Nat Lang Semantics (2015) 23:249–270
DOI 10.1007/s11050-015-9115-y

Constraining the derivation of alternatives

Tue Trinh[1] · **Andreas Haida**[2]

Published online: 13 June 2015

Abstract Inferences that result from exhaustification of a sentence S depend on the set of alternatives to S. In this paper, we present some inference patterns that are problematic for previous theories of alternatives and propose some structural constraints on the derivation of formal alternatives which derive the observations.

Keywords Alternatives · Exhaustification · Implicature · Focus · Symmetry

1 Introduction

One way to describe the strengthened meaning of a sentence S, i.e. the conjunction of its literal meaning and its implicature, is to say that S can be parsed as **exh**(A)(S), where A is a set of alternatives of S and **exh** an exhaustifying operator which maps A and S to true iff S is true and the elements of a subset of A, defined relative to S and A, are false (cf. Krifka 1995; Fox 2007a; Chierchia et al. 2012; Magri 2009, 2011; among others).[1]

[1] Following Fox (2007a,b), we assume N(S, A) in (1) to be the set of "innocently excludable alternatives in A given S," defined as follows:

(i) $N(S, A) := \cap \{A' \mid A' \text{ is a maximal subset of A such that } \{S\} \cup \{\neg S' \mid S' \varepsilon A'\} \text{ is consistent}\}$

✉ Andreas Haida
andreas.haida@hu-berlin.de

Tue Trinh
tuetrinh@uwm.edu

1 Department of Linguistics, University of Wisconsin - Milwaukee, Curtin Hall 535, P.O. Box 413, 3243 N. Downer Ave., Milwaukee, WI 53211, USA

2 Department of English and American Studies, Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany

(1) $\text{exh}(A)(S) = S \wedge \bigwedge\{\neg S' \mid S' \varepsilon N(S, A)\}$

Under this perspective, predicting the strengthened meaning of S involves predicting what is in A. For example, we predict (2a) to have the strengthened meaning in (2b) if we predict (2c).

(2) a. exh(A)(John did some of the homework)
b. John did some of the homework ∧ ¬John did all of the homework
c. A = {John did some of the homework, John did all of the homework}

As has been noted, this approach to implicature amounts to assimilating it to association with focus (cf. Krifka 1995; Fox 2007b; Fox and Katzir 2011). The definition of **exh** is essentially that of the exhaustifying operator **only**, modulo the fact that **only**(A)(S) presupposes the prejacent instead of asserting it.[2]

(3) only(A)(S) = exh(A)(S) if $\models_c$ S, undefined otherwise

For present purposes, we can ignore the difference between presupposition and assertion and regard **only** simply as the overt counterpart of **exh**. The strengthened meaning of (2) can then be identified with the literal meaning of (4).[3]

(4) John only did some of the homework

We will consider **exh**(A)(S) and **only**(A)(S) two different instances of the more general phenomenon of exhaustification of S, and use "EXH" as a cover term for **exh** and **only**. The aim of this paper is to propose a characterization of A which accounts for inference patterns of EXH(A)(S) that pose a challenge for other proposals. An example of such patterns is the paradigm below.

(5) Bill went for a run and didn't smoke. John (only) went for a run.
Inference: ¬[John went for a run and didn't smoke]

(6) Bill passed some of the tests and failed some. John (only) passed some of the tests.
*Inference: ¬[John passed some of the tests and failed some]

Footnote 1 continued
Informally, then, N(S, A) is the intersection of all maximal subsets of A whose elements can be consistently negated in conjunction with S. We will leave it to the reader to do the computations to see how this definition works with respect to the examples discussed in the text.

The operator **exh** can be seen as a notational device expressing pragmatic reasoning on a sentence (cf. Sauerland 2004; Spector 2006), or alternatively as a grammatical device which is syntactically represented and interpreted by compositional semantics. We believe the second option is correct, in light of arguments made elsewhere concerning free choice disjunction (Fox 2007a), Hurford's Constraint (Chierchia et al. 2012), modularity (Magri 2009, 2011), and intermediate implicatures (Sauerland 2012). The precise status of **exh**, however, will not be essential for what follows.

[2] The 'prejacent' is the sentential argument of **exh**.

[3] Embedding (4) under **exh** will have no semantic effect, so we may assume that sentences are by default construed with **exh** (cf. Fox 2007a). This raises the issue of the distributional differences between **only** and **exh**, which we will also have to ignore (cf. Crnič 2012 and references therein).

While (5) can be understood to mean that what is true of Bill is not true of John, (6) cannot.[4] Specifically, (5) can imply that it is not the case that John went for a run and didn't smoke (i.e. that John smoked) but (6) cannot imply that it is not the case that John passed some of the tests and failed some (i.e. that John passed all of the tests). As far as we know, this contrast has not been noted in the literature, and is not predicted by any theory of alternatives on the market.

The paper is structured as follows. The rest of Sect. 1 sets up the background for the subsequent discussion. Section 2 introduces the central empirical puzzle which motivates our account. Section 3 presents a resolution of this puzzle. Section 4 extends the hypothesis proposed in Sect. 3 to solve other problems. Section 5 is the summary.

1.1 Relevance

Rooth (1992) notes that EXH(A)(S) licenses different inferences depending on which constituent of S is focused (i.e. F-marked). He consequently proposes that A be regarded as F(S) ∩ C, where F(S) is the set of formally defined alternatives of S and C the set of pragmatically salient sentences. Rooth defines F(S) as follows.

(7) *Formal alternatives* (Rooth 1992)
F(S) = {S′ | S′ is derivable from S by replacement of F-marked constituents with expressions of the same semantic type}

Fox and Katzir (2011; henceforth F&K) argue that C should be identified with the set of relevant sentences, where relevance is to be closed under negation and conjunction.[5]

(8) *Closure conditions on C* (Fox and Katzir 2011)
(i) If p ε C, then ¬p ε C.
(ii) If p ε C and q ε C, then p ∧ q ε C.

1.2 Symmetry

The conception of A as F(S) restricted by relevance leads to a problem with Rooth's definition of F(S): the so-called "symmetry problem."[6] As an illustration, consider (9), with the parse indicated.

[4] The sequence in (6) is admittedly odd. Our intuition is that the reason for its oddness is precisely that it cannot mean that what is true of Bill is not true of John. To the extent that our intuition is correct, the problem at hand can be framed as one of explaining the contrast in acceptability between (5) and (6).

[5] These two closure conditions follow from the assumption that to be relevant is to distinguish exclusively between answers to the "question under discussion." More explicitly, let Q be the question under discussion and Ans(Q)(w) be the set of answers to Q that are true in w. For a proposition p to be relevant, it must hold that p makes no distinction between w and w′ if there is no distinction between Ans(Q)(w) and Ans(Q)(w′), i.e. it must hold that p(w) = p(w′) if Ans(Q)(w) = Ans(Q)(w′) (cf. Groenendijk and Stokhof 1984; Lewis 1988; von Fintel and Heim 1997).

[6] The symmetry problem was first formulated in the context of the discussion of scalar implicatures (Kroch 1972; von Fintel and Heim 1997). However, it generalizes to association with focus, as shown in Fox and Katzir (2011).

(9) John (only) has $three_F$ chairs
EXH(A)(S), where S = John has three chairs

By hypothesis, both S′ = **John has four chairs** and S″ = **John has exactly three chairs** are in F(S), as **four** and **exactly three** are both of the same semantic type as **three**. And given that S is relevant, S′ is in C if and only if S″ is in C.[7] Thus, A = F(S) ∩ C contains either both S′ and S″ or neither of these alternatives, which means that (9) cannot license ¬S′ as an inference, because this inference requires that A contain S′ but not S″. In actuality, however, this is exactly the inference we draw from (9): the sentence clearly implies that John has three chairs and not four.

Some other examples of the same sort are given below. In each case, the two alternatives S′ and S″ are such that either both are predicted to be in A or none is. And in each case, ¬S′ is the observed inference.

(10) John (only) [did some of the homework]$_F$
S′ = John did all of the homework, S″ = John did some but not all of the homework

(11) John (only) [talked to Mary or Sue]$_F$
S′ = John talked to Mary and Sue, S″ = John talked to Mary or Sue but not both

Such alternatives as S′ and S″ in the examples above are called ‘symmetric alternatives’. More explicitly, we say that S′ and S″ are symmetric alternatives of S if (i) S′ ∧ S″ is a contradiction and (ii) S′ is relevant iff S″ is relevant. The “symmetry problem” is this: A is predicted to contain either both S′ and S″ or none of these, while the facts require that A contain S′ but not S″. This problem is solved by “breaking symmetry,” i.e. by redefining A in such a way that it can contain S′ but not S″.

1.3 Complexity

Building on Katzir (2007, 2008), F&K advance a theory of alternatives which starts from the intuition that what distinguishes between the symmetric alternatives above is their structural complexity: S″ is more complex than S while S′ is not. To exclude S″ from A, then, F&K propose to reduce the set of formal alternatives from the Roothian set, which we will henceforth denote with ‘$F_{Rooth}(S)$’, to a proper subset of this set which contains only structures that are (considered by the discourse participants to be) “no more complex” than the prejacent.

(12) *Formal alternatives* (F&K)
$F(S) = F_{Rooth}(S) \cap \{S' \mid S' \preceq_c S\}$

The relation ‘$x \preceq_c y$’ holds between linguistic expressions in general and is to be understood as ‘x is no more complex than y in discourse context c.’ Here is the definition:

[7] Thus, if S′ ε C, then ¬S′ ε C, which means S ∧ ¬S′ ε C, which means S″ ε C, as S″ ≡ S ∧ ¬S′. The same argument can be made with S″ in place of S′.

(13) *Complexity metric* (F&K)

a. $E' \preceq_c E$ if $E' = T_n(\ldots T_1(E)\ldots)$, where each $T_i(x)$ is the result of replacing a constituent of x with an element of SS(E,c), the substitution source of E in c

b. SS(E,c) = {x | x is a lexical item} ∪ {x | x is a constituent uttered in c}

Basically, E′ is no more complex than E if E′ can be derived from E by a series of substitution transformations, each of which applies to an input x and replaces one constituent of x with a lexical item or a constituent uttered in the context.[8] The reader can verify for herself that F&K's complexity metric correctly predicts S′, but not S″, to be a formal alternative of S in (9), (10), and (11). This metric also accounts for facts beyond those we have considered, such as the contrast between (14) and (15).

(14) Yesterday it was (only) warm. Today it is warm and sunny with gusts of wind.
Inference: ¬Yesterday it was hot
Inference: ¬Yesterday it was warm and sunny with gusts of wind

(15) Yesterday it was (only) warm.
Inference: ¬Yesterday it was hot
*Inference: ¬Yesterday it was warm and sunny with gusts of wind

While (14) implies that it was not hot and also that it was not sunny with gusts of wind, (15) only licenses the inference that it was not hot. This is because in the case of (14), alternatives can be generated by replacing **warm** with **hot**, taken from the lexicon, and with **warm and sunny with gusts of wind**, taken from the discourse context, while in the case of (15) the second replacement is not an available option. Thus, what is uttered in the context affects the domain of EXH and hence the inference of the exhaustified sentence. Several additional supporting data are given in Katzir (2007) and Fox and Katzir (2011). Here is one more example, taken from Fox and Katzir (2011) with slight modification.

(16) Detective A concluded that the robbers stole the book and not the jewelry. Detective B (only) concluded that they stole the book.
Inference: ¬Detective B concluded that the robbers stole the book and not the jewelry
Inference: ¬Detective B concluded that the robbers stole the book and the jewelry

The first inference comes about by way of the alternative in (17a), which is generated by replacing **the book** in the prejacent with **the book and not the jewelry**. The second inference comes about by way of the alternative in (17b), which is derived from the first alternative, (17a), by replacing **not the jewelry** in (17a) with **the jewelry**.

(17) a. Detective B concluded that they stole the book and not the jewelry
b. Detective B concluded that they stole the book and the jewelry

[8] It follows from this definition that $E \preceq_c E$, i.e. that an expression is no more complex than itself.

Thus, F&K assume that alternatives can be derived by "successive replacement" in the sense that a constituent α can be replaced by another constituent β which can itself be altered by further replacements.[9]

A prominent advantage of F&K's theory is that it derives the right alternatives for disjunction in the same manner as other alternatives are derived, i.e. by substitution. Consider (18).

(18) John is (only) required to read the book or do the homework.
Inference: ¬John is required to read the book
Inference: ¬John is required to do the homework

The alternatives necessary for these inferences are derived by replacing **read the book or do the homework** in the prejacent with **read the book** in one case and with **do the homework** in the other.

2 Identifying the problem: unexpected symmetry breaking

In the theory of alternatives proposed by F&K, symmetry is broken exclusively in F(S): the problematic symmetric alternative is excluded from A = F(S) ∩ C by conditions imposed on F(S), not on C. This aspect of F&K's proposal reflects a conviction underlying most semantic analyses which appeal to alternatives.

(19) *Standard view on symmetry*
Symmetry can only be broken formally.

Let us, at this point, introduce the main empirical puzzle that this paper sets out to resolve. This puzzle can be divided into two parts. The first concerns data such as (20a–c), which seem to contradict the standard view on symmetry.

(20) a. Bill went for a run and didn't smoke. John (only) went for a run.
Inference: ¬[John went for a run and didn't smoke]
b. Bill works hard and doesn't watch TV. John (only) works hard.
Inference: ¬[John works hard and doesn't watch TV]
c. Bill is tall and not bald. John is (only) tall.
Inference: ¬[John is tall and not bald]

All of these sentences license the inference that what is true of Bill is not true of John. For example, (20a) implies that it is not the case that John went for a run and didn't smoke (i.e. that he smoked). Given what has been said, however, this inference is licensed only if symmetry can be broken outside of F(S). Thus, (20a) has the form EXH(A)(S) where S = **John went for a run**. The formal alternatives of S in this case include both S′ = **John went for a run and didn't smoke** and S″ = **John went**

[9] This is what we take F&K to mean by "successive replacement," as we see no other way to derive (17b) (see Katzir 2007, p. 679; Fox and Katzir 2011, pp. 97, 103). Of course, this does not mean the attested inferences require such an understanding of "successive replacement," as we will argue below.

for a run and smoked, which are symmetric alternatives.[10] The attested inference, however, is that John smoked, which is only possible if A contains S′ but not S″. Thus, we have a case where F(S) contains symmetric alternatives while A does not, i.e. a case where, on F&K's definition of F(S), symmetry is broken outside of F(S).

The same holds for (20b) and (20c). In (20b), the set of formal alternatives to S = **John works hard**, which is the prejacent of EXH, contains both **John works hard and doesn't watch TV** and its symmetric counterpart **John works hard and watches TV**. Given that EXH(A)(S) implies that it is not the case that John works hard and doesn't watch TV (i.e. that he watches TV), A must be such a restriction of F(S) as to contain the first but not the second alternative. In other words, A must be a symmetry breaking restriction of F(S). Similarly, the set of formal alternatives to the prejacent of EXH in (19c), which is S = **John is tall**, contains both **John is tall and not bald** and the symmetric counterpart **John is tall and bald**, but the inference attested for EXH(A)(S), namely that it is not the case that John is tall and not bald (i.e. that he is bald), is only possible if A contains the second but not the first alternative, i.e. if symmetry is broken outside of F(S).

The observations in (20), then, suggest that the standard view on symmetry is false. It turns out, however, that there is another set of data which has a very similar structure to that in (20) but which leads to the exact opposite conclusion, namely that the standard view on symmetry is correct. This is the second part of the puzzle.

(21) a. Bill ate exactly three cookies. John (only) ate three cookies.
*Inference: ¬John ate exactly three cookies
b. Bill fathered children and no twins. John (only) fathered children.
*Inference: ¬[John fathered children and no twins]
c. Bill passed some of the fitness tests and failed some. John (only) passed some of the fitness tests.
*Inference: ¬[John passed some of the fitness tests and failed some]

In (21a), the prejacent of EXH is S = **John ate three cookies**, and the set of formal alternatives of S, F(S), includes both **John ate four cookies** and **John ate exactly three cookies**. The fact that EXH(A)(S) cannot imply that it is not the case that John ate exactly three cookies (i.e. that John ate more than three cookies) means that F(S) cannot be restricted to a set A which contains **John ate exactly three cookies** but not its symmetric counterpart. In (21b), the prejacent of EXH is S = **John fathered children**, and F(S) contains both **John fathered twins** and its symmetric counterpart **John fathered children and no twins**. The fact that EXH(A)(S) cannot imply that John fathered twins means that A cannot eliminate one alternative from F(S) but not the other. Similarly, the prejacent of EXH in (21c) is S = **John passed some of the fitness tests**, and F(S) contains both **John passed all of the fitness tests** and its symmetric counterpart **John passed some of the fitness tests and failed some**. The fact that EXH(A)(S) cannot imply that it is not the case that John passed some of the fitness

[10] The first alternative S′ is generated by replacing **went for a run** in S with **went for a run and didn't smoke**, taken from the discourse context. The second alternative S″ is generated by replacing **didn't smoke** in S′ with **smoked**, also taken from the discourse context. As S is relevant and $S' \equiv S \wedge \neg S''$ and $S'' \equiv S \wedge \neg S'$, S′ is relevant iff S″ is.

tests and failed some (i.e. that he passed all of the fitness tests) means that A cannot break symmetry by restricting F(S) to a set which contains the second alternative to the exclusion of the first.

Here, then, is the problem facing us. We have a definition of F(S) and a definition of A which together imply that if F(S) contains two symmetric alternatives S' and S'', A must contain both S' and S'' and consequently EXH(A)(S) cannot imply $\neg S'$. In this section, we have examined cases where F(S), by definition, contains two symmetric alternatives S' and S''. In some of these cases, EXH(A)(S) can imply $\neg S'$ while in some others it cannot. The task, therefore, is to establish a distinction between these two sets of cases which matches the attested inference patterns.

3 Solving the problem

3.1 First attempt: relying on pragmatic scales alone

One strategy that readily comes to mind is to appeal to the notion of a "pragmatic scale" (cf. Klinedinst 2004). To illustrate, let us consider (22a) and (22b), as representatives of cases where symmetry can and cannot be broken outside of F(S), respectively.

(22) a. Bill went for a run and didn't smoke. John (only) went for a run.
Inference: ¬[John went for a run and didn't smoke]
b. Bill passed some of the fitness tests and failed some. John (only) passed some of the fitness tests.
*Inference: ¬[John passed some of the fitness tests and failed some]

Given common knowledge, it seems much easier to construct an evaluative scale on which running and smoking is ranked lower than running and not smoking, for example one which measures the degree of healthiness, than it is to construct a scale on which passing all of the tests ranks lower than passing some but not all of the tests. Suppose, then, that symmetry can in principle be broken in C, or equivalently that EXH(A)(S) can in principle license $\neg S'$ as an inference even though F(S) contains both S' and its symmetric counterpart.[11] In addition, suppose we say that in such discourse contexts as those in (22), EXH(A)(S) can only license $\neg S'$ as an inference if $S \wedge \neg S'$ is ranked lower than $S \wedge S'$ on some pragmatic scale. It will then seem that we can account for the contrast between (22a) and (22b). Thus, the prejacent S of EXH in (22a) is **John went for a run** and the relevant alternative S' is **John went for a run and didn't smoke**. The reason EXH(A)(S) can imply $\neg S'$ in this case would be that there is a pragmatic scale measuring the degree of healthiness on which $S \wedge \neg S'$, which says that John went for a run and smoked, is ranked lower than $S \wedge S'$, which says that John went for a run and didn't smoke. For (22b), where S = **John passed some of**

[11] Let it already be said at this point that we will end up abandoning this attempt of having symmetry be broken in C and will come back to assuming F&K's claim that symmetry can only be broken in F. Note that if we were to pursue the hypothesis that symmetry can in principle be broken in C, we would have to address F&K's empirical arguments to the contrary, most probably by showing that the relevant pragmatic scales are missing in the cases that F&K discuss.

the fitness tests and S′ = **John passed some of the fitness tests and failed some**, EXH(A)(S) cannot license ¬S′ as an inference because there is no pragmatic scale on which S ∧ ¬S′, which says that John passed all of the tests, is ranked lower than S ∧ S′, which says that John passed some but not all of the tests.

An account along this line is also supported by the following contrast.[12]

(23) a. Bill went for a run and didn't smoke. John (only) went for a run.
Inference: ¬[John went for a run and didn't smoke]
b. Bill went for a run but didn't lift weight. John (only) went for a run.
*Inference: ¬[John went for a run and didn't lift weight]

We can see that changing (23a) into the minimally different (23b) results in the absence of the inference that what is true of Bill is not true of John. As the change consists solely in replacing the "positive" property of not smoking with a "negative" property of not lifting weight (and the concomitant replacement of **and** with **but**), we have evidence that the availability of an appropriate pragmatic scale, however the details of the account will be worked out, plays a role in licensing the type of inferences observed in such sentences as (23a).

It turns out, however, that the pragmatic scale approach, while it promises to account for the contrast in (23), seems hopeless as an explanation for the difference between the two sets of cases presented in the last section. Specifically, it turns out that even if we make available, and salient, a pragmatic scale on which passing all of the fitness tests is ranked lower than passing some and failing some, (22b) still cannot license the inference that it is not the case that John passed some and failed some of the tests. Thus, imagine a situation where everyone has to take military fitness tests and those who pass all of these tests must join the army and go to the frontline to be slaughtered by the enemy while those who fail some of the tests can stay home and go to college. In this draft dodging context, a scale which measures the degree of luck will be one on which passing all of the military fitness tests is ranked lower than passing some and failing some of these tests. But it seems that even this context cannot support the relevant inference for (22b).

(24) Bill has once again been dealt a better hand than John. He passed some of the military fitness tests and failed some, while John (only) passed some of the tests.

We believe that (24) cannot license the inference that it is not the case that John passed some of the military fitness tests and failed some, i.e. that he passed all of the tests.

Let us now look at the other examples in (21) to convince ourselves that the pragmatic scale account does not hold water. For (21a), imagine a "gluttony rehab" context where eating fewer cookies means making more progress. In this context, there is a pragmatic scale on which eating more than three cookies is ranked lower than eating exactly three cookies, namely one which measures how far a person is from being a glutton. Again, it appears that even this context does not help in getting the following

[12] Apparently, the absence of the inference that John is different from Bill with respect to running and lifting weights makes (23b) sound a bit odd (see footnote 4).

sentence to imply that John ate three but not exactly three cookies, i.e. that he ate more than three cookies.

(25) Bill is making more progress at the gluttony rehab center than John. Today he ate exactly three cookies, while John (only) ate three cookies.

A similar argument can be made for (21b). Imagine a fictional "genetic lab" context where scientists try to engineer the most "representative" male *Homo sapiens*, i.e. one which can father children but which does not have the biological defect of always fathering (identical) twins.[13] In this context, a scale measuring how similar an engineered sample is to the majority of naturally born humans would rank a specimen which fathered twins lower than one which fathered children but no twins. However, it seems that even this context fails to enable the following sentence to imply that John fathered twins.

(26) Bill turned out to be a better specimen than John. He fathered children and no twins, while John (only) fathered children.

We conclude that a solution to the problem identified in Sect. 2 in terms of pragmatic scales is not tenable.[14] Let us try another route.

3.2 Second and final attempt: syntactic atomicity

Let us start with a reformulation of F&K's theory which will facilitate the presentation of the problem it faces. Given the assumption that the prejacent is always relevant and that the set C of relevant sentences is closed under negation and conjunction, F&K's hypothesis that $A = F(S) \cap C$ amounts to the hypothesis that A is a set which satisfies the following three conditions.[15]

(27) *Conditions on A* (F&K)
a. $A \subseteq F(S)$
b. S is in A.
c. There is no S' in $F(S)\backslash A$ such that S' is in the Boolean closure of A.

Informally speaking, A is the set of all and only those sentences in F(S) that are considered relevant, which means that if a sentence φ is in A, then every sentence in

[13] Let us, for argument's sake, assume that the male parent can be responsible for monozygotic twinning.

[14] More precisely, the discussion in this section shows that the availability of an evaluative scale on which $S \wedge \neg S'$ is ranked lower than $S \wedge S'$ is not sufficient for the inference $\neg S'$ to arise. Still, the contrast in (23) suggests that the availability of such a scale is a necessary condition: the inference does not arise in the case of (23b) because the conjunction **but** in the context sentence presupposes or implicates a scale on which running and not lifting weights ($S \wedge S'$) is ranked lower than running and lifting weights ($S \wedge \neg S'$).

[15] Note that F&K actually end up proposing a revision of the last condition: there is no S' in $F(S)\backslash A$ such that $EXH(A)(S')$ is in the Boolean closure of A. The revision is motivated by a problem with disjunction, specifically the problem of ensuring that if a disjunction is relevant then both disjuncts are (cf. Fox and Katzir 2011 for details). It turns out, however, that this problem can be solved by other means (cf. Spector 2010). The difference between the original and the revised condition on A is immaterial for our discussion.

F(S) that becomes relevant by virtue of φ and other sentences in A being relevant must also be in A.

We now turn to a more explicit discussion of the challenge confronted by F&K's theory, i.e. the cases of symmetry being broken outside of F(S). Let us again take (20a), repeated in (28) below, as a representative of these cases.

(28) Bill went for a run and didn't smoke. John (only) went for a run.
Inference: ¬[John went for a run and didn't smoke]

It follows from F&K's definition of formal alternatives that F(S), with S = **John went for a run**, is the set in (29).[16]

(29) $\{\mathbf{run}, \mathbf{smoke}, \neg\mathbf{run}, \neg\mathbf{smoke}, \mathbf{run} \wedge \neg\mathbf{smoke}, \mathbf{run} \wedge \mathbf{smoke}\}$

The attested inference of **John went for a run**, namely that John smoked, requires that the domain A of EXH, assuming it is a subset of F(S) which contains S, be any of the three sets in (30). Thus, EXH(30a)(**run**) = EXH(30b)(**run**) = EXH(30c)(**run**) = $\mathbf{run} \wedge \mathbf{smoke}$.

(30) a. $\{\mathbf{run}, \mathbf{run} \wedge \neg\mathbf{smoke}, \neg\mathbf{run}, \neg\mathbf{smoke}\}$
b. $\{\mathbf{run}, \mathbf{run} \wedge \neg\mathbf{smoke}, \neg\mathbf{run}\}$
c. $\{\mathbf{run}, \mathbf{run} \wedge \neg\mathbf{smoke}\}$

The problem, however, is that none of the sets in (30) satisfies the last condition on A in (27), given that F(S) = (29).[17] To see this, let X be any of these sets. It will then hold that (29)\X contains $\mathbf{run} \wedge \mathbf{smoke}$ which, being equivalent to $\mathbf{run} \wedge \neg(\mathbf{run} \wedge \neg\mathbf{smoke})$, is in the Boolean closure of X.

[16] To be explicit, we assume the following constituent structures for the two sentences in question. The reader can verify for herself that (29) is the set of formal alternatives of (i-b), according to F&K's definition.

(i)

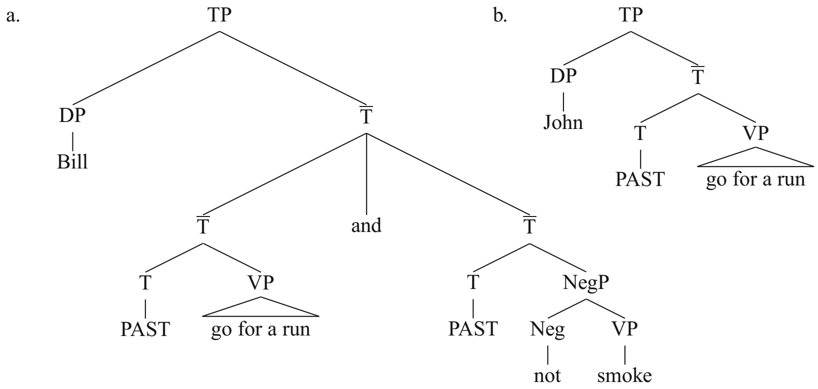

Importantly, F&K predict that neither $\neg(\mathbf{run} \wedge \mathbf{smoke})$ nor $\neg(\mathbf{run} \wedge \neg\mathbf{smoke})$ is a formal alternative of (i-b), as none of these can be generated by successively replacing constituents of (i-b) with constituents of (i-a) or (i-b).

[17] However, note in advance that the last condition in (27) can be satisfied if **smoke** is excluded from A: **smoke** is not in the Boolean closure of the sets in (30b) and (30c).

3.2.1 The Atomicity constraint

The solution we propose to the problem created by the example above for F&K's theory consists in constraining the derivation of formal alternatives in such a way that F(S) does *not* contain **run** ∧ **smoke**. We will present our proposal in two steps: (i) presenting the constraint on F(S) which excludes **run** ∧ **smoke** from this set; (ii) showing how this constraint works with respect to the example above and other examples that we have discussed. Let us begin with the first step.

What we want is to exclude **run** ∧ **smoke** from F(S), with S = **John went for a run**.[18] One way to achieve this goal is to impose on F(S) the condition that all of its members be derivable from S in at most one step. This would exclude the derivation of **run** ∧ **smoke** from the prejacent, which necessarily involves two steps. This derivation is shown in (31). By assumption, the relevant substitution source includes the $\overline{\mathrm{T}}$ constituent α = **went for a run and didn't smoke** and the VP constituent β = **smoke**, both of which have been uttered in the context.[19]

(31) Derivation of **John went for a run and smoked**:

0.	**John went for a run**	the prejacent
1.	**John went for a run and didn't smoke**	$\overline{\mathrm{T}}$ / α
2.	**John went for a run and smoked**	NegP / β

Imposing on F(S) such a condition, however, will cost us an elegant solution to the notorious problem of multiple disjunctions, i.e. sentences of the form (A **or** (B **or** C)), which by default license the inference that exactly one disjunct is true. This inference follows from (the definition of EXH and) the assumption that (A **and** B), (A **and** C), and (B **and** C) are all formal alternatives of (A **or** (B **or** C)). But none of these formal alternatives can be generated in one step: each of them requires a step which replaces **or** with **and** and another step which reduces the ternary coordination to a binary one, in either order.[20] For example, (A **and** B) is derived from (A **or** (B **or** C)) by first replacing (B **or** C) with B and then replacing **or** in the result with **and**.

We believe this is good reason for not imposing on F(S) the condition that its elements be derived in at most one step. The question, then, is what makes **run** ∧ **smoke** special other than the plurality of steps its derivation involves. Our answer will require the introduction of a new concept, that of a "syntactically atomic expression," which in turn requires a brief discussion of what we have been calling the "substitution source."

The substitution source can be naturally thought of as a set of expressions provided by the discourse context for the contruction of formal alternatives. Thus, the syntactic derivation of alternatives "begins" with the substitution source in much the same way as

[18] Note that the computations we are looking at for S = **John went for a run** are all performed in a context where the sentence **Bill went for a run and didnt smoke** has been uttered.

[19] See footnote 16. We write "A/B" to denote the result of replacing A in the previous line with B.

[20] More precisely, the generation of these alternatives requires a step which reduces a structure containing two binary coordinations (i.e. (A or (B or C))) to a structure containing only one binary coordination (e.g. (A or B)). Note that a genuine ternary structure such as (A or B or C) cannot be reduced to (A or B) using F&K's procedures. We thank Roni Katzir for pointing out this subtle difference.

the syntactic derivation of sentences begins with the numeration (cf. Chomsky 1995). One manifest difference between these two contructs, of course, is that whereas the numeration contains only lexical items, the substitution source contains both lexical items and complex phrases. Interestingly, the core of our solution to the problem at hand will be the hypothesis that this difference is only phonological, not syntactic. Specifically, we propose the following.

(32) *Atomicity*
Expressions in the substitution source are syntactically atomic.

Basically, the Atomicity constraint states that expressions in the substitution source cannot have their proper parts replaced by other expressions in the derivation. One way to implement this idea is to say that every expression in the substitution source is formally marked with a feature, AT, which makes it syntactically atomic, i.e. makes its internal structure invisible and thus inaccessible to syntactic rules. The numeration and the substitution source are now similar in the sense that both contain only syntactically atomic expressions.

Returning to our example, we now say that the substitution source of **John went for a run** in the relevant discourse context includes $\alpha = [_{AT}$ **went for a run and didn't smoke**] and $\beta = [_{AT}$ **smoke**], with each expression now explicitly marked as being syntactically atomic. Let us look at the derivation of **run** $\wedge$ **smoke** again.

(33) Derivation of **John went for a run and smoked**:

0.	**John went for a run**	the prejacent
1.	**John** $[_{AT}$ **went for a run and didn't smoke]**	$\overline{\text{T}}$ / α
2.	**John** $[_{AT}$ **went for a run and** $[_{AT}$ **smoked]]**	*NegP / β

As we can see, the second step in (33) violates Atomicity: it replaces a proper part of a syntactically atomic expression with another expression. This means that **run** $\wedge$ **smoke** cannot be derived from the prejacent, i.e. that it is not in F(S)—which is the result we want.[21]

3.2.2 *Deriving the symmetry breaking cases*

Let us first consider the "symmetry breaking" cases, i.e the examples in (20), starting with the one we have been discussing in detail above.

(34) Bill went for a run and didn't smoke. John (only) went for a run.
Inference: ¬[John went for a run and didn't smoke]

With Atomicity imposed on F&K's concept of formal alternatives, F(**John went for a run**) is now the set in (35), as neither **¬run** nor **run** $\wedge$ **smoke** can be derived without violating the Atomicity constraint.

[21] The reader is invited to verify for herself that Atomicity does not create a problem with multiple disjunctions, i.e. that we can derive (A **and** B), (A **and** C), and (B **and** C) from (A **or** (B **or** C)) without violating Atomicity.

(35) F(S) = {**run**, **smoke**, ¬**smoke**, **run** ∧ ¬**smoke**}

The set A, which is the domain of EXH, can now be (36), which satisfies all of the conditions in (27): (i) (36) ⊆ F(S), (ii) **run** ε (36), and (iii) no member of F(S)\(36) = {**smoke**, ¬**smoke**} is in the Boolean closure of (36).

(36) A = {**run**, **run** ∧ ¬**smoke**}

Consequently, EXH(A)(S) = **run** ∧ ¬(**run** ∧ ¬**smoke**), which is the empirically correct result.

Let us now turn to the other symmetry breaking examples in (20) to convince ourselves that the theory works for these cases as well.

(37) Bill works hard and doesn't watch TV. John (only) works hard.
Inference: ¬[John works hard and doesn't watch TV]

It follows from Atomicity that the set of formal alternatives to S = **John works hard** in this case is (38a). The domain of EXH can then be (38b), and exhaustification gives us (38c), which is the empirically correct result.

(38) a. F(S) = {**work hard**, **watch TV**, ¬**watch TV**, **work hard** ∧ ¬**watch TV**}
b. A = {**work hard**, **work hard** ∧ ¬**watch TV**}.
c. EXH(A)(S) = **work hard** ∧ ¬(**work hard** ∧ ¬**watch TV**).

The last example in (20), repeated in (39), works similarly, see (40).

(39) Bill is tall and not bald. John is (only) tall.
Inference: ¬[John is tall and not bald]

(40) a. F(S) = {**tall**, **bald**, ¬**bald**, **tall** ∧ ¬**bald**}
b. A = {**tall**, **tall** ∧ ¬**bald**}
c. EXH(A)(S) = **tall** ∧ ¬(**tall** ∧ ¬**bald**)

3.2.3 The symmetry preserving cases

It remains to show that Atomicity does not allow us to break symmetry in the "symmetry preserving" cases, i.e. in the examples in (21). Let us examine these, beginning with (21a), repeated in (41).

(41) Bill ate exactly three cookies. John (only) ate three cookies.
*Inference: ¬[John ate exactly three cookies]

The definition of F(S) by F&K, constrained by Atomicity, implies that (42) is the set of formal alternatives to S = **John ate three cookies** in this context. Note that application of Atomicity in this case is vacuous in the sense that the non-derivable alternative **exactly four** can be excluded from A without violating the closure condition in (27c).

(42) F(S) = {**three**, **exactly three**, **four**}

To get the non-attested inference, we need the domain of EXH to be (43).

(43) {**three**, **exactly three**}

However, (43) does not qualify as A: F(S)\(43) contains **four**, which is in the Boolean closure of (43), as **four** $\equiv$ **three** $\wedge$ $\neg$**exactly three**. Consequently, we predict that (41) cannot license the inference that it is not the case that John ate exactly three cookies, which is the correct prediction.[22]

The next example, (21b), is repeated in (44).

(44) Bill fathered children and no twins. John (only) fathered children.
*Inference: $\neg$[John fathered some children and no twins]

It follows from Atomicity imposed on F(S) that the set of formal alternatives to S = **John fathered some children** is (45).[23]

(45) F(S) = {**children**, **twins**, **children** $\wedge$ **no twins**, **no twins**}

To get the non-attested inference, we need the domain of EXH to be (46).[24]

(46) {**children**, **children** $\wedge$ **no twins**}

Again, (46) does not qualify as A, since F(S)\(46) contains **twins** which, being equivalent to **children** $\wedge$ $\neg$**(children** $\wedge$ **no twins)**, is in the Boolean closure of (46). Thus, we predict that (44) cannot license the inference that it is not the case that John fathered children and no twins, i.e. that he fathered twins. This is the correct prediction.

The last example in (21) is repeated in (47).

(47) Bill passed some of the military fitness tests and failed some. John (only) passed some of the tests.
*Inference: $\neg$[John passed some of the military fitness tests and failed some]

[22] Note that (41) can, with some effort on the part of the hearer, be understood as implying that John ate exactly three cookies, in which case the sequence would sound odd by virtue of not being able to establish a difference between Bill and John (see footnote 4). One interpretation of this fact is that the effort in question is one of ignoring the contextually salient alternative (**exactly three**), keeping only the lexical alternative (**four**) in F(S). It seems that in general, such an effort can be made, i.e. that inferences based on contextually salient alternatives are optional (see footnotes 16 and 23 in Fox and Katzir 2011). Another example illustrating the same point, brought to our attention by Jacopo Romoli (p.c.), is (i).

(i) Last year, some of my students passed the test. This year, in the same way, not all of them passed it.

Romoli points out, and we agree, that (i) can have the implicature that some of my students passed the test this year. Again, this implicature is predicted to arise if F(**not all**) contains **not some**, generated by lexical replacement, but does not contain the contextually salient alternative **some**. (Adding **in the same way** makes the sequence better, apparently because it then becomes clear that the two sentences are not meant to express a contrast between last year and this year.)

In this paper, we abstract away from the issue of optionality of contextually salient alternatives and discuss the examples under the assumption that the relevant hearer does *not* make an effort to ignore such alternatives.

[23] Without Atomicity, the set of formal alternatives to S would be (45) $\cup$ {**no children**}.

[24] The non-attested inference can also be obtained with another set being the domain of EXH, namely {**children**, **children** $\wedge$ **no twins**, **no twins**}. However, this set has the same problem as (46), as the reader can easily verify.

The set of formal alternatives to **John passed some of the tests**, under the assumption that F(S) is constrained by Atomicity, is (48).

(48) F(S) = {**pass some**, **pass all**, **fail some**, **fail all**, **pass some** ∧ **fail some**}

To get the non-attested inference, we need the domain of EXH to be (49).[25]

(49) {**pass some**, **fail some**, **fail all**, **pass some** ∧ **fail some**}

However, (49) does not qualify as A, as F(S)\(49) contains **pass all**, which, being equivalent to both **pass some** ∧ ¬**fail some** and **pass some** ∧ ¬(**pass some** ∧ **fail some**), is in the Boolean closure of (49). Consequently, we make the correct prediction that (47) cannot imply that it is not the case that John passed some of the tests and failed some, i.e. that he passed all of the tests.

3.2.4 An apparent problem

An example which might appear problematic for the Atomicity account is (16), taken from F&K and repeated in (50).

(50) Detective A concluded that the robbers stole the book and not the jewelry. Detective B (only) concluded that they stole the book.
Inference: ¬Detective B concluded that the robbers stole the book and not the jewelry
Inference: ¬Detective B concluded that the robbers stole the book and the jewelry

We have presented, in the same way we think F&K did, the second inference of (50) as a consequence of (51) being a formal alternative of the prejacent **Detective B concluded that the robbers stole the book**.

(51) Detective B concluded that the robbers stole the book and the jewelry

Given Atomicity, it turns out that (51) cannot be derived from the prejacent in the context provided by (50), as the derivation involves replacing a proper part of a syntactically atomic expression with another expression, as shown in (52). The substitution source in this case includes α = [$_{AT}$ **the book and not the jewelry**] and β = [$_{AT}$ **the jewelry**].

(52) Derivation of (51):

0.	**... stole the book**	the prejacent
1.	**... stole [$_{AT}$ the book and not the jewelry]**	**the book** / α
2.	**... stole [$_{AT}$ the book and [$_{AT}$ the jewelry]]**	***not the jewelry** / β

[25] There are two other sets which will get us the non-attested inference, namely {**pass some**, **fail some**} and {**pass some**, **pass some** ∧ **fail some**}. However, both fail to qualify for A, for the same reason that (49) does.

However, there is another way for the inference in question to come about: by EXH negating the alternative in (53), which *can* be derived from the prejacent, namely by replacing the direct object **the book** in (51a) with $\beta = [_{\text{AT}}$ **the jewelry**].

(53) Detective B concluded that John stole the jewelry

Under the standard assumption that propositional attitude verbs such as **conclude** denote (restricted) universal quantifiers over worlds and thus distribute over conjunction (cf. Hintikka 1969), if Detective B did not conclude that John stole the jewelry, then he did not conclude that John stole the book and the jewelry.[26]

4 Extending the proposal

4.1 Indirect implicature

In this section we will consider how the proposal we just made can be extended in some natural fashion to derive some observations that have posed problems, specifically overgeneration problems, for Neo-Gricean theories of scalar implicature. First, we should note that the system, as it is, can already account for a puzzle, presented in Romoli (2012a), which relates to the well-known observation that **[not[all]]**, by default, implicates **[some]** (Atlas and Levinson 1981; Horn 1989; Sauerland 2004), as exemplified in (54).[27]

(54) The committee didn't pass all of my students.
Inference: ¬[The committee didn't pass some of my students]

Romoli points out, correctly, that F&K's structural characterization of alternatives predicts the attested inference of (54) to be non-existent.[28] Thus, suppose the analysis of the prejacent S = **the committee didn't pass all of my students** is (55).

[26] A curious fact to note is that (i), which is the modal-less variant of (50), seems to have the obligatory inference that John did not steal the jewelry, i.e. that John is no different from Bill, and hence sounds a little odd (see footnote 4).

(i) Bill stole the book but not the jewelry. John (only) stole the book.

Given what we have said, this fact means that the set of alternatives to **John stole the book** in this context has to be A = {**book**, **jewelry**} and cannot be A′ = {**book**, **book** ∧ ¬**jewelry**, **jewelry**} or A″ = {**book**, **book** ∧ ¬**jewelry**}: A′ would license no inference about the jewelry and A″ would license the inference that John stole the jewelry. The reason why the set of alternatives cannot be A′ is most likely that this set contains only elements which cannot be negated by EXH. We have no real answer as to why A″ is not available as an option. We suspect the reason to be that the relevant pragmatic scale is missing (see the discussion in Sect. 3.1). Whether or why that hypothesis is true is a question which we will have to leave for future research.

[27] We let "[α[β]]" stand for sentences in which α c-commands β.

[28] More precisely, the problematic prediction is made by F&K's theory in conjunction with the assumption, empirically motivated and built into the definition of EXH, that the computation of scalar implicatures involves the negation of not only logically stronger but also logically independent alternatives (Spector 2006).

(55) $[_{TP}$ PAST $[_{NegP}$ not $[_{VP}$ the committee pass all of my students]]]

F&K's theory predicts both S′ = **the committee didn't pass some of my students** and S″ = **the committee passed some of my students** to be formal alternatives of (55): the first derived by replacing **all** with **some** and the second derived by replacing NegP with VP followed by replacing **all** with **some**. The problem, of course, is that S′ and S″ are symmetric alternatives, which means that (54) is predicted not to license the inference that the committee passed some of my students, which is the attested inference of (54). On the contrary, the facts follow straightforwardly from Atomicity, which simply rules out S″ as a formal alternative of S, because its derivation involves replacing a subpart of a syntactically atomic expression, namely VP, with another.[29] Thus, Atomicity circumvents the overgeneration problem that F&K's theory faces in this particular case, and in cases where a strong scalar item is embedded under negation.[30]

4.2 The switching problem

Romoli (2012b) discusses another case of overgeneration which has to do with structures in which a weak scalar item embeds a strong one, as exemplified in (56).

(56) Some of my students did all of the readings.
*Inference: ¬[All of my students did some of the readings]

Intuitively, (56) does not have the implicature that it is not the case that all of my students did some of the readings, i.e. that some of my students did none of the readings. This implicature is predicted if S′ = **all of my students did some of the readings** is a formal alternative of the prejacent S = **some of my students did all of the readings**.[31] Thus, what we need is for S′ not to be a formal alternative of S. Unfortunately, Atomicity does not rule out S′ as a formal alternative of S, since the former can be derived from the latter without any AT-marked expression having its proper parts replaced by another expression, as (57) shows. Note that the substitution source includes α = $\textbf{some}_{AT}$ and β = $\textbf{all}_{AT}$.

[29] Note that it is also impossible to derive S″ from S by replacement of **all** with **some** followed by replacement of NegP with the newly formed VP = **the committee pass some of my students**. The reason is that this VP is *not* an element of the substitution source: it is neither a lexical item nor a constituent that has been uttered in the context.

[30] Note, incidentally, that the success of our proposal in dealing with the puzzle of indirect implicature provides support for the definition of structural simplification in terms of substitutions alone (as in F&K and in the present paper) rather than in terms of substitutions and deletions (as in Katzir 2007). With substitutions alone, NegP can only be simplified to eliminate negation by substituting VP for NegP, but then, Atomicity mandates that VP will be used as is, with no possibility of replacing **all** with **some**. This accounts for the unavailability of S″. If deletions were allowed, there would be a way to derive the problematic S″ without violating Atomicity: simply delete **not** and then replace **all** with **some**. We thank Roni Katzir for drawing our attention to this point.

[31] Because S′ is logically independent of S and EXH negates logically independent alternatives (see (1)).

(57) Derivation of S′ from the prejacent S:
0. **Some of my student did all of the reading** the prejacent
1. **Some of my student did some$_{AT}$ of the reading** **all** / α
2. **All$_{AT}$ of my student did some$_{AT}$ of the reading** **some** / β

The puzzle, then, is that **[some[all]]** does not implicate ¬**[all[some]]**, although **[all[some]]** is a (logically independent) formal alternative of **[some[all]]**. We might imagine a condition which has the effect of preventing **some** and **all** to "switch places" in the derivation of a formal alternative. But an explanation along this line would founder in the face of the following example.[32]

(58) None of my students did all of the readings.
Inference: All of my students did some of the readings

Under the standard assumption that **none**, at the relevant level of analysis, is to be decomposed into sentential negation taking scope over **some**, i.e. as **[not[some]]** (cf. Penka 2011 and references therein), the observation in (58) can be represented as follows.

(59) [not [some of my students did all of the readings]]
Inference: ¬[not [all of my students did some of the readings]]

This inference is predicted by our account, as we predict **[not[all[some]]]** to be a formal alternative of **[not[some[all]]]** (the derivation of the former from the latter is the same as (57) modulo the fact that every sentence is embedded under negation). The problem, then, is that **[not[some[all]]]** implicates ¬**[not[all[some]]]** but **[some[all]]** does not implicate ¬**[all[some]]**, even though in both cases, the relevant alternative is derivable from and logically independent of the prejacent. A more informal way to frame the problem is the following: certain facts suggest that **some** and **all** can "switch places" only if both are embedded under negation, while we predict that negation makes no difference (hence the title of this subsection).

The solution to this problem, therefore, will consist in preventing the derivation of **[all[some]]** from **[some[all]]** without preventing the derivation of **[not[all[some]]]** from **[not[some[all]]]**. We propose to do this by adding to Atomicity the following constraints on the syntactic derivation of formal alternatives.[33]

[32] According to an anonymous reviewer, the inference in (58) is "not compelling." We agree that this inference is rather weak, but believe that it is possible, concuring with the claim made in Romoli (2012a) based on experiments done in Chemla (2009). However, our aim here is to convince the reader that to the extent that the inference in question is licensed, it can be accounted for by an extension of our proposal which looks rather natural.

[33] Romoli (2012b) proposes a solution to the "switching problem" which has in common with the solution we set out below the stipulation that (i) replacement of scalar items should proceed from the bottom up and (ii) no such replacement may be weakening. The reader is invited to consult Chapter 6 in Romoli (2012b) for details. Romoli does not make use of AT-marking or any similar device, which we argue to be necessary to explicate the notion of "replacement from the bottom up." Furthermore, we do not see how his approach to the switching problem can be generalized to account for the other facts which we have discussed so far.

(60) *Constraints on the derivation of formal alternatives*[34]
  a. Only non-AT-marked expressions are replaceable.
  b. The most deeply embedded replaceable expression must be replaced first.
  c. No replacement may yield a sentence which is logically weaker than the prejacent.

Let us see how these constraints work. First, consider the case where the prejacent is **[not[some[all]]]**. We predict that the following derivation is possible.

(61)	**[not[some[all]]]**	the prejacent
	[not[some[some$_{AT}$]]]	**all / some$_{AT}$**
	[not[all$_{AT}$[some$_{AT}$]]]	**some / all$_{AT}$**

Note that none of the steps result in a sentence logically weaker than the prejacent. Now consider the case where the prejacent is [some[all]]. Logically, there are two ways to derive the problematic [all[some]] from it; either as in (62) or as in (63).

(62)	**[some[all]]**	the prejacent
	[some[some$_{AT}$]]	*__all / some$_{AT}$__
	[all$_{AT}$[some$_{AT}$]]	**some / all$_{AT}$**
(63)	**[some[all]]**	the prejacent
	[all$_{AT}$[all]]	*__some / all$_{AT}$__
	[all$_{AT}$[some$_{AT}$]]	**all / some$_{AT}$**

It turns out that both derivations are excluded by the constraints in (61). In (62), the second step yields a sentence weaker than the prejacent, and in (63), the second step replaces a constituent which is *not* the most deeply embedded replaceable expression since **all** is replaceable and c-commanded by **some**.

We thus have a solution to the second overgeneration problem as well. It can be verified (though we won't do so here) that the addition of the constraints in (61) to our proposal does not affect how it accounts for the facts discussed in the previous sections.

5 Summary

We have proposed a characterization of formal alternatives which does justice to certain observations, some of them novel, which pose a challenge for existing theories of alternative generation. The structural constraints we propose on F(S) show intriguing parallels between the derivation of formal alternatives and the derivation of sentences. Both involve an initial set of expressions to be used, the "substitution source" in the case

[34] While (60a) might seem more arbitrary than (60b), which says that the derivation of alternatives, like many other syntactic processes, proceeds "bottom up," and (60c), which makes clear functional sense, it is not entirely without rationale: barring replacement of AT-marked expressions can be understood as barring vacuous replacement steps. Thus, (60a) rules out replacement of A with B followed by replacement of B with C, which amounts to the same as replacement of A with C.

of formal alternatives and the “numeration” in the case of sentences. Expressions in the substitution source are treated as syntactically atomic, making this construct essentially a sort of numeration, albeit one that is contextually determined. The derivation of formal alternatives must proceed from the bottom up, with the relevant rule applying to more deeply embedded constituents before applying to less deeply embedded ones. This condition makes the syntactic derivation of formal alternatives strikingly similar to the syntactic derivation of sentences, prompting the question of whether/to what extent the former might be a “cooptation” of the latter. These are questions that we hope will be addressed in future research.

Acknowledgments We thank Luka Crnič, Danny Fox, Manfred Krifka, Clemens Mayr, Hazel Pearson, Sophie Repp, and Jacopo Romoli for fruitful discussion. We thank Roni Katzir, who reviewed this paper non-anonymously, for his many valuable critiques, and thank an anonymous reviewer for several helpful suggestions. Special thanks go to Danny Fox for recommending substantial changes to an earlier version of this paper. The present work is supported by a grant from the Deutsche Forschungsgemeinschaft (Sonderforschungsbereich 632, Projekt A2).

References

Atlas, J.D., and S.C. Levinson. 1981. *It*-clefts, informativeness, and logical form: Radical pragmatics. In *Radical pragmatics*, ed. P. Cole, 1–61. New York: Academic Press.
Chemla, E. 2009. Presuppositions of quantified sentences: Experimental data. *Natural Language Semantics* 17(4): 299–340.
Chierchia, G., D. Fox, and B. Spector. 2012. The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In *Semantics: An international handbook of natural language meaning*, ed. P. Portner, C. Maienborn, and K. von Heusinger, 2297–2332. Berlin: de Gruyter.
Chomsky, N. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.
Crnič, L. 2012. Focus particles and embedded exhaustification. *Journal of Semantics* 30(4): 533–558.
Fox, D. 2007a. Free choice disjunction and the theory of scalar implicatures. In *Presupposition and implicature in compositional semantics*, ed. U. Sauerland and P. Stateva, 71–120. Basingstoke: Palgrave-Macmillan.
Fox, D. 2007b. Pragmatics in linguistic theory. MIT class notes.
Fox, D., and R. Katzir. 2011. On the characterization of alternatives. *Natural Language Semantics* 19: 87–107.
Groenendijk, J., and M. Stokhof. 1984. Studies on the semantics of questions and the pragmatics of answers. PhD thesis, University of Amsterdam.
Hintikka, J. 1969. Semantics for propositional attitudes. In *Philosophical logic*, ed. J.W. Davis et al., 21–45. Dordrecht: Reidel.
Horn, L. 1989. *A natural history of negation*. Chicago: The University of Chicago Press.
Katzir, R. 2007. Structurally-defined alternatives. *Linguistics and Philosophy* 30: 669–690.
Katzir, R. 2008. Structural competition in grammar. PhD thesis, MIT.
Klinedinst, N. 2004. Only scalar *only*. Handout presented at the ‘Presupposition and Implicature’ workshop, Paris.
Krifka, M. 1995. The semantics and pragmatics of polarity items. *Linguistic Analysis* 25(3–4): 209–257.
Kroch, A. 1972. Lexical and inferred meanings for some time adverbials. *Quarterly Progress Reports of the Research Laboratory of Electronics* 104: 260–267.
Lewis, D. 1988. Relevant implication. *Theoria* 54(3): 161–174.
Magri, G. 2009. A theory of individual-level predicates based on blind mandatory scalar implicatures. *Natural Language Semantics* 17(3): 245–297.
Magri, G. 2011. Another argument for embedded scalar implicatures based on oddness in downward entailing environments. *Semantics and Pragmatics* 4(6): 1–51. doi:10.3765/sp.4.6.
Penka, D. 2011. *Negative indefinites*. Oxford: Oxford University Press.
Romoli, J. 2012a. A problem for the structural characterization of alternatives. Unpublished manuscript, Harvard University.

Romoli, J. 2012b. Soft but strong: Neg-raising, soft triggers, and exhaustification. PhD thesis, Harvard University.

Rooth, M. 1992. A theory of focus interpretation. *Natural Language Semantics* 1: 75–116.

Sauerland, U. 2004. Scalar implicatures in complex sentences. *Linguistics and Philosophy* 27(3): 367–391.

Sauerland, U. 2012. The computation of scalar implicatures: Pragmatic, lexical or grammatical? *Language and Linguistic Compass* 6: 36–49.

Spector, B. 2006. Scalar implicatures: Local or global? Exhaustivity and Gricean reasoning. Unpublished manuscript, Institut Jean Nicod, Paris.

Spector, B. 2010. Scalar implicatures and grammar alternatives for focus and scalar implicatures (based on Fox and Katzir). Handout for ESSLLI, Copenhagen.

von Fintel, K., and I. Heim. 1997. Pragmatics in linguistic theory. MIT class notes.