders in which they took place.

(16) a. They got married and had a child
b. They had a child and got married

And in the domain of implicature, it has been argued by one of us (Singh, 2008b) that linear order corresponds to the strength of the scalar items that appear in disjunctive sentences like (17):[15]

(17) a. []He talked to Sue or to both Sue and Mary
b. [#]He talked to both Sue and Mary or to Sue

We think that it would be short-sighted to confine a left-to-right mechanism to presupposition, especially since it appears to be irrelevant there, while phenomena across the pragmatic board seem to have access to such a mechanism. For our purposes we can appeal to a non-redundancy condition, as proposed by Horn (1972:ex. 2.12):[16]

(18) INCREMENTAL NON-REDUNDANCY: The second conjunct Q of a conjunction $P\&Q$ must not follow from the first conjunct P

[13]The precise presupposition that would project in this context is a matter of some controversy. We chose to use *HWAM!* to test for the weakest presupposition that could in principle project here.

[14]We use the term 'discourse' in an imprecise way. Much of what we are about to see has been argued, from Cohen (1971) onward, to take place at local and more mechanical levels. We trust that no confusion should arise.

[15]See Fox and Spector (2008) for a different perspective on such sentences, though their characterization preserves the dependence on linear order.

[16]See van der Sandt (1992) and Schlenker (2008b,a) for recent discussion.

References

Abusch, Dorit, and Mats Rooth. 2004. Empty-domain effects for presuppositional and non-presuppositional determiners. In *Context-dependence in the analysis of linguistic meaning*, 7–27. Elsevier.

Beaver, David, and Emiel Krahmer. 2001. A partial account of presupposition projection. *Journal of Logic, Language and Information* 10:147–182.

Chemla, Emmanuel. 2008. Similarity: Towards a unified account of scalar implicature, free choice permission and presupposition projection. Ms., Ecole Normale Supérieure and MIT, January 2008.

Cohen, Laurence Jonathan. 1971. Some remarks on Grice's views about the logical particles of natural language. In *Pragmatics of natural language*, ed. Yehoshua Bar-Hillel, 50–68. Reidel.

Diesing, Molly. 1992. *Indefinites*. Cambridge, Mass.: MIT Press.

von Fintel, Kai. 1999. NPI licensing, Strawson entailment, and context dependency. *Journal of Semantics* 16:97–148.

von Fintel, Kai. 2004. Would you believe it? The King of France is back! Presuppositions and truth-value intuitions. In *Descriptions and beyond*, ed. Marga Reimer and Anne Bezuidenhout. Oxford University Press.

Fox, Danny. 2007a. Free choice disjunction and the theory of scalar implicatures. In *Presupposition and implicature in compositional semantics*, ed. Uli Sauerland and Penka Stateva, 71–120. Palgrave-Macmillan.

Fox, Danny. 2007b. Schlenker 2006. Lecture notes for 24.954 Pragmatics, December 2007.

Fox, Danny, and Benjamin Spector. 2008. Economy and embedded exhaustification. Handout of a talk presented at the Syntax-Semantics Reading Group at MIT, May 2008.

George, Benjamin. 2008. Predicting presupposition projection: Some alternatives in the strong Kleene tradition. Ms., UCLA, March 2008.

Geurts, Bart. 2007. Existential import. In *Existence: Semantics and syntax*, ed. Ileana Comorovski and Klaus von Heusinger, 253–271. Dordrecht: Springer.

Grice, Paul. 1989. *Studies in the way of words*. Cambridge, MA: Harvard University Press.

Groenendijk, Jeroen, and Martin Stokhof. 1984. On the semantics of questions and the pragmatics of answers. In *Varieties of formal semantics*, ed. Fred Landman and Frank Veltman, Vol. 3 of *Groningen-Amsterdam Studies in Semantics*, 143–170. Dordrecht: Foris.

Heim, Irene. 1982. The semantics of definite and indefinite Noun Phrases. Doctoral Dissertation, University of Massachusetts, Amherst.

Heim, Irene. 1983. On the projection problem for presuppositions. In *WCCFL*, 114–125.

Heim, Irene. 1990. E-type pronouns and donkey anaphora. *Linguistics and Philosophy* 13:137–178.

Heim, Irene, and Angelika Kratzer. 1998. *Semantics in generative grammar*. Malden, MA: Blackwell.

Horn, Laurence. 1972. On the semantic properties of the logical operators in English. Doctoral Dissertation, UCLA.

Horn, Laurence. 1989. *A natural history of negation*. Chicago: University of Chicago Press.

de Jong, F., and Henk Verkuyl. 1985. Generalized quantifiers: The properness of their strength. In *Generalized quantifiers in natural language*, ed. Johan van Benthem and Alice ter Meulen, 21–43. Dordrecht: Foris.

Karttunen, Lauri. 1973. Presuppositions of compound sentences. *Linguistic Inquiry* 4:169–193.

Lappin, Shalom, and Tanya Reinhart. 1988. Presuppositional effects of strong determiners: A processing account. *Linguistics* 26:1021–1037.

Link, Godehard. 1983. The logical analysis of plurals and mass terms: A lattice-theoretic approach. In *Meaning, use, and interpretation of language*, ed. R. Bäuerle, C. Schwarze, and A. von Stechow, 303–323. De Gruyter.

Partee, Barbara Hall, and Mats Rooth. 1983. Generalized conjunction and type ambiguity. In *Meaning, use and interpretation of language*, ed. R. Bäuerle, C. Schwarze, and Arnim von Stechow, 362–383. de Gruyter.

Penka, Doris. 2007. Negative indefinites. Doctoral Dissertation, Universität Tübingen, Tübingen.

Peters, Stanley. 1979. A truth-conditional formulation of Karttunen's account of presupposition. *Synthese* 40:301–316.

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

van der Sandt, Rob. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9:333–377.

Schlenker, Philippe. 2008a. Be articulate: A pragmatic theory of presupposition projection. Ms., To appear in Theoretical Linguistics, April 2008.

Schlenker, Philippe. 2008b. Local contexts. Ms., Institut Jean-Nicod and NYU, February 2008.

Shanon, Benny. 1976. On the two kinds of presuppositions in Natural Language. *Foundations of Language* 14:247–249.

Singh, Raj. 2008a. Modularity and locality in interpretation. Doctoral Dissertation, MIT, Cambridge, Mass.

Singh, Raj. 2008b. On the interpretation of disjunction: Asymmetric, incremental, and eager for inconsistency. To appear, Linguistics and Philosophy, May 2008.

Soames, Scott. 1982. How presuppositions are inherited: A solution to the projection problem. *Linguistic Inquiry* 13:483–545.

Soames, Scott. 1989. Presupposition. In *Handbook of philosophical logic*, ed. Dov Gabbay and Franz Guenther, Vol. IV, 553–616. Dordrecht: Reidel.

Spector, Benjamin. 2006. Aspects de la pragmatique des opérateurs logiques. Doctoral Dissertation, Université de Paris 7, Paris.

Strawson, Peter Frederick. 1952. *Introduction to logical theory*. London: Methuen.