# Symmetry resolution and blocking[1]

Bernhard SCHWARZ — *McGill University*
Michael WAGNER — *McGill University*

**Abstract.** We propose a blocking condition that limits the possible effects of exhaustification: exhaustifying a sentence $\phi$ cannot output a meaning that could be expressed as the basic, non-exhaustified, meaning of a sentence that is no more complex than $\phi$. We propose that this blocking condition provides a solution to the so-called symmetry problem. We compare our solution to the proposal in Katzir (2007) and Fox and Katzir (2011), which instead prevents exhaustification from excluding alternatives that are more complex than the assertion. In support of our blocking condition, we argue that Katzir and Fox's complexity filter does not actually solve the symmetry problem in full, and in fact is incompatible with exhaustification data. We also argue against a central auxiliary assumption that Katzir and Fox's account appeals to, viz. the assumption that symmetry cannot be resolved by context.

**Keywords:** exhaustification, symmetry problem, complexity constraint, blocking.

## 1. Introduction

Exhaustification is a process that can tacitly strengthen basic sentence meanings. Prominent instances of this effect are so-called scalar implicatures, such as those often associated with uses of the existential determiner *some*. For example, (1) has a basic meaning that is merely existential, that there are players who responded, which we will state as in (2). But (1) can, and typically does, carry the stronger total meaning stated in (3). That is, (1) can be understood as also implying that there were players who did not respond, the second conjunct of (3). The presence of this implication is confirmed, for instance, by its felicitous denial, as in (4).

(1)     Some players responded.

(2)     Basic meaning of (1):
        SOME(P)(R)

(3)     Available strengthened meaning of (1):
        SOME(P)(R) $\wedge$ ¬ALL(P)(R)

(4)     A:   Some players responded.
        B:   No, all of them did!

It is this sort of strengthening that is commonly said to arise from *exhaustification*. Its nature is a matter of debate. In particular, while the mapping from basic to strengthened meaning is classically attributed to Gricean pragmatics, it has also been credited to a grammatically encoded operator (e.g., Chierchia et al. 2011). However, since the choice between these options

does not bear on the arguments below, we will refrain here from committing to one view or another.

Reflecting this lack of commitment, we will often use the term *strengthening* to refer to the phenomenon in question, which might be more obviously theory neutral than *exhaustification*. But sometimes we will also refer to the inferences added by strengthening as *implicatures*.

In the broadest characterization, strengthening takes the form in (5). That is, the strengthened meaning of sentence $\phi$, $\text{Exh}_{\text{Excl}}(\phi)$, is the conjunction of the basic meaning of $\phi$, $\text{Basic}(\phi)$, with the negations of the members of a set of propositions *Excl*. We will also refer to *Excl* as the *exclusion set*, and to its members as the *excluded propositions*.

(5)    Strengthened meaning:
       $\text{Exh}_{\text{Excl}}(\phi) = \text{Basic}(\phi) \wedge \bigwedge \{\neg p : p \in \text{Excl}\}$

The format in (5) can be illustrated with the equality in (6). That is, the strengthened meaning of sentence (1) given in (3) can be understood as due to strengthening relative to an exclusion set that contains just the proposition $\text{ALL}(P)(R)$, that all players responded.

(6)    $\text{Exh}_{\{\text{ALL}(P)(R)\}}(1) = \text{SOME}(P)(R) \wedge \neg\text{ALL}(P)(R)$

A central task for a theory of strengthening is to explain limitations on the range of actually observed sentence meanings of the format in (5). To illustrate one important type of limitation, we can continue to focus on sentence (1). Consider a hypothetical exclusion set that only contains $\neg\text{ALL}(P)(R)$, the proposition that *not* all players responded. As stated in (7), strengthening (1) relative to this exclusion set would yield $\text{ALL}(P)(R)$, the proposition that all players responded. As recorded in (8), however, such a strengthened meaning for (2) is unavailable. This verdict can be confirmed with the observation that the implication that this strengthening would add cannot be felicitously denied in response to (1). That is, in contrast to (4), the attempted denial in the dialogue in (9) is unequivocally infelicitous.

(7)    $\text{Exh}_{\{\neg\text{ALL}(P)(R)\}}(1) = \text{SOME}(P)(R) \wedge \neg\neg\text{ALL}(P)(R) = \text{ALL}(P)(R)$

(8)    Unavailable strengthened meaning of (1):
       $\text{ALL}(P)(R)$

(9)    A:   Some players responded.
       B:   #No, not all of them did!

Why is it that the basic meaning of (1) can be strengthened with the conjunct $\neg\text{ALL}(P)(R)$, yielding (3), but not with $\text{ALL}(P)(R)$, yielding (8)? More generally, for any case where the basic meaning of a sentence $\phi$ can be strengthened with the conjunct $\neg p$, as in (10a), but not with the conjunct p, as in (10b), the question is on what grounds the former is available while the latter is not. Since the two strengthenings in (10) are in a transparent sense *symmetric*, this problem is known as the *symmetry problem*. First observed as a challenge to the Gricean reasoning in Kroch (1972), the symmetry problem has been central in analyses of exhaustification.

(10)    a.    $\text{Basic}(\phi) \wedge \neg p$
        b.    $\text{Basic}(\phi) \wedge p$

In the dominant approach, the symmetry problem is solved by constraining the membership of

the exclusion set Excl. On this view, sentence (1) permits the strengthening in (3), but not the one in (8), because for the interpretation of (1), the exclusion set may contain just ALL(P)(R), but cannot contain ¬ALL(P)(R). More generally, if for a given sentence $\phi$, the strengthening (10a) is attested while (10b) is not, this is assumed to be due to a constraint which relative to $\phi$, allows for the exclusion set to contain just p, but prevents it from containing ¬p.

In a prominent school of thought, this approach is developed by constraining the exclusion set with reference to the *linguistic form* that would be required to express its members. The idea is that for strengthening of $\phi$, any proposition in the exclusion set must be the basic meaning of a sentence, a *linguistic alternative* to $\phi$, that is no more complex than $\phi$ itself. Horn (1972), for example, postulates lexical scales, today often called 'Horn scales', Levinson (1983) invokes an 'economy' condition, and Horn (2000) a 'brevity' condition. Katzir 2007 and Fox and Katzir 2011 formalize a notion of complexity that requires formal alternatives to be no more complex than the asserted one. This formal *complexity filter on Excl* (CF) can be stated as in (11).

(11)     Complexity Filter on Excl:

         $\text{Excl}(\phi) \subseteq \{\text{Basic}(\psi)$: $\psi$ is no more complex than $\phi\}$

CF has the intended effect for sentence (1). The unavailable strengthened meaning (8) could potentially arise for (1) from excluding the proposition ¬ALL(P)(R). This proposition is the basic meaning of, for example, the sentence in (12). However, due to the presence of negation *not* in the structure, (12) is more complex than (1). In fact, it seems that *any* sentence whose exclusion would yield the unavailable strengthening (8) is more complex than (1). If so, then CF correctly prevents this strengthening.

(12)     Not all players responded.

In this paper, we will explore an alternative solution to the symmetry problem. We will pursue the idea that rather than by constraining Excl, the problem is solved through a condition that can block potential strengthened meanings as *permissible outputs* of Exh. To introduce the approach, consider again (1), and let us juxtapose it with sentence (13). Note that the basic meaning of (13) is identical to the unavailable strengthened meaning for (1) in (8). We now propose that since (13) is moreover no more complex than (1), its existence *blocks* the use of (1) in the strengthened meaning in (8), rendering this strengthening unavailable.

(13)     All players responded.

So we propose that a potential strengthened meaning for a sentence $\phi$ is unavailable if this potential strengthened meaning amounts to the *basic* meaning of a sentence that is no more complex than $\phi$. This *blocking constraint* (BC) is a condition on the output of Exh, rather than a condition on Excl, an input to Exh. We can state this constraint as in (14).[2]

(14)     Blocking constraint (BC)

         $\text{Exh}_{\text{Excl}}(\phi)$ is blocked if: $\exists\psi[\text{Basic}(\psi) = \text{Exh}_{\text{Excl}}(\phi) \wedge \psi$ is no more complex than $\phi]$

We can think of BC as a constraint that requires strengthening to be in the service of abbreviation. That is, strengthening must serve to *reduce complexity* of linguistic form that would

---

[2]Geurts (2011: ch. 6) appeals to complexity for listener's selection of cells from a partition of the speaker's "intentional state". A detailed comparison is deferred to future work.

otherwise be required to express a given meaning. We might accordingly think of BC as an economy constraint rooted in the Gricean Manner maxim: Even though sentence (15) has SOME(P)(R) $\wedge$ ¬ALL(P)(R) as its basic meaning, it is visibly more complex than (1) itself. It in fact seems that *any* sentence with this basic meaning is more complex than (1). If so, then, the available strengthening in (3) is permitted by BC.

(15)    Some but not all players responded.

The notion of complexity that BC appeals to may well be the very same notion as the one invoked by CF. But crucially, CF appeals to complexity in constraining the *input* to exhausti-fication, while BC references complexity in constraining permissible exhaustified *outputs*. For basic data like (1), the effects of CF and BC coincide. Not only do both correctly exclude the unavailable strengthening in (8), but both also permit the available strengthening in (3).

Our argument in favor of BC in the following has two sides. In section 2, we argue that CF cannot be maintained as a constraint on Excl, and hence BC is needed to account for basic cases like (1). In section 3, we then examine cases that are compatible with CF, but where CF is insufficient. These are cases where an unavailable implicature would be derivable based on alternatives that are not more complex than the asserted sentence. These examples will also reveal shortcomings of several auxiliary assumptions made in the account based on CF in Katzir (2007) and Fox and Katzir (2011), which we explore in section 4.

## 2. The case against the complexity filter

The symmetry problem introduced in Section 1 arises when a sentence can convey one of two conceivable contradictory implicatures but not the other. CF aims to address this question, but is formulated as a general constraint on Excl, not just for situations in which symmetry arises. The effects of CF are therefore predicted to be observable even in cases where embedding removes the potential contradiction. To start, consider (16) the exchange adapts an example used in Fox and Katzir (2011) for a different purpose (see footnote 3 below).

(16)    A: Tell me what exactly the detective has concluded up to this point
        B: The detective concluded that Jed stole some of the books.

Sentence (16B) is logically compatible with the simultaneous exclusion all sorts of alternatives that can arise from altering the predicate *stole some of the books*. For example, the truth of (16B) is compatible with the exclusion of each of the alternatives listed in (17). Note that this compatibility is crucially enabled by the embedding under *conclude*. Even though the embed-ded clauses in (17a) and (17b) are complementary, the embedding under *conclude* ensures that (17a) and (17b) as a whole can both be false at the same time.

(17)    a.    The detectives concluded that Jed stole all of the books.
        b.    The detectives concluded that Jed did not steal all of the books.
        c.    The detectives concluded that Jed did not steal the jewels.
        d.    The detectives concluded that Jed did not steal the jewels unless they were hidden in the books

Note that the sentences (17b)–(17d) are more complex than (16B), and that it does not seem possible to avoid the added complexity while expressing the same propositions. The complexity

filter CF therefore predicts that those propositions cannot be in the exclusion set. But this prediction seems incorrect. It seems possible for (16B) to be strengthened by the negation of each of these propositions. We found that each of those implicatures can be felicitously denied in a response to (16B), perhaps by a third interlocutor, as in (18).

(18)    a.    C: No! The detective also concluded that Jed did not steal all of the books!
        b.    C: No! The detectives also concluded that Jed did not steal the jewels!
        c.    C: No! The detectives also concluded that Jed did not steal the jewels unless they were hidden in the books

The possibility of these responses shows that, at least in a context in which the details matter, as established in our case by sentence (16A), the proposition expressed by the more complex alternatives in (17b) can be in Excl. Importantly, while this is incompatible with CF, it is compatible with BC. To illustrate, sentence (19) is more complex than (16B), and the proposition that it expresses does not seem to permit a less complex encoding. Hence, nothing blocks a strengthening of (16B) that excludes just the proposition expressed by (17b). It is clear, in fact, that a strengthened meaning that excludes all of the propositions expressed in (17) cannot be expressed as a basic meaning without an obvious increase in complexity relative to (16B), hence such a strengthening is permitted by BC.[3]

(19)    The detectives concluded that Jed stole some but not all of the books.

It is important to acknowledge that, while we take (16) to show that CF cannot be maintained, the strengthening observed for sentences with embedding sometimes *is* constrained in a way that seems aligned with CF. Fox and Spector (2018: 15) discuss cases of embedding under a universal operator that seem to pattern differently from (16). In contrast to the reply in (18a), the one on (20B) seems infelicitous. This reply should be felicitous if (20A) could be understood as excluding potential alternatives like those in (21). But since the presence of negation renders those more complex than (20A), the infelicity of (20B) is in line with CF.

(20)    a.    A:  You're required to buy some of the furniture.
        b.    B: # No! We are not allowed to buy all of it.

---

[3]The example in Fox and Katzir (2011) that (16) adapts, given in (i), involves association with *only*. Fox and Katzir use (i) to support the claim that more complex alternatives can be excluded by *only* if they arise from substitutions of material that is *explicitly provided in the context*, here by the first sentence of (ia) and the second sentence of (ib).

(i)    a.    Detective A only concluded that the robbers [stole the books]$_F$. Detective B concluded that the robbers [stole the books but not the jewelry]$_F$.
        b.    Detective B concluded that the robbers [stole the books but not the jewelry]$_F$. Detective A only concluded that the robbers [stole the books]$_F$.
        Inference: ¬ Detective A concluded that the robbers stole the books but not the jewelry.
        Inference: ¬ Detective A concluded that the robbers stole the books and the jewelry.

In Fox and Katzir (2011), the implicit claim is that such contextual support is *necessary* to enable to the exclusion of the relevant alternatives. However, they do not actually consider controls where the relevant sort of contextual support is missing. In fact, aligned with our intuitions about (16), observations in Hirsch (2024) indicate that such support is not in fact needed (cf. also Hirsch and Schwarz 2023). Even without being continued by the second sentence, the first sentence in (ia) could well be taken to entail that Detective A did *not* conclude that the robbers did not steal the jewelry. Likewise for (ib).

(21)  a.  You're not allowed to buy all of the furniture.
      b.  You're required to not buy all of the furniture.

However, given our interpretation of (16), a different explanation must be sought. While blocking sets limits on what can be excluded, it is clearly not the only constraint on strengthening. Which implicatures that are possible in principle are actually attributed to a speaker generally depends on context. We suggest that the infelicity in this case has to do with the fact that the question of whether there is an additional requirement not to buy all furniture seems contrived without further context, and hence failing to addressing this question is not taken to be meaningful here. In support of this suggestion, (22) shows that a context that is expected to motivate B's pedantic objection indeed renders it more natural (as does adding an additive). In (23), the pendantic objection seems justified even more clearly, as the stakes are raised further, and here the relevant implicature in fact seems quite unexceptional.[4]

(22)  a.  B: Did you listen to the instructions? Tell me exactly what we have to do.
      b.  A: We're required to buy some of the furniture.
      c.  B: No! We are (also) not allowed to buy all of it.

(23)  a.  B: Did you listen to the instructions? Tell me exactly how to diffuse the bomb.
      b.  A: One has to press some of the buttons.
      c.  B: No! You (also) must not press all of them.

Moreover, while the examples incompatible with CF presented so far all invoke contextual support in a dialogue, this is not a general feature of such data. Similar inferences can arise quite naturally without much context. The example in (24) is a case in point (cf. also so-called conditional perfection, see, e.g., Horn 2000). The felicity of objections like those in (25) suggests that strengthening can exclude alternatives of arbitrary complexity.

(24)  To get a cookie, you must first eat an apple.

(25)  a.  Wait, I thought I would also have to eat an orange!
      b.  Wait, I thought I would also have to either eat an orange or take a walk!
      c.  Wait, I thought I would also have to not use my phone for one hour!

Taken together, these examples suggest that CF is not viable as a general constraint on the exclusion set: Strengthening based on alternatives that should be screened out by CF are often attested once contradictions are avoided by eliminating symmetry through embedding.

If CF is not viable, this means that we need a different solution to the symmetry issue posed by cases like (1). We saw in the previous section that BC potentially provides such a solution. In the following, we will further motivate BC based on cases of symmetry in which CF cannot possibly provide a solution, since the problematic formal alternatives are no more complex than what is asserted.

---

[4]The example in (i), from our handout distributed at the conference, may be harder to fix, but even this example seems to improve if the context asks for an exact report on what speaker A knows.

(i)  A:  Every week, some players responded.
     B: # Wait, I thought (it's also true) that every week some of them did not respond!

## 3. The case for blocking

CF prevents the unwanted strengthening in (8) for (1) by exploiting a limitation of English relative to the structure of (1). In contrast to ALL, ¬ALL is not lexicalized. Expressing ¬ALL therefore requires a syntax that transparently encodes negation, but adding negation to (1) would increase complexity. It is in virtue of this increase that CF here has the intended effect. Notably, then, if the meaning ¬ALL *were* lexicalized, then CF would not help.

Exploring this line of thought, we will now present a challenge to CF as a general solution the symmetry problem. We will discuss cases where unwanted meanings are carried by acceptable sentences which lack the complexity that would be needed for CF to screen them out. We refer to the such sentences as *simplex threats*. In each case, the unwanted meaning that the simplex threat supports, while potentially output by Exh, is correctly blocked by BC. Hence these simplex threats furnish an argument for the presence of BC as a constraint on exhaustification.

### 3.1. Three cases of simplex threats

Before delving in, let us clarify that while we seem to be the first to interpret the simplex threats discussed below as support for a blocking condition, our discussion crucially builds on previous work. The simplex threat based on expressed negation that we will present (Section 3.1.1) is a variant of a case discussed in Romoli (2013). Like we do for our case, Romoli interpreted his example as a challenge for CF. The simplex threats with antonyms (Section 3.1.2) are related to data in Breheny et al. (2018), and directly inspired by cases in Buccola et al. (2022). Again, like we do for our variants, Breheny et al. and Buccola et al. used their data to challenge CF. Finally, the case of contextually supported simplex threats (Section 3.1.3) is centrally discussed in Katzir (2007) and Fox and Katzir (2011). We will discuss their additional assumptions aimed at diffusing this simplex threat (and potentially the others as well) in Section 4.

### 3.1.1. Simplex threats enabled by expressed negation

The first type of simplex threat we present arises in certain data where an overt negation interacts scopally with another operator (cf. Romoli 2013). Here we focus on a case where negation is interpreted in scope of an existential indefinite, as in (26).[5]

(26)    Some players have not responded.

The basic meaning of this sentence can be given in (27) (where D is the domain of individuals). As stated in (28), we observe an asymmetry in the intuitively available strengthenings much like the one that we described for (1). Sentence (26) can be understood as conveying that some but not all players responded, see (28a), but not as conveying that none of them did, see (28b).

(27)    Basic meaning of (26):
        SOME(P)(D–R)

(28)    a.    Available strengthened meaning of (26):
              SOME(P)(D–R) ∧ ¬ALL(P)(D–R) = ¬ALL(P)(R) ∧ SOME(P)(R)

---

[5]This example has an advantage over Romoli's in that it cannot be explained by the additional 'Atomicity' constraint proposed in Trinh and Haida (2015) to save the CF account (which we do not have the space to discuss here).

b.  Unavailable strengthened meaning of (26):
    SOME(P)(D–R) ∧ ¬SOME(P)(R) = NO(P)(R)

CF fails to derive this asymmetry. Sentence (29a) below, which permits a surface scope reading with the basic meaning NO(P)(R), is no more complex than (26). Therefore, since (28a) is the strengthened meaning of (26) that can result from, say, the exclusion set {NO(P)(R)}, CF correctly permits this strengthening. However, sentence (29b), with the basic meaning SOME(P)(R), is not more complex than (26), either. Therefore, since (28b) is the strengthening for (26) that would result from, say, the exclusion set {SOME(P)(R)}, CF permits this unwanted strengthening as well.

(29)    a.  All players have not responded.
        b.  Some players have responded.

If CF does not capture the asymmetry in (28), then what does? We propose that the blocking condition BC has the intended effect. BC correctly prevents (28b) as a strengthened meaning for (26). It is blocked due to the existence of a sentence that is no more complex than (26) and that has (28b) as its basic meaning. For example, (29a) is such a sentence. We moreover suggest that, in contrast, BC permits the strengthened meaning in (28a). To be sure, there are sentences that have (28a) as their basic meaning. Example (30) is a case in point. However, (30) is more complex than (26). We submit that this is in fact true for all sentences with the basic meaning (28a). If so, then BC correctly permits (28a) as a strengthened meaning for (26).

(30)    Some but not all players have responded.

The crucial feature of example (26) that distinguishes it from our original case in (1) is that it already contains negation: It is on those grounds that a symmetric pair of linguistic alternatives can be derived by removing rather than adding complexity. We now turn to a case in which simplex threats arise by simple substitutions, based on antonymic content vocabulary.

### 3.1.2. Simplex threats enabled by antonyms

Another type of simplex threat can arise in cases where a predicate that appears in the scope of another operator has a lexical antonym (cf. Breheny et al. 2018, Buccola et al. 2022). Here we consider the case where a predicate of this sort appears in the scope of an existential indefinite.

(31)    Some players were inside.

The basic meaning of (31) is given in (32). As stated in (33), we once again observe an asymmetry in the intuitively available strengthenings. (31) can be understood as conveying that some but not all players were inside, see (33a), but not as conveying that all were inside, see (33b).

(32)    Basic meaning of (31):
        SOME(P)(In)

(33)    a.  Available strengthened meaning of (31):
            SOME(P)(In) ∧ ¬ALL(P)(In)

        b.  Unavailable strengthened meaning of (31):
            SOME(P)(In) ∧ ¬¬ALL(P)(In) = ALL(P)(In)

It may initially appear as though CF captured this asymmetry, as CF may seem to apply to (31) in the same way it applies to the basic example (1). To begin, since sentence (34a), with basic meaning ALL(P)(In), is no more complex than (31), CF permits the exclusion set to be, for example, {ALL(P)(In)}. Hence it permits exhaustification that yields the available strengthened meaning in (33a). At first sight, CF seems to also correctly prevent the undesirable inference. After all, while sentence (34b) transparently has ¬ALL(P)(In) as its basic meaning, the presence of syntactic negation renders (34b) more complex than (31).

(34)   a.   All players were inside.
       b.   Not all players were inside.

However, closer inspection reveals CF to be insufficient to prevent the exclusion set from containing ¬ALL(P)(In). The reason is that this same proposition is also expressed by a different sentence, one that crucially is no more complex than (31), viz. (35), where *inside* is replaced by its antonym *outside*:

(35)   Some players were outside.

Sentence (35) transparently expresses the proposition SOME(P)(Out). Since the complementarity of the antonyms *inside* and *outside* guarantees the equality of the two sets Out and D–In, this basic meaning is equivalent to SOME(P)(D–In) and hence to ¬ALL(P)(In). Therefore, CF does not prevent an exclusion set containing this proposition after all. Allowing for, say, {¬ALL(P)(In)} as the exclusion set, CF after all fails to prevent the unavailable strengthened meaning in (33b). In other words, sentence (35) is a simplex threat.

The blocking condition BC captures the asymmetry in (33). BC correctly blocks (33b) as a strengthened meaning for (31), due to the existence of a sentence like (34a), which is no more complex than (31) and has (33b) as its basic meaning. While there are sentences with the basic meaning (33a), we submit that they all are more complex than (31). For example, both of the sentences in (36) have (33a) as their basic meaning. But given their complexity relative to (31), their existence does not stand in the way of (31) having this same meaning as a result of strengthening. We submit that this extends to all sentences with the basic meaning (33a). If so, then BC correctly allows for (31) to acquire this meaning from strengthening.

(36)   a.   Some but not all players were inside.
       b.   Some players were inside and some were outside.

The example in (31) illustrates that symmetric linguistic alternatives can evade the grasp of CF in virtue of featuring predicates that have antonyms. By introducing semantic negation without adding syntactic complexity, replacing such a predicate with its antonym can give rise to a simplex threat. We will next attend to yet another type of simplex threat, cases where a problematic alternative should be permitted in virtue of features of the linguistic context.

### 3.1.3. Simplex threats enabled by context

Under a construal of the notion of syntactic complexity that seems initially plausible and that we have tacitly assumed so far, a structure $\beta$ counts as no more complex than a structure $\alpha$ only if $\beta$ can be obtained from $\alpha$ by merely deleting constituents of $\alpha$ or replacing terminal nodes. However, based on an observation in Matsumoto (1995), Katzir (2007) suggests that for

the purposes of the theory of exhaustification, this condition is too demanding. Consider the conjunctive sentence given by the concatenation of the clause in (37a) with (37b).

(37)    a.    It was warm yesterday ...
          b.    ... and it is a little bit more than warm today.

While the basic meaning of (37a) merely implies that it was warm yesterday, Matsumoto observed that as part of the conjunctive sentence, (37a) intuitively permits a strengthened meaning that implies in addition that yesterday it was *not* a little bit more than warm. Hence strengthening can exclude the proposition expressed by the linguistic alternative to (37a) stated in (38).

(38)       It was a little bit more than warm yesterday ...

Under CF, this is surprising since (38) appears to have more structure than (37a). In fact any sentence with the same basic meaning as (38) seems to have more structure than (37a). This leads Matsumoto to conclude that a complexity constraint on the exclusion set is not viable. However, Katzir suggests that Matsumoto's observation can be reconciled with CF by properly construing the notion of complexity that CF is taken to reference. Katzir proposes that for the purposes of CF, a constituent that appears in the linguistic context of an utterance is treated like a lexical item, in the sense that CF applies *as though* this constituent lacked internal structure. Hence in the case at hand, given that the constituent *a little bit more than warm* appears in the linguistic context of (37a), substituting it for *warm* will not actually be taken to incur an increase of complexity for the purposes of CF. If so, then (38) counts as no more complex than (37a), reconciling CF with the observed strengthening.

Assuming the revised understanding of complexity motivated by Matsumoto's example, consider now the sentence given by the concatenation of (39a) with (39b), which adapts similar examples discussed in Katzir (2007) and Fox and Katzir (2011).

(39)    a.    Last week, some of them responded, ...
          b.    ... and this week (too), not all of them responded.

We take it that, with *them* referring to the players, (39) can give rise to the very same strengthening asymmetry that we initially detected for sentence (1). Even in the context of the continuation (39b), (39a) can be understood as conveying that some but not all players responded last week, as in (40a). This meaning for (39a) can be brought out clearly by including the additive particle *too* in (39b), thereby forcing the strengthened interpretation for (39a) in (40a). In contrast, as indicated in (40b), with or without the additive particle, (39a) certainly cannot be understood as conveying that last week all of the players responded.

(40)    a.    Available strengthened meaning of (39a):
             SOME(P)(R) $\wedge$ ¬ALL(P)(R)

          b.    Unavailable strengthened meaning of (39a):
             SOME(P)(R) $\wedge$ ¬¬ALL(P)(R) = ALL(P)(R)

Under the amended construal of complexity, *both* of the linguistic alternatives to (39a) in (41) qualify as no more complex than (39a). In particular, this is true for (41b), on the grounds that the constituent *not all of them responded* appears in the linguistic context (39b). CF therefore permits the unwanted strengthening in (40b) alongside the attested one in (40a). Hence the alternative (41b) constitutes yet another case of a simplex threat.

(41)  a.  Last week, all of them responded …
      b.  Last week, not all of them responded …

On the other hand, the asymmetry in (40) is once again captured by BC. BC applies to this case in the very same way as it does to our initial example in (1). The unavailable strengthening in (40b) is blocked by sentence (41a), for example, given that (41a) has (40b) as its basic meaning. The available strengthening in (40a) correctly evades blocking. While sentence (42) has (40a) as its basic meaning, it is more complex than (39a). Assuming that this is true for all sentences with this basic meaning, BC correctly permits the strengthening in (40a).

(42)  Last week, some but not all of them responded …

To recap, we have now seen three types of simplex threats, cases where CF fails to exclude alternatives that would yield unavailable strengthenings. We also saw that in each case, the unavailable strengthening is correctly blocked under BC. In the absence of an alternative explanation for why the potential strengthenings in question are unavailable, this furnishes an argument that BC is needed as a constraint on the output of exhaustification.

How conclusive is this argument? We will next scrutinize a possible alternative approach to the simplex threats that we presented, emerging from proposals in Katzir (2007) and Fox and Katzir (2011). We will argue that the alternative account does not in fact successfully remove the challenge posed by simplex threats. Moreover, we will argue that a crucial assumption that the alternative is based on, which we will call the *Stalemate Assumption*, cannot be accurate.

## 4. Symmetric stalemate to the rescue?

Our argument for the blocking condition BC could perhaps be considered conclusive if it were safe to assume that CF would be the only constraint on the exclusion set. However, it is of course possible that further constraints on Excl can be motivated, and that these constraints offer another account of our simplex threat data. We will now discuss one attempt to supplement CF in this way, viz. a cluster of assumptions motivated in Katzir (2007) and Fox and Katzir (2011).

Katzir and Fox propose that in the grammar of exhaustification, CF operates alongside another central principle that regulates the membership of Excl. The principle holds that in cases where grammatical constraints on linguistic alternatives permit two alternatives that express symmetric propositions, it is not possible for just one of those two propositions to be included in Excl, even in cases in which context would motivate choosing one over the other. In the terms of Katzir and Fox, context can never *break* symmetry. In other words, the claim is that symmetry leads to a stalemate between two mutually contradictory inferences.[6]

(43)  Stalemate Assumption:
      Symmetric alternatives permitted by CF cannot be selectively omitted from Excl

Naturally, when CF is supplemented with the Stalemate Assumption, the predictions for our simplex threat data change. The central feature of those data is precisely that there are two

---

[6]Under this formulation, symmetric alternatives must either all be included or all be excluded from the exclusion set—a stalemate cannot be broken by context. If one assumes with Fox (2007) that exclusives in fact ignore alternatives that are not 'innocently excludable,' then the stalemate assumption could instead simply say that no symmetric alternatives can be omitted from Excl.

symmetric alternatives which satisfy CF. The prediction is now that in the presence of symmetric alternatives, *none* of the the potential inferences will be available. Concretely, as matters stand, in each of (28), (33), and (40), the available strengthening in (a) is now prevented as much as the unavailable strengthening in (b).

For our third simplex threat case, the case invoking the linguistic context, Fox and Katzir (2011) expressly appeal to the Stalemate Assumption to avoid the undesirable implicature. To see how this works, consider again the coordinated sentence given by (44a) and (44b), repeated from (39), as well as the two possible alternatives to (44a) in (45a) and (45b), repeated from (41).

(44)   a.   Last week, some of them responded, ...
       b.   ... and this week (too), some of them responded.

(45)   a.   Last week, all of them responded ...
       b.   Last week, not all of them responded ...

Recall that, given the baseline data discussed in Matsumoto (1995), the contextual support from (44b) should allow for (45b) to serve as an alternative to (44a), deriving the undesirable inference. The Stalemate Assumption, however, ensures that if both alternatives are available, *neither* inference can be derived. Refining the assumptions about the effect of context, Fox and Katzir (2011: fn. 16) suggest that including the meanings of contextually provided complex alternatives in the exclusion set is merely optional. This refinement has the intended effect for (44a). If we do not include (45b) in Excl, exhaustification can output the strengthened meaning in (40a); if we include both, neither implicature is derived.

But can this strategy be extended to the other two types of simplex threats that we have seen? Once we assume the Stalemate Assumption, the prediction for these cases is now that we should see no strengthening—so the problematic strengthening based on the simplex threat is correctly ruled out. But how to permit the strengthening that *is* available?

The strategy would again have to be to somehow eliminate the undesirable alternative. Fox and Katzir point to one possible route. They propose that in addition to being constrained by CF, the form of linguistic alternatives is shaped by F(ocus)-marking. The possible alternatives to a given sentence are assumed to differ in this sentence only with regard to F-marked constituents. Put differently, modulo F-marked constituents, an alternative to a given sentence is required to be an exact copy of that sentence. Certain alternatives can then be avoided by assuming a particular F-placement. To show this, we return to the relevant examples (26) and (31), repeated in (46a) and (46b), respectively.

(46)   a.   Some players have not responded.
       b.   Some players were inside.

For sentence (46a), the problematic pair of symmetric alternatives is stated again in (47a-i), repeated from (29). Likewise, the relevant symmetric alternatives to (46b) are stated in (47b), repeated from (34a) and (35).

(47)   a.   (i)    All players have not responded.
            (ii)   Some players have responded.

       b.   (i)    All players were inside.
            (ii)   Some players were outside.

Suppose now that the sentences in (46) carry F-marking that includes *some* in an F-marked constituent, but not the negation *not* or the antonymic predicate *inside*. For example, suppose for concreteness that both structures only feature narrow F-marking on *some*, as in (48).

(48)  a.  [Some]_F players have not responded.
      b.  [Some]_F players were inside.

Under Fox and Katzir's proposal, such F-marking ensures that that alternatives can differ from those sentences in the position of *some* but nowhere else. Therefore, while the sentences in (47a-i) and (47b-i) qualify as alternatives to (48a) and (48b), respectively, the simplex threats in (47a-ii) and (47b-ii) do not. With the simplex threats so eliminated by grammar, the strengthenings in (28a) and (33a) are correctly predicted to be available. That is, it is correctly predicted that (46a) can be understood as excluding (47a-i), to convey that some but not all players have responded, and (46b) as excluding (47b-i), to convey that some but not all players were inside.

However, Fox and Katzir's proposal also makes predictions about the sentences in (46) that are not in fact consistent with intuitions. The account is too permissive in some ways, and too restrictive in others. On the one hand, the proposal allows for the potential strengthenings in (28b) and (33b) to arise under conditions that support structures with appropriate F-marking, viz. structures where *not* and *inside* are included in F-marked constituents, while *some* is not. For example, (28b) and (33b) would be expected to arise from the structures in (49), with narrow focus on *not* and *inside*. Since there is in fact no reason to expect that such F-marking is categorically excluded, it is predicted that the strengthenings in (28b) and (33b) can be observed under certain conditions. However, as implied by our initial characterization of the data, intuitions are clear that these strengthenings are not in fact available under any conditions. In fact, even with a prosody favoring the F-marking in (48), the possibility of the implicature based on substituting *some* persists. Apparently, the assumption that scalar implicatures are constrained by limiting substitutions to F-marked content is incorrect.

(49)  a.  Some players have [not]_F responded.
      b.  Some players were [inside]_F.

On the other hand, Fox and Katzir's proposal also leads one to expect that the availability of the strengthenings in (28a) and (33a) depends on F-marking. Specifically, strengthening should be obviated by F-marking that leads to a symmetric stalemate. Under conditions that force both *some* and *not* to be included in an F-marked constituent in (46a), both (47a-i) and (47a-ii) would qualify as alternatives to (46a). Likewise, both (47b-i) and (47b-ii) would qualify as alternatives to (46b) under conditions that force both *some* and *inside* to be included in an F-marked constituent in (46b). Under such conditions, a stalemate should ensue, and hence we should not see either of the strengthenings in (28) and (33), respectively. It is predicted, for example, that neither strengthening is observed under conditions that force narrow focus on *some* and *not* or *inside*, as in (50), or that force broad focus on the entire sentence, as in (51).

(50)  a.  [Some]_F players have [not]_F responded.
      b.  [Some]_F players were [inside]_F.

(51)  a.  [Some players have not responded]_F.
      b.  [Some players were inside]_F.

This prediction too seems incompatible with intuitions. As hard as one may try to force the sort of F-marking shown in (50) and (51), it seems that the strengthening asymmetries described in (28) and (33) persist, with (28a) and (33a) continuing to be available. In particular, these asymmetries are still intuited when the sentences in (46) are produced with prosodic focus on *some* and *not* or *inside*, in an attempt to force the double narrow F-markings in (50). And, paralleling an observation in Romoli (2013), they also persist in responses to a wh-question like *What happened?*, which should favor the broad F-markings in (51).

In sum, the strategy of appealing to the Stalemate Assumption and F-marking to avoid the simplex threat fails on two counts: First, we showed direct counter-evidence to the proposal that the exclusion set is constrained by F: focus placement that should make undesirable alternatives available and desirable ones unavailable do not change which implicature is available in cases of symmetry. Second, we saw that focus marking that should make *both* alternatives available fails to lead to a stalemate. An alternative explanation is needed, and we propose that the blocking condition BC can fill this gap.

Our observations warrant a further conclusion. Much of the recent literature on the symmetry problem assumes the Stalemate Assumption and takes the goal of any successful solution to the symmetry problem to offer a way of evading a stalemate that would otherwise be expected. It is particularly notable therefore that in the absence of a formal way to exclude the simplex threats via assumptions about grammar, our data appear to be incompatible with Stalemate Assumption itself. What if the assumption is actually incorrect? What if what the symmetry problem actually calls for is a solution that avoids the otherwise expected symmetric availability of two opposite implicatures? Under this perspective, symmetry does not need to be 'broken' to allow an inference, rather, the task is simply to exclude unattested strengthenings.

There is in fact more direct empirical evidence against the Stalemate Assumption. Consider the following dialogue, based on a similar example in Breheny et al. (2018):

(52)    What do you do to protect your teeth?
        Kid 1: I brush my teeth. (And) I don't eat candy.
        Kid 2: I brush my teeth (too).

Here Kid 2 can be taken to imply that they eat candy—presumably through the denial of the alternative *I don't eat candy*. As noted in Breheny et al. (2018), this is problematic for CF, since the alternative *I eat candy* constitutes what we labeled a simplex threat. But we think that cases like (52) raise a more general problem for accounts that try to restrict the exclusion set: If we are still harboring hopes that a new, more sophisticated formal criterion will disqualify the problematic alternative in (52), the following dialogue should convince us otherwise:[7]

(53)    What do you do to protect your teeth?
        Kid 1: I brush my teeth. (But) I don't floss.
        Kid 2: I brush my teeth (too).

It seems to us that here Kid 2 can be taken to imply that they do not floss either. In fact, this by itself seems unremarkable and as it should be, since it can be derived simply by denying the

---

[7]See Trinh (2018), e.g., for a proposal that disallows dropping a negation in an alternative, which would remove the simplex threat in (52), but is in compatible with (53), which precisely requires dropping negation in a an alternative.

alternative *I floss*. But in combination with (52), it shows that a formal explanation is not just seem unlikely here, it is uncalled for. What tips the balance toward an implicature based on the negated or non-negated alternative is context: The question in (52) and (53) makes reporting *not eating candy* and *flossing* directly relevant and therefore omitting them meaningful—but the same is not true of their negations. And this is clearly not because of formal aspects of these alternatives, but because of what we know about how flossing and eating candy affect our teeth. Taken together, the two dialogues suggest that the Stalemate Assumption is incorrect.

Nothing about the solution in terms BC rides on the Stalemate Assumption, and BC has the further virtue of remaining silent on cases like (52) and (53): There is no simpler expression to convey the strengthened meaning, whether we base the implicature on the negated or non-negated alternative. Which one, if any, a speaker is more likely to have in mind is resolved by pragmatic reasoning. Note that if we abandon the Stalemate Assumption, we no longer need to worry about the possibility that both alternatives are in the exclusion set, since a pragmatic account, maybe based on the Gricean Maxim of Quality, will preempt contradictory utterances.[8]

But aren't there clear cases of 'de facto' stalemates between symmetric inferences? How can we explain those, if we abandon the Stalemate Assumption? Consider the following example:

(54)  a.  A:  I have a yellow bicycle.
      b.  B: # That's not true! You have a cheap bicycle.

Here B's objection seems contrived, but why? The sentence should have the following two alternatives, and—unless we make the Stalemate Assumption—shouldn't we now expect that an implicatures based on either of them is viable?

(55)  a.  Alternative 1: I have an expensive bicycle
      b.  Alternative 2: I have a cheap bicycle

Note first that B's response would be fine in a context in which (for some contrived reason) everyone knows that for a bike, being yellow implies being expensive. The existence of the formal lexical alternatives of equal complexity (here: *cheap* and *expensive*) does not automatically create a stalemate. This is incompatible with Stalemate Assumption, which predicts that (54) should be infelicitous in *any* context.

For sure, without a rich context (54) is quite odd. But appealing to a formal stalemate condition to explain this seems unnecessary. There is a plausible pragmatic reason for why A cannot be taken to court by B for having implied that their bicycle is not cheap. If A arbitrarily adds the alternative *expensive* to the exclusion set (and we see no reason why they should not be able to do so in principle), this is not reliably recoverable by B, and hence B is not entitled to object to the implications of this choice on the grounds of what A asserted. Of course A would not even try this, knowing that their message would not be recoverable, and resort to a different way to convey this meaning instead. Contexts that remove the arbitrariness from the choice of alternatives by justifying one inference over another remove the infelicity.

These examples illustrate that BC is not the only constraint on strengthening. Pragmatic reasoning constrains what is included in the exclusion set in a given context. This should not be

---

[8]So we take the observation that alternatives in the exclusion set need to be innocently excludable not as a constraint on Exh, but as a pragmatic effect.

surprising, given that conversational implicatures more generally are obviously context dependent. They also provide evidence against the Stalemate Assumption. If it is indeed incorrect, this further supports our argument that BC is needed to solve the symmetry problem. As we saw, this closes potential alternative avenues to explain simplex threats.

## 5. Conclusion

We have argued that the Blocking Constraint (BC), itself maybe derivable from the Gricean Manner maxim, provides a potential solution to the symmetry problem. We presented evidence that an alternative account in terms of a syntactic Complexity Filter (CF) on alternatives is questionable, since implicatures based on complex alternatives that CF should rule out are in fact often available. BC was furthermore motivated by cases where an account in terms of CF is impossible, since the problematic alternative are simplex. Finally, we have argued that the Stalemate Assumption, which plays a crucial role in supporting the CF account and could support alternative strategies to deal with simplex threats, is incorrect, as well as the assumption that the exclusion set is a subset of the alternatives relevant for focus.

Formal constraints such as CF assume that sets of alternatives, or at least those excluded by Exh, are sets of linguistic expressions. BC too is sensitive to the linguistic resources of a language, but it is an output condition on strengthening, rather than a constraint on the alternatives that strengthening is based on, and does not require an exclusion set with formal alternatives. We have only looked at a limited number of cases in this paper, but based on these it seems to us that the members of the exclusion set might in fact be completely free and not subject to formal or semantic constraints, although they must be recoverable in context and may be subject to other pragmatic effects.

We saw that BC correctly predicts that in cases of symmetry in which *neither* strengthened reading is expressible without added complexity, *both* implicatures can in principle be available. BC makes a further prediction: If in a case of symmetry *both* strengthened readings are expressible without added complexity, then *neither* implicature should be available, because each one should be blocked by its competitor. Swanson (2010) presents examples which he argues provide exactly this scenario: Two stronger but mutually exclusive strengthened meanings of an expression have lexicalised counterparts. Let's consider one of these examples:

(56)    Going to confession is permitted.

(57)    a.    Going to confession is optional.
        b.    Going to confession is required.

Sentence (56) can implicate (57a), but not (57b). It can convey that confession is not required, hence optional, but not that confession is not optional, hence required. Swanson uses this and related examples to show that CF cannot be the only constraint on Excl. In our terminology, *Going to confession is optional* poses a simplex threat. We add here the observation that the example is also a counterexample to the Stalemate Assumption, since the inference in (57b) is possible although it has a symmetric counterpart.

If Swanson is correct and the lexical meaning of *optional* lexically encodes *permitted and not required*, then (56) also falsifies BC: BC correctly predicts that the assertion with *require* blocks

the unavailable strengthening of *permit* to convey 'required;' but the assertion with *optional* should block the available strengthening that confession is *not* required. We contend, however, that *optional* does not actually have this meaning. If it did, both of the following statements should be equally infelicitous, but only the second is outright contradictory:[9]

(58)    a.   ?Doing drugs is optional, and in fact it is not permitted
          b.   #Doing drugs is optional, and in fact it is required

This suggests that the lexical meaning of *optional* is close in meaning to *not required*, and only conveys *permitted* in certain (possibly most) contexts as an implicature. One informant called (58a) 'cheeky,' indicating that it is hard not to draw the implicature. But there is a clear contrast to (58b). That an implicature is at play is also supported by the observation that the inference tends to vanish under negation, leading to the implausible reading that taking drugs is required:

(59)    a.   Taking drugs is not optional.
          b.   Taking drugs is non-optional

The relation between *optional* and *permitted* is then parallel to the relation between *some* on the one hand and *not all* (and *some not*) on the other: The strengthened versions of these are truth-conditionally equivalent, but neither basic meaning is equivalent to the strengthened reading of the other, and hence BC correctly predicts that each of them can be strengthened.

This diffuses the challenge for BC, but not for CF and the Stalemate Assumption: Even if the basic meaning of *optional* does not entail *permitted*, the alternative poses a simplex threat for CF. The case of *optional* shows that, contrary to what is sometimes assumed (starting with Horn 1972), expressions that play an equivalent role to *not all/some not* vis-á-vis their alternatives can be lexicalized. It does not, as we had worried, constitute a failure of blocking.

Blocking effects are generally quite stubborn. We saw that *some* cannot convey *all* even if we use focus to make alternatives salient that should in principle support this implicature. And yet Bonnefon et al. (2009) point out cases that do look like true failures of blocking:

(60)    a.   A: What impression did I make during dinner?
          b.   B: Some thought you drank too much.

(61)    a.   Some but not all thought that you drank too much.
          b.   Everyone thought you drank too much.

We may well infer here that B believes that everyone thought that A drank too much. This could be derived by exclusion of the alternative in (61a). This is of course the very strengthening both Blocking and CF set out to prevent. The example seems to falsify both. We think think there is a way, however, to make sense of this under the Blocking account: Bonnefon et al. (2009) plausibly argue the choice of *some* is a 'politeness strategy,' used to avoid hurting the feelings of the addressee. Under the blocking account we can make sense of this if politeness and other social constraints can take certain utterances out of the picture, preempting otherwise expected blocking effects. This could open the door for *some* to convey *everyone* in (60).[10]

---

[9]The following example found online supports this point:

(i)      Breaking the law is optional and the consequential punishment is known so what's the problem?

[10]Portner (2024), however, argues that this inference is not due to the same mechanism that derives scalar implicatures (see also Swanson 2017). We defer a more thorough discussion to another occasion.

# References

Bonnefon, J.-F., A. Feeney, and G. Villejoubert (2009). When some is actually all: Scalar inferences in face-threatening contexts. *Cognition 112*, 249–258.

Breheny, R., N. Klinedinst, J. Romoli, and Y. Sudo (2018). The symmetry problem: current theories and prospects. *Natural Language Semantics 26*, 85–110.

Buccola, B., M. Križ, and E. Chemla (2022). Conceptual alternatives: Competition in language and beyond. *Linguistics & Philosophy 45*, 265–291.

Chierchia, G., D. Fox, and B. Spector (2011). The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In C. M. Paul Portner and K. von Heusinger (Eds.), *Handbook of Semantics*. de Gruyter, New York.

Fox, D. (2007). Free choice and the theory of scalar implicatures. In U. . P. S. U. Sauerland (Ed.), *Presupposition and Implicature in Compositional Semantics*, pp. 537—586. New York: Palgrave Macmillan.

Fox, D. and R. Katzir (2011). On the characterization of alternatives. *Natural Language Semantics 19*(1), 87–107.

Fox, D. and B. Spector (2018). Economy and embedded exhaustification. *Natural Language Semantics 26*(1), 1–50.

Geurts, B. (2011). *Quantity Implicatures*. Cambridge: Cambridge University Press.

Hirsch, A. (2024). Constraining alternatives. Invited talk at the 25th NY-St. Petersburg Institute of Linguistics, Cognition, and Culture; January 15, 2024 (joint work with Bernhard Schwarz).

Hirsch, A. and B. Schwarz (2023). Type disambiguation and logical strength. In S.-Y. Lam and S. Ozaki (Eds.), *Proceedings of 53d annual meeting of the North East Linguistic Society (NELS 53)*. GLSA (Graduate Linguistics Student Association), Department of Linguistics, University of Massachusetts.

Horn, L. (1972). *On the semantic properties of logical operators in English*. Ph. D. thesis, University of California, Los Angeles.

Horn, L. R. (2000). From if to iff: Conditional perfection as pragmatic strengthening. *Journal of Pragmatics 32*(3), 289–326.

Katzir, R. (2007). Structurally-defined alternatives. *Linguistics and Philosophy 30*(6), 669–690.

Kroch, A. (1972). Lexical and inferred meanings for some time adverbs. *Quarterly progress reports of the research laboratory of electronics 104*, 260–267.

Levinson, S. (1983). *Pragmatics*. Cambridge University Press.

Matsumoto, Y. (1995). The conversational condition on horn scales. *Linguistics and Philosophy 18*(1), 21–60.

Portner, P. (2024). Social relations in semantics and pragmatics. Colloquium talk at McGill University.

Romoli, J. (2013). A problem for the structural characterization of alternatives. *Snippets 27*.

Swanson, E. (2010). Structurally defined alternatives and lexicalizations of XOR. *Linguistics and Philosophy 33*(1), 31–36.

Swanson, E. (2017). Omissive implicature. *philosophical topics 45*(2), 117–138.

Trinh, T. (2018). Keeping it simple. *Natural Language Semantics 26*, 111–124.

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