A note on the cardinalities of sets of scalar alternatives

Salvador Mascarenhas
Institut Jean-Nicod, Département d'Etudes Cognitives,
ENS, EHESS, CNRS, PSL University

May 2021

Accepted for publication by *Journal of Semantics*, May 2021

Acknowledgments: I thank Benjamin Spector, Emmanuel Chemla, Philippe Schlenker, Philip Johnson-Laird, and Maria Aloni and two anonymous reviewers for detailed and extremely helpful feedback. This work was funded in part by *Agence Nationale de la Recherche* grants ANR-18-CE28-0008 (LANG-REASON, PI: the author) and ANR-17-EURE-0017 (FrontCog, Department of Cognitive Studies, Ecole Normale Supérieure).

1 Introduction

Scalar implicatures as in (1) are cancellable inferences that appear to arise from considering certain alternative sentences.

(1) I saw John or Mary.
Alternative: I saw John and Mary.
Scalar Implicature: The speaker did not see both John and Mary

When first described by Paul Grice (1967, published in 1989), scalar implicatures were taken to be the result of processes of pragmatic inference. In a nutshell: starting from what the speaker of (1) *literally said*, the hearer will consider what stronger statements the speaker *chose not to utter*, and conclude that the speaker believes all such statements to be false, as long as their falsehood is compatible with what she literally said. More recent approaches reject the pragmatic dimension of the phenomenon, proposing unpronounced operators that derive scalar implicatures as a matter of literal meaning rather than the result of pragmatic inferences (Chierchia et al. 2012).

In this note, I show that there is one technical component, universally employed in formal theories of scalar implicature, that has very puzzling consequences under the assumption that these theories ought to be psychologically tenable. The issue can be summarized as follows. Every modern theory of scalar implicature makes crucial use of a set of scalar alternatives, sentences that are in a precise sense related to the sentence uttered by a speaker. These are the alternatives that a hearer will take into consideration when thinking of what the speaker could have said but chose not to. They are the alternatives that unpronounced operators in non-pragmatic approaches must access in order to deliver truth conditions for sentences containing them. I show that the cardinalities of these sets increase at very fast rates, and moreover that, even for sentences with a relatively small number of coordinated clauses, the cardinalities of alternative sets are very large numbers. If the theories of alternatives considered here are making claims about psychological processes, then these claims are very difficult to square with what we know about alternatives in grammar and in reasoning. On the other hand, if these theories are to be taken as mathematical idealizations or theories of pragmatic competence, then it becomes necessary to investigate what psychologically tenable heuristics might implement this competence. It does not follow from the facts I report here that modern formal approaches to scalar implicature are doomed. Instead, I point out a collection of puzzling and previously unnoticed facts that require answers.

2 Alternative sets and their cardinalities

I take the work of Chierchia (2004), Sauerland (2004), Spector (2007), and Fox (2007) to be representative examples of formal approaches to scalar implicature. All of these theories assume or propose some mechanism that generates formal alternatives as a function of a source and possibly a context. I restrict my attention to explicit proposals about the set of formal alternatives that have not been falsified at the time of writing.[1] There are three approaches in the literature that fit this criterion.

2.1 Positive propositions

While making no commitments about the inner workings of the alternative-generating procedure, Spector's (2007) theory makes use of the set of *positive propositions* that can be constructed from a source S. The notion is simple: given a source S with a set A of atomic propositions, the set $P^+(A)$ of positive propositions based on A is the closure of A under conjunction and disjunction. Thus, for a simple source $S = a \vee b$, we get the alternative set

$$P^+(\{a,b\}) = \{a, b, a \wedge b, a \vee b\}.$$

The set of positive propositions based on a set of atoms A corresponds to the set of non-constant monotonic (increasing) Boolean functions of $|A|$-many variables. A Boolean function is monotonic if switching any of its inputs from false to true can never bring about a switch from true to false in the output. Thus, conjunction and disjunction are monotonic functions in this sense, but not negation or implication. Monotonic Boolean functions on n variables are precisely those that can be defined with the Boolean meet and join. The "non-constant" proviso simply excludes the *verum* and *falsum* functions from the list.[2]

The number of monotonic Boolean functions on n variables $M(n)$ was defined by Dedekind in 1897 (see for example Kleitman 1969). Accordingly, the set of all $M(n)$

[1]For example, the classical theory by (Horn 1972) is an explicit theory of alternatives in terms of substitutions of connectives from lexically determined scales, but it fails to generate alternative sets rich enough to derive implicatures as in (8), an example discussed at the end of this article.

[2]See also for example the article "Monotone Boolean function" in the online *Encyclopedia of Mathematics* (Springer) for a connection between monotone Boolean functions and disjunctive normal forms without negations, i.e. positive propositions.

for $n \in \mathbb{N}$ is known as the set of *Dedekind numbers*. The sequence grows very rapidly. As of early 2021, no closed-form expression for $M(n)$ is known, and exact values for $M(n)$ are only known for $n \leq 8$. The set of positive propositions with n atoms inherits the mathematical unwieldiness of Dedekind numbers, since, for A a set of atomic propositions, $|P^{+}(A)| = M(|A|) - 2$.

2.2 Syntactic substitution approaches

2.2.1 Sauerland alternatives

Sauerland (2004) proposes a syntactic substitution mechanism where only terminal nodes corresponding to binary connectives are the targets of substitutions. For σ a parse tree, the set of Sauerland alternatives consists of all possible trees created by substituting for each binary connective in σ a connective from the set $\{\wedge, \vee, L, R\}$. The connectives L and R stand for *left* and *right* respectively, and take their names from their semantics: $[\![\varphi\ L\ \psi]\!] = [\![\varphi]\!]$, and conversely for R.

To allow for straightforward comparison with Spector's (2007) positive propositions, I consider only trees σ that contain no repetitions of atomic propositions. This way we can express the cardinality of a set of alternatives as a direct function of the number of distinct atoms in both approaches. For a σ with n atoms, we find $n-1$ positions for binary connectives. Since we have four possible connectives to plug into each of these positions ($\{\wedge, \vee, L, R\}$), we find 4^{n-1} Sauerland alternatives for a tree with n atoms.[3] Sauerland alternatives therefore grow at an exponential rate.

2.2.2 Katzir and Fox complexity

Katzir (2007) and Fox and Katzir (2011) propose a different mechanism for generating formal alternatives based on a notion of structural complexity. As explained by Katzir (2007), there are contexts where, in order to derive the observed scalar implicatures, a non-monotonic expression must be compared with a formal alternative that is structurally simpler than it without being its scale-mate. For example:

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

An utterance of (2) suggests that the speaker does not find it unlikely that at least three semanticists will be in attendance. This means that (3) must be a formal alternative to (2). But the quantifier in (3) is not a scale-mate of the one in (2). The relevant relation between the two sentences seems to be one of complexity: (3) is less complex than (2), and therefore it is an alternative relevant for pragmatic computations.

Katzir's (2007) theory is given in a fully explicit way. First we need to define a substitution source. This will contain all elements that can legally substitute elements of the

[3] These are *syntactic* alternatives, they may well include alternatives that are equivalent while being syntactically distinct. Indeed, Spector (2016) has shown that there are $\frac{3^n-1}{2}$ equivalence classes of Sauerland alternatives. Consequently, for all cases except the trivial $n = 1$, the equivalence classes of Sauerland alternatives will be smaller than the syntactic alternatives.

speaker's utterance when generating alternatives. The rationale is that nothing in the substitution source can be *more* complex than the original utterance.

(4) Substitution source:
Let σ be a parse tree. The substitution source for σ is the union of the lexicon of the language with the set of all subtrees of σ.

Next, we define a relation of structural complexity on syntactic structures. Finally, we define the set of structural alternatives in terms of structural complexity.

(5) Structural complexity:
Let σ, σ' be parse trees. If we can transform σ into σ' by a finite series of deletions and replacements of constituents in σ with constituents of the same category taken from the substitution source of σ, we say that σ' is at most as complex as σ, in symbols $\sigma' \preceq \sigma$.

(6) Structural alternatives:
Let σ be a parse tree. The set of structural alternatives for σ is defined as $A(\sigma) := \{\sigma' : \sigma' \preceq \sigma\}$.

The alternative sets produced by this procedure are surprisingly large even for very small numbers of atoms in the source. In order to make discussing concrete examples easy, let us apply this theory to a propositional language, where in particular each binary connective is immediately dominated by a three-way branching node with further trees in the left and right branches, and a binary connective in the middle. Then, even if we exclude the lexicon (other propositional atoms) from the substitution source (4), a source $\sigma = a \wedge b$, with structure $[[a] \wedge [b]]$, will have 20 syntactically distinct alternatives. As a reference point, there are 16 possible propositions that can be formed out of two atoms, and only four positive propositions. It is instructive to go over a simple example of the procedure at work, so I will presently prove that there are indeed 20 alternatives for this source.

Ignoring the connective for now, there are three nodes in $a \wedge b$ that can be the target of a substitution: the root of the tree, the left node, and the right node. There are three subconstituents in the source: a, b, and $a \wedge b$. This means:

1. Substitutions at the root return three possible trees: a, b, and $a \wedge b$, in other words the set of subconstituents in the substitution source themselves.

2. Substitutions at the left and right nodes (still ignoring the connective) amount to all ordered pairs that can be formed from elements of the substitution source. The substitution source has three elements, so we get $3^2 = 9$ possible substitutions. For example, and using $\otimes$ as a placeholder for the binary connective we are still ignoring at this stage: $[a \otimes a]$, $[b \otimes a]$, $[(a \wedge b) \otimes b]$.

3. These nine possible substitutions each have one slot for a connective, represented as $\otimes$ in the examples just given. This is the only possible target for substitutions of connectives. Assuming two possible connectives ($\wedge, \vee$) for each of these nine trees, we get $2 \cdot 9 = 18$ distinct trees.

4. One of the trees generated in step 3. is $a \wedge b$, from $[a \otimes b]$ in step 2. with substitution

of $\wedge$ for $\otimes$. This formula was already counted in step 1. So we remove this tree from the count of step 1., and get a total number of $2 \cdot 9 + 2 = 20$ distinct trees.

In the example I just gave, I did not use the deletion transformation made available by the definition in (5). This was legitimate because any deletions of nodes from a source $a \wedge b$ will be identical to substitutions of one of the terminal nodes for the root of the tree. Deletions in this case were subsumed under a special case of substitutions. This is not so in the general case, and indeed tallying up all of the alternatives that are generated by the exact procedure in (5) is harder than it might seem. Since the precise number is not of intrinsic interest, I will instead establish a clear and intuitive lower bound for these syntactic alternatives. I make the following simplifying assumptions.

1. Atomic propositions occur only once in a parse tree. This assumption serves the same purpose as above, in the case of Sauerland alternatives. It allows for a straightforward comparison with the semantic approach of Spector (2007).

2. I adapt the definition in (5) to outlaw (A) deletions and (B) substitutions of elements for points in the tree other than terminal nodes. That is: for σ the source, we say that σ' is a less-complex transformation of σ just in case σ' is the result of substituting subconstituents of σ for terminal nodes of σ. Note that this algorithm is a restriction of the original. It will always generate a subset of the alternatives generated by the original procedure. Indeed, in all but the trivial case, it will generate a proper subset of what the original procedure generates.

With these assumptions, we can show that we get, for a source with n atoms,

$$(2n-1)^n \cdot 2^{n-1}$$

distinct syntactic alternatives.[4] Here is why:

1. There are n terminal nodes and $2n-1$ subconstituents to be substituted with repetitions for any of those terminal nodes.[5] Whence $(2n-1)^n$.

2. There are $n-1$ positions in which to place one of two binary connectives, so there are 2^{n-1} possible strings of connectives to be plugged in, for each of the connective-less tree-skeletons given by the term in 1. above. Whence the multiplication by 2^{n-1}.

[4]This expression provides a lower bound for the theories of alternative generation by Katzir (2007) and Fox and Katzir (2011). It is an open question exactly how individual variants of these theories will fare. For example, Trinh and Haida (2015) propose a constraint on these theories that potentially will prune elements from the substitution source, making it smaller, without ever adding elements to it. Consequently, at least in some cases, Trinh and Haida (2015) will generate fewer alternatives than the original theories. Yet, since I have shown that the syntactic alternatives are an exponential function of the number of elements in the substitution source, the theory by Trinh and Haida (2015) will still generate sets of alternatives that grow exponentially with the size of the substitution source.

[5]This is perhaps easiest to see for a simple propositional language exclusively with binary connectives. We can define a function s recursively that returns the number of subconstituents of a formula in this language: $s(p) = 1$ for p an atom, $s(\varphi \otimes \psi) = s(\varphi) + s(\psi) + 1$ for $\otimes$ a binary connective. We can also define a function a that returns the number of atoms (i.e. terminal nodes) in a formula: $a(p) = 1$, $a(\varphi \otimes \psi) = a(\varphi) + a(\psi)$. It's easy to prove by induction on the complexity of formulas that $s(\varphi) = 2a(\varphi) - 1$.

3 Discussion

Table 1 summarizes the findings in this article. It includes the total number of propositions made out of n atoms merely to provide an anchor that will be familiar to the reader.

Table 1: Number of alternatives by procedure for a source with 2, 3, 4, and n atoms.

	2	3	4	n
Propositions	16	256	65,536	2^{2^n}
Positive propositions	4	18	166	Dedekind numbers: $M(n)-2$
Sauerland (2004)	4	16	64	4^{n-1}
Katzir (2007) lower bound	18	500	19,208	$(2n-1)^n \cdot 2^{n-1}$

For all three algorithms, the size of the output set of alternatives increases fast as a function of the number of atoms in the input. This is particularly striking even for small numbers in Table 1 in the case of Katzir's (2007) theory. Indeed, in this theory the number of predicted alternatives grows at an exponential rate. While less visible within the range of small numbers in Table 1, the Sauerland algorithm is still clearly an exponential function of the length of the input. The positive-propositions approach fares no better, and grows at an exponential rate as well (Kleitman 1969).[6]

Thus, the three theories of alternative generation considered here have outputs that grow at least as an exponential function of the size of the input. These three theories are quite distinct, and they constitute, as far as I can tell, the most recent explicit algorithms for generating alternatives. Since these theories are representative of the state of the art, it is significant that all three algorithms are intractable.[7]

3.1 Algorithmic-level accounts and the values for small numbers

The values seen in Table 1 for small values of n are already rather large. This is again less obvious for the Sauerland and positive-propositions approach, which seem manageable even for $n = 4$ when compared to the syntactic approach. I submit that these values for small numbers are interesting in their own right, irrespective of considerations of relative growth.

[6] As mentioned before, no closed-formed expression is known for calculating Dedekind numbers. Kleitman (1969) attributes the earliest proof of a lower-bound for Dedekind numbers to Gilbert (1954). Gilbert found that the nth Dedekind number is greater than or equal to $2^{\binom{n}{n/2}}$, where $\binom{a}{b}$ notates a binomial coefficient. Later work improves on this lower bound (Kleitman 1969; Kleitman and Markowsky 1975), but the earlier lower bound is already clearly exponential.

[7] Ideally, one would want to show that the problem of generating alternatives is in itself intractable, not just that our current best algorithms are inefficient. Unfortunately, such a result is not forthcoming. There is simply nowhere near enough of a consensus in the field about precisely what the empirical scope of theories of scalar implicature ought to be. Consequently, we cannot articulate a general computational-level theory of alternative generation of which the three theories considered here would be algorithmic-level instantiations. Instead, these three theories constitute competing computational-level theories themselves.

The output of an algorithm may grow as an exponential function of the input without ipso facto being cognitively untenable. It suffices that only small enough numbers ever occur in practice. However, there is good reason to think that we need for our theories to handle a range of numbers that already produce alternative sets that are far too big.

For concreteness, Table 1 tells us that the most economical alternative-generating algorithm will produce 64 alternatives to a sentence as in (7) below.

(7) Ann and Bill, or Carl and Dan will come to the party tonight.

While it is clear that (7) is not a very easy sentence to parse, it is by now well established that speakers of English can process and reason systematically on the basis of sentences with this kind of structure. Walsh and Johnson-Laird (2004) studied deductive reasoning problems with premises with four atomic propositions along the lines of (7). Moreover, the mistakes that humans make when reasoning with and about sentences like (7) constitute precisely the kind of inference-making behavior that we would expect if they were interpreting (7) with a strong scalar implicature to the effect that "*Only* Ann and Bill, or else *only* Carl and Dan will come to the party tonight" (Mascarenhas 2014). This means that it is entirely fair to expect our theories of scalar implicature to apply to sentences with four atoms as in (7). It also suggests that we can ask the same questions about sentences with five or more atoms, where the numbers of predicted alternatives become even greater.

The literature on the psychology of reasoning offers theories of reasoning with and about alternative propositions. In particular, mental model theory (Johnson-Laird 1983) is an account of the human faculty of reasoning whose hallmark is precisely a theory of how humans entertain and manipulate alternative propositions, called alternative mental models. It is reasonable to ask whether the formal alternatives in theories of scalar implicature fit in the mental models picture of reasoning with and about alternative mental models.

The short answer is no. Mental model theory has established correlations between the number of alternative mental models under consideration and humans' performance in reasoning tasks. Every study within this paradigm indicates that humans can efficiently reason with/about rather small numbers of alternative mental models at any given point in time, with five to seven concurrent mental models being the limit.[8] These results cannot be squared with the predictions of even the most economical of current theories applied to a sentence as in (7).

This observation is immediately problematic for pragmatic accounts, which propose that an online inferential process is responsible for scalar implicatures. If this process is just general-purpose reasoning, then it flies in the face of what the only extant account of reasoning with alternatives from psychology has found. If this process is in fact not general-purpose reasoning, but part of a module specially dedicated to pragmatic reasoning, then it seems to have properties dramatically different from those of general-purpose reasoning, a puzzling fact in its own.

[8] For a general discussion of some of these results, see Johnson-Laird (2008). For correlations between the number of mental models under consideration and reaction times for very small numbers (one vs. two models), see for example Walsh and Johnson-Laird (2004).

The situation is different for grammatical approaches. Results from mental-model theory are of no consequence here, since grammatical approaches propose that the faculty of language is responsible for the relevant computations, and not a reasoning faculty of a general or special kind. Indeed, when framed within the faculty language instead of reasoning, the puzzle I raise here becomes a very general question, since alternatives just like those in scalar implicature play a role in the semantics of "only" and other related operators. What specific mechanisms allow the faculty of language to handle exponentially growing numbers of alternatives so seamlessly?

3.2 Alternative generation at the computational level

To the best of my understanding, the theories reviewed here were not conceived as algorithmic-level theories. Instead, they were designed as part of computational-level theories of pragmatics. At this level of analysis, exponential blowup is not necessarily a damning feature, provided that these computational-level theories can be coupled with algorithmic-level implementations that are tractable.

Indeed, we know that these computational-level theories can in principle be dramatically reduced in (algorithmic) practice. The syntactic approach for one will generally produce great amounts of syntactically distinct but equivalent alternatives. For example, a source like $a \lor b \lor c$ will generate the alternatives a, $a \land a$, and $a \land a \land a$.

The positive propositions approach, due to its semantic nature, does not include equivalent-but-distinct alternatives, yet it still produces alternatives that are altogether idle in the computation of scalar implicatures.

It is useful to consider a simpler variant of (7) above to illustrate how both theories generate unnecessary alternatives. Sentences as in (8), first discussed in the formal pragmatics literature by Spector (2007), have the implicature in (9).

(8) Either John and Mary or Bill will come to the party. $(a \land b) \lor c$
(9) Neither John nor Mary will come, or else Bill won't come. $(\neg a \land \neg b) \lor \neg c$

For $n = 3$ as in the example in (8), we get 17 equivalence classes of syntactic alternatives in Katzir's (2007) theory, versus 18 for Spector's (2007) positive propositions approach (see Mascarenhas 2014 for a detailed discussion). The alternative present in the positive propositions approach but not in the syntactic theory is $(a \land b) \lor (a \land c) \lor (b \land c)$. In words: "at least two of a, b, and c." First, I point out that this alternative is not required to derive the observed implicature in (9). In fact, we actually need only one alternative to derive the implicature, namely $(a \lor b) \land c$, which is present in both theories, and whose negation is the observed implicature in (9). Second, the missing alternative does not give rise to an observed implicature. The implicature would be that only one of a, b, and c is true, which contradicts the literal meaning of (8). Consequently, while the syntactic theory generates *far* more alternatives than the positive-propositions theory for sentences as in (8), an algorithmic-level implementation of this theory that only considered equivalence classes of substitutions would generate *fewer* alternatives than its competitor.

But finding such an implementation is far from trivial, for note that it is not enough to simply prune the set of alternatives after the fact, or to ask during the computation, for each new potential alternative, whether it is equivalent to an alternative already included in the output. Either of these strategies still requires that the algorithm generate astronomical numbers of alternatives, even if it then decides not to include them in the output.[9]

Zooming out, the theories reviewed here may be right as computational-level accounts. But plausible algorithmic-level implementations do not exist at the time of writing of this article, and designing them is far from a straightforward task.

4 Conclusion

Modern theories of scalar implicature make crucial use of sets of alternatives. In this article I showed how the three most precisely defined alternative-generating procedures in the literature (a) generate very large sets even for the case of small inputs and (b) generate sets whose size increases (at least) exponentially as a function of the input. These facts highlight a puzzling state of affairs that deserves investigation. If these theories of alternative-generation should have algorithmic-level psychological import, then one is at a loss trying to integrate them within the broader existing research on cognitive mechanisms that deal with reasoning and with alternatives. If on the other hand they are to be taken only as theories of pragmatic competence, then we must ask what mechanisms implement this description of competence; what heuristics does the mind use that allow it to consider only a manageable subset of these large collections of alternatives, while deriving the observed implicatures.

References

Chierchia, G. (2004). Scalar implicatures, polarity phenomena, and the syntax/semantics interface. In A. Belleti (Ed.), *Structures and beyond.* Oxford: Oxford University Press.

Chierchia, G., Fox, D., & Spector, B. (2012). The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In P. Portner, C.

[9] A strategy that is in principle available is for the algorithm to start from a manageable set of contextually available alternatives and ask for each of those potential alternatives whether it is generated by the general theory of alternatives. The first step in finding out whether this is a solution to the challenge I raise in this article is to determine the complexity of the problem for a non-deterministic machine. This is because plausibly the problem would now best be framed in terms of checking whether a particular alternative is in the set of alternatives, rather than generating the full set in itself. I do not have an answer to this question, which would require a complete complexity result for each of the algorithms discussed here in the non-deterministic case, rather than deterministic lower bounds as I establish.

Be that as it may, a reviewer points out that the idea of starting with an independent set of candidate alternatives runs against arguments by Fox and Katzir (2011) that context does not freely restrict the alternatives to scalar implicature. In sum, it is in principle possible that alternatives aren't generated fully and then winnowed, but rather one starts from a small set of contextual alternatives and then tests each of those to see if it can be generated. But it is an open question whether even this process is tractable, and it is unclear whether a complete theory of scalar implicature can be built from such a theory of alternatives.

Maienborn, & K. von Heusinger (Eds.), *Semantics: An international handbook of natural language meaning*. Berlin: Mouton de Gruyter.

Fox, D. (2007). Free choice disjunction and the theory of scalar implicature. In U. Sauerland & P. Stateva (Eds.), *Presupposition and implicature in compositional semantics* (pp. 71–120). Pelgrave McMillan.

Fox, D., & Katzir, R. (2011). On the characterization of alternatives. *Natural Language Semantics*, *19*, 87–107.

Gilbert, E. N. (1954). Lattice-theoretic properties of frontal switching functions. *Journal of Mathematicl Physics*, *33*, 57–67.

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

Horn, L. (1972). *On the semantic properties of the logical operators in English* (PhD thesis). UCLA.

Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge: Cambridge University Press.

Johnson-Laird, P. N. (2008). Mental models and deductive reasoning. In L. Rips & J. Adler (Eds.), *Reasoning: Studies in human inference and its foundations* (pp. 206–222). Cambridge: Cambridge University Press.

Katzir, R. (2007). Structurally-defined alternatives. *Linguistics and Philosophy*, *30*, 669–690.

Kleitman, D. (1969). On Dedekind's problem: The number of monotone boolean functions. In *Proceedings of the american mathematical society* (Vol. 21, pp. 677–682).

Kleitman, D., & Markowsky, G. (1975). On Dedekind's problem: The number of isotone Boolean functions II. *Transactions of the American Mathematical Society*, *213*, 373–392.

Mascarenhas, S. (2014). *Formal semantics and the psychology of reasoning: Building new bridges and investigating interactions* (PhD thesis). New York University.

Sauerland, U. (2004). Scalar implicatures in complex sentences. *Linguistics and Philosophy*, *27*, 367–391.

Spector, B. (2007). Scalar implicatures: Exhaustivity and Gricean reasoning. In M. Aloni, P. Dekker, & A. Butler (Eds.), *Questions in dynamic semantics*. Elsevier.

Spector, B. (2016). Comparing exhaustivity operators. *Semantics & Pragmatics*, *9*(11), 1–33.

Trinh, T., & Haida, A. (2015). Constraining the derivation of alternatives. *Natural Language Semantics*, *23*, 249–270.

Walsh, C., & Johnson-Laird, P. N. (2004). Coreference and reasoning. *Memory and Cognition*, *32*, 96–106.