

Constraining alternatives¹

Aron HIRSCH — *University of Maryland*

Bernhard SCHWARZ — *McGill University*

Abstract. The focus operator *only* asserts that alternatives to its prejacent are false. But, what alternatives does *only* exclude? There must be *constraints* on alternatives. Fox and Katzir (2011) propose that alternatives are formed in the syntax, and are constrained by grammar such that they cannot be structurally more complex than the prejacent. We argue that this constraint is too restrictive. We take a new look at complex alternatives, and conclude that they can be attested, with their distribution regulated pragmatically by the Question Under Discussion (after Roberts 1996/2012, Beaver & Clark 2008, Katzir 2023).

Keywords: alternatives, *only*, the symmetry problem, Question Under Discussion

1. Introduction

As defined in (1), *only* applies to a proposition, *p*, presupposes that *p* is true, and asserts that alternatives to *p* are false. In this paper, our concern is with the alternative set (ALT). We will consider how the make-up of the alternative set is determined.

(1) Defining *only*

$$\llbracket \text{only} \rrbracket^{\text{ALT}} = \lambda p_{\text{st}} . \lambda w : p(w) . \forall p' \in \text{ALT} [p \not\subseteq p' \rightarrow \neg p'(w)]$$

The challenge is that there must be *constraints* on what propositions are contained in ALT. In this paper, we will focus on one of these constraints, and consider how it is introduced. Is it imposed by the grammar? Or, does it arise pragmatically through interaction with the discourse context? The issue can be introduced with the *symmetry problem*.

1.1. The symmetry problem

Rooth (1985, 1992) proposed that alternatives are formed by replacing the focus with any element of the same semantic type. Yet, by itself, this algorithm predicts *too many* alternatives. To illustrate, we will consider (2), with the LF in (3).

(2) Al only [cooked]_F.

(3) [_{TP} only [_{VP} Al [cook]_F]]

The VP is focused, and denotes a property of type $\langle e, \text{st} \rangle$. Accordingly, by Rooth's algorithm, alternatives can be formed by replacing the denotation of *cook* with any other property of that type. One is the property expressed by *clean*. Replacing COOK with CLEAN yields (4): that Al

¹For discussion, we thank Luis Alonso-Ovalle, Sarah Boukendour, Danny Fox, Nina Haslinger, Masha Polinsky, Paolo Santorio, Viola Schmitt, Michael Wagner, and Alexander Williams. We have presented portions of this project at NELS 53 and SuB 29, as well as in talks at the NYI Winter School, Yale University, the University of Stuttgart, and UMD. Schwarz gratefully acknowledges financial support from the Social Sciences and Humanities Research Council of Canada (Insight Grant 435-2019-0143).

cleaned. If this alternative is in ALT, *only* would negate it to yield the entailment in (5): that Al did *not* clean. That is an intuitive entailment of the sentence in (3).

- | | |
|--|---|
| (4) Positive alternative
CLEAN(Al) | (5) Entailment
\neg CLEAN(Al) |
|--|---|

However, there are replacements which would result in unattested alternatives. Most crucially, consider the alternative in (6). In addition to the property expressed by *clean*, there is the one expressed by *not clean*. If COOK were replaced by \neg CLEAN, the alternative in (6) results: that Al did not clean. If that were in ALT, it would be negated to yield an entailment that Al actually *did* clean. That entailment is not intuited, and contradicts (5).

- | | |
|---|------------------------------------|
| (6) Negative alternative
\neg CLEAN(Al) | (7) Entailment
CLEAN(Al) |
|---|------------------------------------|

The alternatives in (4) and (6) can be thought of as ‘symmetric’ in the sense that one is the negation of the other, so jointly excluding both of them would result in contradictory inferences. The alternative set must be constrained so that it contains the positive alternative, but not the negative one. The target ALT is flagged in (8).

- (8) **Target ALT**
ALT = { CLEAN(Al), \neg CLEAN(Al), ... }

The question is, then: *how* is ALT constrained? Since Rooth (1992), it has been recognized that two factors play a role in determining ALT: grammar and context. The recipe is schematized in (9). The grammar generates a set of ‘formal alternatives’, indicated as FA. C is a set of alternatives provided by the context. ALT is the intersection of the two sets. That is, ALT is a contextually determined subset of the formal alternatives.

- (9) **Recipe for ALT**
ALT = FA \cap C

In positing constraints on alternatives, then, there are two possibilities: constraints could come from either grammar or context. In addressing the symmetry problem, the traditional approach has been to posit constraints on the grammatical computation.

1.2. Structural complexity

An influential proposal for how to break symmetry in the grammar is due to Fox & Katzir (2011), which builds on earlier work on scalar implicatures by Katzir (2007). Their proposal is that formal alternatives are constrained based on their syntax. In particular, alternatives are computed as syntactic objects by replacing the focused constituent with constituents of *equal* or *lesser* structural complexity. Consider, again, the structure for (2).

- (2) **LF for (2)**
[_{TP} only [_{VP} Al [cook]_F]]

Constraining alternatives

In syntactic terms, *cook* is a simplex VP. The positive alternative would have the syntax in (10), where *cook* is replaced by *clean*, which is equally complex. In the negative alternative, by contrast, *cook* would be replaced by *not clean*, and that is more complex, due to the negation. As such, just the positive alternative is generated in FA, and so only it can be in ALT. Symmetry is broken in the grammar in favor of the positive alternative, as attested.²

(10) **Positive alternative**

[_{VP} AI [clean]]

(11) **Negative alternative**

[_{VP} AI [not [clean]]]

Fox & Katzir's (2011) proposal, based on structural complexity, offers a simple and principled solution to the symmetry problem. However, our aim is to build a case that focus alternatives are *not* restricted by structural complexity, despite initial appearances.

1.3. Plan for the paper

We will begin the paper by identifying cases where the alternative set for *only does* contain alternatives which, in syntactic terms, would be more complex than the prejacent. These data will lead to a nuanced empirical picture where complex alternatives are available *sometimes*, but not *always*. To account for the full data set, our tactic will be to shift the burden in constraining complex alternatives from the *grammar* to the *context*.³

We will propose that formal alternatives are computed in the grammar irrespective of their structural complexity. The presence of complex alternatives in ALT is, then, determined by a feature of the discourse, in particular the Question Under Discussion (the 'QUD'). The idea that the QUD plays a role in constraining the alternative set for *only* is not novel, and has been pursued in Roberts (1996/2012), Beaver & Clark (2008), and Katzir (2023), among other works. Our contribution in the present paper is to argue that the QUD should account for data previously attributed to a grammatical restriction on complexity.

The discussion will proceed as follows. In Section 2, we will present challenges to Fox & Katzir's constraint. In Section 3, we will show how the QUD analysis can be applied to break symmetry. In Sections 4 and 5, we will show that the QUD analysis also predicts those cases where more complex alternatives *are* included in ALT. Section 6 concludes.

2. Challenging complexity

The first step is to put forward challenges to Fox & Katzir's proposal that complex alternatives are banned in the grammar. To achieve that, we will bring out cases where ALT does include alternatives which would be structurally more complex than the prejacent. Here, we will focus on two cases, with the central results as previewed in (12).

²Fox & Katzir (2011) allow that a complex alternative can be present in FA, but only if that alternative is explicitly provided as a constituent in the linguistic context, and does not result in FA containing symmetric alternatives. No data considered in the present paper will meet both of these conditions.

³We will use the term 'complex' alternatives in a theory neutral way to refer to alternatives which would be complex if expressed as a syntactic object. As discussed later in the paper, the approach that we pursue does not depend on alternatives being formed as syntactic objects. Similarly, we will refer to propositions or properties as 'positive' or 'negative' based on whether their syntactic expression would contain negation.

- (12) a. Alternatives can contain disjunction. (*Case 1*)
 b. Alternatives can contain negation. (*Case 2*)

In Case 1, we will establish that ALT can include an alternative which is more complex because it contains coordination, in particular *disjunction*. In Case 2, we show that alternatives can even contain negation itself in certain environments. By identifying challenges for Fox & Katzir’s constraint, we join with several recent works on scalar implicatures, in particular Trinh & Haida (2015), Breheny et al. (2018), and Schwarz & Wagner (2024).

2.1. Case 1: disjunction

Our target in Case 1 is to establish that ALT for *only* can include a complex alternative which contains *disjunction*. In order to bring out the intuition, it will be useful to introduce a modal into the test example. We will work with the example in (13).

- (13) Al only has to [cook]_F.

To start, we will consider the meaning that would be predicted for (13) if it were analyzed in accord with the structural complexity constraint. We consider *two* alternatives allowed by the constraint. As in the earlier example in (2), focus occurs on the simplex VP *cook*. To derive (14a), *cook* would be replaced by *clean*, and to derive (14b), the replacement is *shop*. Hence, the alternatives say that Al has to clean, and that he has to shop, respectively.

- (14) **Alternatives (by complexity)**
 a. \square CLEAN(Al)
 b. \square SHOP(Al)

These alternatives should each be in ALT, and *only* will negate them to deliver the entailments in (15): that Al does *not* have to clean, and that he does *not* have to shop. These are indeed intuitive entailments of the sentence. However, our aim is to show that the predicted meaning is still *too weak*. The derived entailments are not enough.

- (15) **Entailments**
 a. $\neg\square$ CLEAN(Al)
 b. $\neg\square$ SHOP(Al)

The crucial scenario is one where there is a *disjunctive* requirement placed on Al. Suppose that the facts at the world (w_1) are as in (16). Al has to cook, and he also has to do one other chore, but he can choose which one to do: Al can *either* clean *or* shop, and it’s up to him.

- (16) **Facts at w_1**
 a. Al has to cook.
 b. Al has to clean or shop (his choice).

Intuitively, the sentence in (13) is *false* in this scenario. We can’t say that Al only has to cook if he also has to clean or shop. But, strikingly, the prediction so far is that (13) should be *true*, since both of the entailments in (15) are verified in this scenario. Because Al can choose the

Constraining alternatives

chore, it's true that he does not *have* to clean, and true that he does not *have* to shop. If (15a) and (15b) were the sole entailments of the sentence, the predicted value is true.

To capture the false judgment, there must be additional alternatives for *only* to exclude. That way, further entailments will be derived, one of which must *not* be verified at w_1 . The crucial alternative is the one in (17), which contains disjunction. This alternative says that AI has to clean or shop, and if excluded, the entailment in (18) would result: that AI does *not* have to clean or shop. That entailment would fail at the world in (16), where there is that disjunctive requirement placed on AI. Hence, if (17) is an alternative, falsity is captured.

(17) **Target alternative**

$\Box [\text{CLEAN}(\text{AI}) \vee \text{SHOP}(\text{AI})]$

(18) **Target entailment**

$\neg\Box [\text{CLEAN}(\text{AI}) \vee \text{SHOP}(\text{AI})]$

Now, the key result is in place to bring out a challenge for a structural complexity constraint. In syntactic terms, the alternative in (17) is more complex than the original prejacent. The LF for (13) is (19). Recall that focus occurs on the simplex VP *cook*. To derive the target alternative syntactically, it would be necessary to replace *cook* with *clean or shop* as in (20). However, due to the coordinate structure, (20) is more complex than (19).

(19) **LF for (13)**

$[_{\text{TP}} \text{ only } [\text{ have } [\text{ AI } [\text{ cook }]_{\text{F}}]]]$

(20) **Alternative (syntax)**

$[\text{ have } [\text{ AI } [\text{ clean or shop }]]]]$

Hence, Fox & Katzir's approach faces an *under-generation* problem. The target alternative should not be in the set of formal alternatives, and thus could not be in ALT. Yet, the alternative is attested. The constraint goes too far in limiting alternatives.

To this point, we have established that one kind of more complex alternative can be in ALT, in particular a disjunctive alternative. Yet, as we discussed at the outset, the central aim of Fox & Katzir's proposal is to solve the symmetry problem, which requires eliminating a complex alternative that contains *negation*. In Case 2, our aim is to establish that the alternatives for *only* can even contain negation itself in certain environments.

2.2. Case 2: negation

The sentence in (21) is adapted from Schwarz & Wagner (2024), and has a crucial feature: that focus occurs on the *propositional complement* of an intensional verb, here the speech report predicate *say*. In this case, we will see that entailments are intuited which diagnose the presence in ALT of an alternative containing negation.

(21) Mary only said that [it rained]_F.

Suppose that the speaker of (21) has gone up to Mary to ask what she knows about the weather, and now has returned to the group. The speaker utters (21). Intuitively, (21) conveys that Mary said that it rained, but did not say *any* other relevant proposition, regardless of polarity. For example, (21) conveys both (22a) and (22b). We can infer that Mary didn't say that it snowed, and also that she didn't say that it didn't snow.

(22) **Entailments**

- a. \neg SAY(Mary, SNOW)
- b. \neg SAY(Mary, \neg SNOW)

To derive the observed inferences, it must be that alternatives are formed not only with positive replacements, but also with replacements containing negation. In the sentence in (21), the F-marked constituent is *it rained*. To derive (23), the replacement is *it snowed*, so the alternative says that Mary said that it snowed. To derive (24), the replacement is *it didn't snow*, so the alternative says that Mary said that it didn't snow, with negation embedded under *say*. If both are in ALT, they will each be negated, deriving the inferences in (22).

(23) **Positive alternative**

SAY(Mary, SNOW)

(24) **Negative alternative**

SAY(Mary, \neg SNOW)

The presence of the negative alternative in ALT poses a further challenge for Fox & Katzir's constraint. In syntactic terms, the sentence has the LF in (25), where focus is shown on the embedded TP. If the negative alternative is formed syntactically, it would have the LF in (26), which is more complex, due to the additional negation head present in (26). The negative alternative should *not* be in the set of formal alternatives, contrary to fact.

(25) **LF for (21)**

[_{TP} only [Mary said that [_{TP} it rained]_F]]

(26) **Alternative in (24) (syntax)**

[Mary said that [_{TP} it did not snow]]

Just like in Case 1, the results in Case 2 establish that the structural complexity constraint put forward by Fox & Katzir is *too* restrictive. It fails to allow for an alternative which is in fact attested, here an alternative with negation in an embedded environment.

2.3. Next step

We are ready to re-evaluate how complex alternatives are constrained. Consider, now, the revised empirical picture. At the outset of the paper, we considered the basic symmetry case, where a complex negative alternative is unattested. Yet, Cases 1-2 have shown that complex alternatives can arise, and can even contain negation itself. Hence, we conclude that complex alternatives are available not *always* — but still *sometimes*.

The overall data set is at odds with Fox & Katzir's constraint by which complex alternatives are disallowed by the grammar. Because complex alternatives can be attested, they must be generated as possible alternatives. Accordingly, we pursue a shift in perspective where complex alternatives are instead constrained by context. Consider again (27).

(27) **Recipe for ALT**

ALT = FA \cap C

As discussed at the outset, ALT reflects the intersection of two sets: FA (from grammar) and C

Constraining alternatives

(from context). We propose that alternatives in FA are generated irrespective of their structural complexity. One option would be to return to Rooth's (1985) analysis, where alternatives are computed based on semantic type. If alternatives are derived in the semantic component, then syntax cannot play a role. Complex alternatives must then be regulated by C.

We will pursue a pragmatic account in two steps. First, in Section 3, we will return to the basic symmetry example, where the negative alternative was unattested, and predict its absence based on pragmatic considerations. Then, in Section 4, we will return to the data from Cases 1-2, and show that attested complex alternatives are predicted in the same line of analysis.

3. Factoring in the question

In pursuing a pragmatic approach to constraining alternatives, we appeal to the Question Under Discussion. Roberts (1996/2012) proposed that discourse is always organized into question-answer pairs, as in (28). The QUD is explicit in some cases, but is often left implicit. As in Hamblin (1973), we take a question to denote the set of its possible answers.

(28) (QUD)
[_{TP} ... only ...]

When it comes to compute the alternative set for *only*, we take it that the C component is anaphoric to the Hamblin set for the QUD, as pursued extensively in Beaver & Clark (2008) (see also Roberts 1996/2012, Katzir 2023). That is, the alternatives in C correspond to the possible answers to the contextual question.

Appeal to the QUD has a key consequence: to predict the alternative set for *only*, it is necessary to determine the question. When no question is explicit, we must determine the *implicit* QUD that would be accommodated. Our aim, then, is to capture where complex alternatives do and do not occur, based on plausible constraints on an accommodated QUD.

3.1. Breaking symmetry

We return to the basic symmetry case in (2), where alternatives must be constrained. The problem, recall, is that (2) licenses the inference in (29a), but not (29b). Since *only* excludes alternatives, the target alternative set is (30). ALT must contain the positive alternative (that Al cleaned), but not its negative counterpart (that Al did not clean).

(2) Al only [cooked]_F.

(29) **Entailments**

- a. $\Rightarrow \neg \text{CLEAN}(\text{Al})$
- b. $\not\Rightarrow \text{CLEAN}(\text{Al})$

(30) **Target ALT**

ALT = { CLEAN(Al), ~~¬CLEAN(Al)~~, ... }

In our approach, alternatives are computed in the grammatical component without regard to structural complexity, so the set of *formal* alternatives would be as in (31). Both the positive alternative, shown in (31b), and the negative alternative, in (31c), are contained in the formal alternative set. However, we suggest that the implicit QUD that would get accommodated by

default supplies only the positive alternative, breaking symmetry.

(31) **Formal alternatives**

FA = { (31a), (31b), (31c), ... }

- a. COOK(AI)
- b. CLEAN(AI)
- c. ¬CLEAN(AI)

A natural way to form an implicit QUD would be to simply replace the focused constituent in the sentence with a simplex *wh*-phrase (as in Büring 2003). That would result in the implicit discourse in (32), where the question in (32a) serves as the QUD.

(32) **Target QUD**

- a. What did AI do?
- b. AI only [cooked]_F.

The *wh*-phrase in (32a) ranges over properties. But, when the question occurs out of the blue, there is an intuition that *what* is further restricted by default to range only over properties which express *positive actions*. That is, (32a) could be paraphrased as ‘What action did AI take?’. By default, then, the domain of *what* is as in (33). In turn, the Hamblin set for the question would contain only positive answers, as made explicit in (34).

(33) **Domain of *what***

{ COOK, CLEAN, ... }

(34) **Hamblin set**

{ (34a), (34b), ... }

- a. COOK(AI)
- b. CLEAN(AI)

The alternative set for *only* will, then, be formed as in (35). While FA contains both positive and negative alternatives, only positive alternatives are in C, as determined by the QUD. Therefore, only positive alternatives will be contained in ALT. ALT contains the positive alternative (that AI cleaned), but not the negative one (that AI did not clean). Symmetry is broken in the intended direction — not in grammar, but in context, due to the QUD.

(35) **Computing ALT**

ALT = FA ∩ C = { COOK(AI), CLEAN(AI), ... }

Now, in principle, there are other possible questions that could be accommodated which would incorrectly supply the *negative* alternative. Yet, we suggest that such questions are ruled out by different constraints. One question to consider as a potential QUD would be (36). In this case, the question contains negation inside the question nucleus. The negation would result in each answer being negative. The Hamblin set would be (37).

Constraining alternatives

(36) **Hypothetical QUD**

What didn't Al do?

(37) **Hamblin set**

{ (37a), (37b), ... }

a. \neg COOK(Al)

b. \neg CLEAN(Al)

If (36) were the QUD, then the alternative set for *only* in (2) would be as in (38). C and, therefore, ALT contains just the negative alternative that Al didn't clean, rather than its positive counterpart. Symmetry is broken, but it is broken in the incorrect direction.

(38) **Computing ALT**

$ALT = FA \cap C = \{ \neg$ COOK(Al), $\boxed{\neg$ CLEAN(Al)}, ... }

However, there are, in fact, general constraints on questions and answers which predict that (36) *cannot* be the QUD. The crucial constraint is the simple requirement of question-answer *congruence*. For our purposes, we can formulate the constraint as (39), based on Rooth (1992). By (39), when a sentence contains *only*, the prejacent of *only* must itself be an answer to the question for the sentence to be congruent to that question.

(39) **Congruence constraint**

[only α] is felicitous, given a QUD Q, only if $\llbracket \alpha \rrbracket \in Q$.

Given the Congruence Constraint, the only way that a question could potentially supply the negative alternative is if it had a Hamblin set such as (40). In (40), the answers are mixed in polarity. There is the positive answer in (40a) (that Al cooked), as well as the negative answer in (40b) (that Al didn't clean). The target sentence would be congruent to the question, given (40a), but the negative alternative would still be supplied to *only*, given (40b).

(40) **Hamblin set (hypothetical)**

{ (40a), (40b), ... }

a. COOK(Al)

b. \neg CLEAN(Al)

We suggest, however, that a question such as (40) would not be accommodated by default as an implicit QUD. To express (40) in an out of the blue context, it would be necessary for the question to be as in (41), where the *wh*-phrase is not just *what*, but rather contains additional syntactic material to restrict the domain of *what* to the properties *cook* and *not clean*. We take it that the accommodation procedure is subject to a constraint such as (42), by which an implicit QUD can preferentially be expressed with a simple *wh*-phrase. Given (42), this constraint would eliminate an accommodated QUD such as (40) from consideration.

(41) **Potential QUD**

?Which among cooking and not cleaning did Al do?

(42) **Simplicity constraint**

An implicit QUD should be expressible with a simple *wh*-phrase (if possible).

In light of the constraints in (39) and (42), the expectation is that the implicit QUD *must* be formed by replacing the focus with a simple *wh*-phrase while maintaining congruence. For (2), the only implicit discourse that would be considered is as in (32). As we saw, the question in (32a) supplies positive answers, breaking symmetry as observed.

(32) **Predicted QUD**

- a. What did AI do?
- b. AI only [cooked]_F.

At this point, we have in place an analysis where complex alternatives are freely generated in the grammar, but can be filtered by the QUD. In the basic symmetry case, the QUD constrains the alternative set so that just the positive alternative is present. Hence, the QUD approach can block a complex negative alternative when it is *not* attested. In the next section, we show that the QUD approach also predicts where complex alternatives *do* occur.

4. The key prediction

To re-iterate, in Fox & Katzir's (2011) approach, complex alternatives are not generated in the grammar. By contrast, we have proposed that complex alternatives are generated, and then regulated by the context. We make the key prediction in (43).

(43) **Key prediction**

Complex alternatives are available when supplied by the QUD.

With (43) in hand, we return to the data from Cases 1-2 in Section 2, where certain complex alternatives *are* attested. Our aim is to establish that those alternatives are predicted to arise, given the expected implicit QUD. That is, some questions do make available certain kinds of complex alternatives, and those are the ones that we observed earlier in the paper.

4.1. Predicting Case 1

The example in (13) is repeated from Case 1. In Case 1, we established that (44) is an alternative for *only*. This is the disjunctive alternative that AI has to clean *or* shop. (44) is generated by the grammar in our approach. Because the grammar forms alternatives without regard to their structural complexity, COOK can be replaced by CLEAN \cup SHOP to form (44). Moreover, we will establish that (44) is supplied by the QUD, as well.

(13) AI only has to [cook]_F.

(44) **Target alternative**

\square [CLEAN(AI) \vee SHOP(AI)]

As before, we take the QUD to be formed by replacing the focused constituent with a simple *wh*-phrase while maintaining congruence. The expected implicit discourse is, then, (45). As before, *what* in (45a) ranges over properties and, intuitively, is restricted by default to range over positive actions. A partial Hamblin set is provided in (45). The proposition that AI has to cook is in the set, satisfying the congruence requirement.

Constraining alternatives

(45) **Implicit QUD**

- a. What does Al have to do?
- b. Al only has to [cook]_F.

(46) **Hamblin set (partial)**

- { (46a), (46b), (46c), ... }
- a. \square COOK(Al)
 - b. \square CLEAN(Al)
 - c. \square SHOP(Al)

The crucial point to consider now is whether the target alternative is in the set. So far, we have assumed that *what* ranges over properties conveying simple positive actions like COOK, CLEAN, or SHOP. In that case, the disjunctive alternative would *not* be a Hamblin answer. However, we want to suggest that *what* not only ranges over such properties, but also the disjunctions thereof. In that case, the *wh*-domain is revised as in (47), and in turn, the Hamblin set for the question expands from (46) as shown in (48).

(47) **Domain of *what***

{ COOK, CLEAN, SHOP, CLEAN \cup SHOP, ... }

(48) **Hamblin set (revised)**

- { (48a), (48b), (48c), ... }
- a. \square COOK(Al)
 - b. \square CLEAN(Al)
 - c. \square SHOP(Al)
 - d. \square [CLEAN(Al) \vee SHOP(Al)]

If (48) is the correct Hamblin set, the disjunctive alternative is supplied by the question. The Hamblin set contains not only the propositions that Al has to clean and that Al has to shop, but also the proposition that Al has to clean *or* shop. (48d) is the target proposition. It is necessary, however, to independently verify that (48d) is in the Hamblin set.

In general, one way to probe into the Hamblin set for a question is to consider intuitions when the question is *embedded*. We will consider, in particular, data where the embedding predicate is *know*. If an individual knows a question, it can be inferred that the individual know at least the information given by all true answers in the Hamblin set for the question. This is shown schematically in (49). Given (49), inferences drawn under embedding can offer a window into what propositions are contained in the Hamblin set.

(49) **Embedding test**

x knows Q (at w) $\Rightarrow \forall p \in Q$ [$p(w) \rightarrow$ KNOW(x , p , w)]

To that end, we consider the example in (50), where the question in (45a) is embedded under *know*. In order to assess the inferences which result, we will suppose, as in Case 1, that the facts at the world are as in (51), with a disjunctive requirement.

(50) Granny knows what Al has to do.

(51) **Facts at w_1**

- a. Al has to cook.
- b. Al has to clean or shop.

Intuitively, for the sentence in (50) to be true at the world in (51), it must be that Granny knows both (51a) and (51b). She must know that Al has to cook, but also the disjunctive information

that he has to clean or shop. (52b) is entailed, as well as (52a).

(52) **Entailments (at w_1)**

- a. \Rightarrow KNOW(Granny, \Box COOK(Al))
- b. \Rightarrow KNOW(Granny, \Box CLEAN(Al) \vee SHOP(Al))

The intuition can be brought out based on continuations. If either one or the other entailment fails, it is natural to object to the sentence on that basis. In (53a), the continuation objects to (52a), and in (53b) the continuation objects to (52b). Depending on Granny’s knowledge state, either continuation can be equally felicitous in the provided scenario. Hence, we conclude that inferences *are* generated about the answer containing disjunction when the question is embedded, confirming that it is in the Hamblin set for the question.

(53) A: Granny knows what Al has to do.

- a. B: No, she doesn’t know that Al has to cook.
- b. B: No, she doesn’t know that Al has to clean or shop.

With the Hamblin set for the question in place, the data with *only* from Case 1 follow for free. Given the implicit discourse in (45) above, the target disjunctive alternative is not only generated by the grammar, but also supplied as an answer to the QUD. Accordingly, it will be a formal alternative in FA and in C — and thus in ALT, as observed. We now turn to Case 2 and show that the complex alternative observed in those data is predicted in kind.

4.2. Predicting Case 2

In Case 2, we considered (21), which has the crucial feature that focus occurs on an *embedded proposition*. The target alternative is (54), where the replacement includes negation. That Mary said it rained has as an alternative that she she said it *didn’t* snow. Again, it is unproblematic to derive that alternative in the grammar in our view, despite its greater complexity. The question becomes whether it is supplied by the QUD.

(21) Mary only said that [it rained]_F.

(54) **Target alternative**
SAY(Mary, \neg SNOW)

In this case, the implicit QUD would be as in (55), derived in the familiar way. The focus is replaced by a simplex *wh*-phrase. As we will see, this question does supply the target alternative as a possible answer in its Hamblin set.

(55) **Implicit QUD**

- a. What did Mary say?
- b. She only said that [it rained]_F.

In (55a), *what* does not range over *properties*, as it did in the earlier data, but rather over *propositions*. Moreover, the intuition is that, by default, *what* ranges over *all* propositions, regardless of their polarity. The question in (55a) intuitively solicits *any* proposition that Mary said, positive or negative. By that intuition, the Hamblin set would be as in (56). The target

Constraining alternatives

alternative, containing negation, corresponds to the answer in (56c).

(56) Hamblin set

- { (56a), (56b), (56c), ... }
- a. SAY(Mary, RAIN)
- b. SAY(Mary, SNOW)
- c. SAY(Mary, ¬SNOW)

To verify that (56c) is in the Hamblin set we can, again, probe for its presence by drawing on intuitions under embedding. The test sentence is (57), with the question in (55) embedded under *know*. The scenario to consider is (58). In addition to saying the proposition that it rained, Mary did also say the negative proposition that it didn't snow.

(57) Al knows what Mary said.

(58) **Facts at w_2**

- a. Mary said that it rained.
- b. Mary said that it didn't snow.

Now, in order for the sentence to be true in this scenario, it must be that Al knows *both* of the propositions that Mary said, the negative one as well as the positive one. The two entailments in (59) are each observed, and if Al fails to have one or the other piece of knowledge, it is felicitous to object to the sentence, as in (60). With embedding, inferences *are* drawn about (56c), and thus, its presence in the Hamblin set is confirmed.

(59) Entailments (at w_2)

- a. \Rightarrow KNOW(Al, SAY(Mary, RAIN))
- b. \Rightarrow KNOW(Al, SAY(Mary, ¬SNOW))

(60) A: Al knows what Mary said.

- a. B: No, he doesn't know that Mary said it rained.
- b. B: No, he doesn't know that Mary said it didn't snow.

The data with *only* in Case 2, then, are correctly predicted. The expected implicit discourse is (55) above, where the target alternative is supplied by the QUD, and thus in ALT. Hence, Cases 1 and 2 have led to convergent results. In each case, the expected implicit QUD provides a particular kind of more complex alternative, and strikingly, that is the alternative that we showed to be attested in ALT earlier in the paper.

4.3. Taking stock

We have pursued a perspective where complex alternatives are generated in FA, and constrained by C. In doing so, we have captured the full set of data presented so far in the paper based on the view that C is anaphoric to an implicit QUD, accommodated in accord with plausible constraints. In the basic symmetry example, the default QUD only supplies positive answers, breaking symmetry in favor of a positive alternative over a complex negative one. But, at the same time, complex alternatives are able to arise. These include a disjunctive alternative in Case 1, and an alternative with embedded negation in Case 2.

To this point, we have presented data without an explicit context. As such, we have had to rely on accommodation of an *implicit* QUD in accounting for the distribution of alternatives. Yet, we can subject the QUD approach to further testing by providing an *explicit* context which manipulates the QUD overtly. That is our aim in the next section.

5. A further test

A general prediction emerges from the QUD analysis. Suppose that a sentence with *only* comes with an explicit QUD. Then, that QUD should provide the alternatives for *only*, regardless of what question would get accommodated by default in an out-of-the-blue utterance. As such, if the explicit QUD supplies complex alternatives, those should be in ALT.

To assess the prediction, we will work further with symmetry data. Consider, again, (2). As we've seen throughout the paper, out of the blue, the attested alternative set is (61). It contains the positive alternative that Al cleaned, but not its negative counterpart.

(2) Al only [cooked]_F.

(61) **Target ALT**

ALT = { CLEAN(Al), =CLEAN(Al) }

To account for ALT, we have made appeal to an implicit question. We suggested that the default QUD that would be accommodated out of the blue supplies only the positive alternative. But, with an explicit QUD, it should be possible to reverse the pattern. Suppose there were an overt question, which supplies the negative alternative, instead of the positive one. Then, just the *negative* alternative should be in ALT. Our aim is to show that manipulating the explicit QUD does impact the direction of symmetry breaking, as predicted.

5.1. Overt QUD

To construct a natural example, we will use a different sentence with the same logical profile as (2). The scenario is provided in (62). Our test case is modified from Trinh & Haida (2015) (see also Breheny et al. 2018, Schwarz & Wagner 2024).

(62) (Al's doctor told him to sleep and not run today.)

a. But, which of these orders did he actually obey?

b. Al only [slept]_F.

In this context, the set up makes salient two propositions, one positive (that Al sleep), and one negative (that Al not run). Crucially, these propositions are not just salient. Rather, they are the Hamblin answers to an explicit QUD. The question in (62a) includes additional syntactic material in the *wh*-phrase, which contains an anaphor to the two propositions (*these orders*). Hence, the Hamblin set for the question in (62a) is as in (63).

(63) **Hamblin set**

{ (63a), (63b) }

a. SLEEP(Al)

b. ¬RUN(Al)

Constraining alternatives

The Hamblin set in (63) contains answers which vary in polarity: (63a) is positive, and (63b) is negative. We considered a question with that profile in Section 3, and suggested that it would not be accommodated out of the blue as a default QUD. Yet, this question is licit as an *overt* QUD, and can supply the alternatives for *only* in (62b). ALT is, then, as in (64). The positive alternative (that Al slept) corresponds to the prejacent of *only* in (62b). As such, the alternative that would get excluded by *only* is the negative one (that Al didn't run). Excluding that alternative yields the inference in (65), that Al *did* run.

(64) **Predicted ALT**

ALT = { SLEEP(Al), \neg RUN(Al) }

(65) **Entailment**

$\neg\neg$ RUN(Al) \Leftrightarrow RUN(Al)

Strikingly, the predicted inference is observed. In the context in (62), (62b) conveys that Al obeyed the doctor's order to sleep, but failed to obey the other order. That is, Al did run. Even in symmetry data, a more complex negative alternative is in ALT when it is an answer to the QUD. The data in (62) thus provide strong support for the QUD approach.⁴

5.2. The QUD is crucial

It bears note that, in contrast to our proposal, Fox & Katzir (2011) took it that context *cannot* affect symmetry breaking. They reached this conclusion based on data similar to (66). (66a) makes salient the negative proposition that they did not steal the jewelry. Yet, nonetheless, that is not present in the alternative set for *only* in (66b). (66b) agrees with (66a) in saying that they did not steal the jewelry last week either. Hence, ALT must be as in (67), containing the positive alternative (that they stole the jewelry) for *only* to exclude.

- (66) a. In today's robbery, they didn't steal the jewelry.
b. In last week's robbery, they only [stole the books]_F.

(67) **Target ALT**

{ STEAL(they, the jewelry), ~~\neg STEAL(they, the jewelry)~~, ... }

We suggest that this observation is, in fact, predicted in our approach. Although the negative proposition is made *salient* in (67), it is not introduced as an *answer* to a QUD. Since (66) does not provide any explicit QUD, a QUD must be accommodated. The expected implicit QUD is as in (68). *What* ranges over properties, and intuitively, has its default restriction to positive actions. The Hamblin set in (69) contains only positive propositions, and the attested alternative set in (67), therefore, follows from the QUD.

⁴Haslinger & Schmitt (2024) suggest that complex alternatives are available, except when they would lead to the symmetry problem. They maintain that symmetry is broken in the grammar, based on complexity. Yet, (62) argues that C is controlling (see also Schwarz & Wagner 2024).

(68) **Implicit QUD**

What did they do in last week's robbery?

(69) **Hamblin set**

{ (69a), (69b), ... }

a. STEAL(they, the books)

b. STEAL(they, the jewelry)

Taken together, the data in (62) and (66) argue that context matters in determining ALT, and that the QUD, in particular, is controlling. A complex negative alternative is present in ALT when supplied by the QUD in (62). At the same time, salience, without QUD manipulation, fails to impact ALT in (67). Our proposal allows for a reversal in the direction of symmetry breaking in (62), while correctly preventing such reversal in (66).

6. Conclusion

In this paper, we have established that focus alternatives can be structurally more complex than the original prejacent. The results presented a challenge to Fox & Katzir's (2011) proposal where the grammar generally disallows complex alternatives. In response, we have pursued a view where alternatives are generated in the grammar irrespective of structural complexity. Complex alternatives are regulated by an implicit or explicit QUD.

We have focused our discussion on data involving content predicates such as *cook* and *clean*. Yet, the problem of how to constrain alternatives also arises with logical vocabulary items. Consider, for instance, the sentence in (70), where the determiner *some* is under focus. (70) conveys that Al ate some of the cookies, *but not all*. As such, ALT must contain the positive alternative in (70a), but not the negative alternative in (70b). If (70b) were in ALT, it would be excluded to derive the inference that Al *did* eat all of the cookies.

(70) Al only ate some_F of the cookies.

a. Al ate all of the cookies.

b. Al ate not all of the cookies.

In considering whether the QUD approach can capture (70), it is necessary to identify the QUD for (70), and to assess whether that QUD breaks symmetry by supplying (70a), but not (70b), as a possible answer.⁵ We must leave logical items to future research.

References

Beaver, David, and Brady Clark. 2008. *Sense and sensitivity: How focus determines meaning*. Wiley-Blackwell.

⁵Schwarz & Wagner (2024) focus on logical cases, and pursue a different pragmatic approach to symmetric breaking, based on *blocking*. In their analysis, a value for ALT is unavailable if it would result in a meaning that could be expressed by a structurally simpler expression. In (70), if ALT contained (70b), but not (70a), (70) would convey that Mary ate *all* of the cookies. The same meaning could be conveyed more simply by 'Mary ate all of the cookies', so that value for ALT is illicit. An account based on blocking would not extend to non-logical cases like (2). If ALT in 'Al only cooked' contained \neg CLEAN(Al), the sentence would convey that Al cooked and also cleaned. There is no simple expression which expresses that conjunction, and so no competitor to block that value for ALT (as required in default contexts). Hence, we do not pursue a blocking account here.

Constraining alternatives

- Breheeny, Richard, Nathan Klinedinst, Jacopo Romoli, and Yasutada Sudo. 2018. The symmetry problem: current theories and prospects. *Natural Language Semantics* 26:85–110.
- Büring, Daniel. 2003. On D-trees, beans, and B-accent. *Linguistics and Philosophy* 26:511–545.
- Fox, Danny, and Roni Katzir. 2011. On the characterization of alternatives. *Natural Language Semantics* 19:87–107.
- Haslinger, Nina, and Viola Schmitt. 2024. Revisiting the role of structural complexity in symmetry breaking. talk at Sinn und Bedeutung 29, Noto (Sicily).
- Katzir, Roni. 2007. Structurally-defined alternatives. *Linguistics and Philosophy* 30:669–690.
- Katzir, Roni. 2023. On the roles of anaphoricity and questions in free focus. *Natural Language Semantics* 1–28.
- Roberts, Craige. 1996. Information structure in discourse: Toward a unified theory of formal pragmatics. *Ohio State University Working Papers in Linguistics* 49:91–136.
- Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5:1–69.
- Rooth, Mats. 1985. Association with focus. Doctoral Dissertation, University of Massachusetts Amherst.
- Rooth, Mats. 1992. A theory of focus interpretation. *Natural Language Semantics* 1:75–116.
- Schwarz, Bernhard, and Michael Wagner. 2024. Resolving symmetry without formal alternatives. In *Proceedings of Sinn und Bedeutung 28*, ed. Kristina Liefke, Daniel Gutzmann, Agata Renans, and Tatjana Scheffler, 829–847.
- Spector, Benjamin. 2008. An unnoticed reading for wh-questions: Elided answers and weak islands. *Linguistic Inquiry* 39:677–686.
- Trinh, Tue, and Andreas Haida. 2015. Constraining the derivation of alternatives. *Natural Language Semantics* 23:249–270.
- Von Stechow, Kai, and Irene Heim. 1997. Class notes on pragmatics. MIT.