Order selection

The logical foundations of exhaustivity

Lucas Fagen
The University of Chicago

1 Roadmap

Exhaustive inferences (aka scalar implicatures) are usually modeled as a form of reasoning about *alternatives*.

> WHAT ARE ALTERNATIVES?
>
> ✗ sentences
>
> ✓ model-theoretic objects
>
> → If alternatives are sentences, there is no connection between the meaning of a sentence and its exhaustive interpretation.

Instead I will develop a direct semantic mapping between literal content and exhaustive interpretation. The proposal improves on mainstream conceptions of alternatives in several respects.

- **Descriptive adequacy**: it makes better predictions about the content of exhaustive inferences.
- **Explanatory adequacy**: it illuminates the relation between semantic content and pragmatic interpretation.

2 Exhaustivity

Consider the following context:
The Linguistics and Computer Science departments at the University of Chicago are hosting a regular weekly series of panel discussion events in which grad students can converse and answer questions about AI, for the benefit and edification of an undergraduate audience. Any grad student from either department is allowed to show up and participate at any time, but nobody is required to participate.

(1) **A**: Who attended the panel?
B: Agnes.
⇝ *nobody else*

In this context, if B's answer is given with falling intonation, it licenses the conclusion that nobody attended other than Agnes.

> This is an extremely well-studied phenomenon that has been examined from pragmatic, syntactic, and psycholinguistic perspectives.
>
> - Semantics, not so much (despite appearances).
>
> There is some debate over whether exhaustivity is pragmatic (Horn, 1972; Gazdar, 1979) or semantic (Chierchia, Fox, and Spector, 2012).
>
> - Traditionally attributed to Grice's maxim of Quantity (Grice, 1967), but this has been challenged.
> - I want to stay neutral on this question.

My point of departure: *is the content of an exhaustive inference predictable from literal semantic content?*

✗ Contemporary theories say NO.

✓ I will argue YES.

To get clear judgments I will only discuss exhaustive inferences that attach to explicit question/answer exchanges.

- If all speech acts are uttered against a backgrounded Question Under Discussion (Roberts, 2012; Ginzburg, 1996) then this is a way to make the QUD explicit.
- The proposal generalizes to cases without an explicit question as long as there is a clear focus/background structure (Krifka, 1992).

3 Direction

One reason to ask a question is to gain information about the extension of a property.

- Someone who asks *who attended the panel?* wants to know who is in the extension of **attended**.
- I will assume that unembedded interrogative clauses either denote properties or make properties dynamically accessible (Krifka, 1992; von Stechow and Zimmermann, 1984).

The DIRECTION of an exhaustive inference describes whether it attributes the question-property or its complement to some individuals.

(2) Who attended the last panel?
  a. Agnes. ⇝ *nobody else*
  b. Some of the linguists. ⇝ *not every linguist; no computer scientists*

These inferences have NEGATIVE DIRECTION.

(3) Who attended the last panel?
  a. Not all of the linguists. ⇝ some of them did
  b. Agnes and Beatrice didn't both attend. ⇝ one of them did.

These inferences have POSITIVE DIRECTION.

(4) Who attended the last panel?
  a. Most linguists but not a lot of computer scientists. ⇝ *not every linguist, some computer scientists*
  b. If Agnes attended, then Beatrice attended. ⇝ *if Beatrice attended, then Agnes attended*

These inferences have MIXED DIRECTION.

The contemporary literature on alternatives attempts to predict direction by restricting the class of sentences that count as alternatives to another sentence.

- The alternatives to *Agnes attended* include *Beatrice attended*, *Charmian attended*, etc.
- The alternatives to *some of the linguists* include *all of the linguists*.

And so on. Exhaustivity amounts to negating the alternatives.

The symmetry problem

This means we need a theory of alternatives, or we can't predict content.

⇝ Why isn't *Beatrice didn't attend* an alternative to *Agnes attended*?

For any exhaustive inference modeled as the negation of q, we need to explain why $\neg q$ isn't an alternative, or else we would get the exact opposite inference from the attested one. This is called the SYMMETRY PROBLEM (Matsumoto, 1995; Katzir, 2007).

- A large amount of contemporary literature is devoted to addressing the symmetry problem, usually by proposing *syntactic* restrictions on the class of licit alternative sentences.
- The symmetry problem is usually stated in a theory-dependent way (presupposing that alternatives are sentences), but everyone needs to explain the directional facts.

The right analysis should say: positive answers have negative direction, negative answers have positive direction, and mixed answers have mixed direction.

4 Extent

The EXTENT of an exhaustive inference describes the set of individuals about whom exhaustive inferences are possible.

(5) Who attended the last panel?
  a. Agnes. ⇝ *nobody else*
  b. Not Beatrice. ⇝̸ *conclude nothing about anyone else*

von Stechow and Zimmermann, 1984 claim that *not Beatrice* implicates *everyone else*. The speakers I consulted agree that this is the exception, not the rule (van Rooij and Schulz, 2004 and Spector, 2007 agree).

Some negative answers do allow exhaustive inferences though.

(6) Who attended the last panel?
  a. Some of the linguists. ⇝ *not every linguist, no computer scientists*
  b. Not a lot of linguists. ⇝ *some linguists*
  ⇝̸ *conclude nothing about computer scientists*

The difference is that the positive answer, but not the negative answer, also allows inferences about computer scientists.

Extent and aboutness

The main challenge in describing extent is defining a coherent notion of "aboutness".

- Positive answers have UNRESTRICTED EXTENT: exhaustivity can target individuals that they aren't about.
- Negative answers have RESTRICTED EXTENT: exhaustivity can only target individuals that they are about.

But what is aboutness?

Also, what is polarity?

- Nonmonotonic answers vary in their extent properties.

(7) Who attended the last panel?
  a. Exactly three linguists. $\rightsquigarrow$ *no computer scientists*
  b. Some but not all of the linguists. $\rightsquigarrow$ *no computer scientists*

These answers have unrestricted extent.

(8) Who attended the last panel?
  a. Most syntacticians but not a lot of phonologists. $\not\rightsquigarrow$ *conclude nothing about computer scientists*
  b. More syntacticians than phonologists. $\not\rightsquigarrow$ *conclude nothing about computer scientists*
  c. If Agnes attended, then Beatrice attended. $\not\rightsquigarrow$ *conclude nothing about anyone else*

These answers have restricted extent.

A theory of alternatives should explain why answers have different extent properties.

→ This is also an instance of the symmetry problem.

5 Generalizations

Here are the empirical patterns I will attempt to explain.

(9) **Claims**
  a. POSITIVE answers have negative direction and unrestricted extent.
  b. NEGATIVE answers have positive direction and restricted extent.
  c. MIXED (nonmonotonic) answers vary in their direction and extent properties.

I claim that these facts can be predicted from content alone, without a supplementary theory of alternative sentences.

6 Formal alternatives

6.1 Scales

Neo-Gricean pragmatics (Horn, 1972; Gazdar, 1979; Hirschberg, 1985; Matsumoto, 1995) addresses the symmetry problem (the problem of predicting content) via LEXICAL SCALES.

- Expressions are conventionally associated with other expressions.
- *Some* is associated with *all*, *or* is associated with *and*, etc.

Only sentences that can be syntactically constructed by substituting scalemates for each other count as alternatives.

Why do scales exist?

- Scales are supposed to be a matter of conventional linguistic knowledge: they are "given to us" (Gazdar, 1979, pg. 58).
- This view lacks explanatory adequacy and is also unfalsifiable.

6.2 The structural theory

Katzir, 2007 proposes a very general syntactic algorithm that derives formal alternatives without lexical scales.

- Any expression can be replaced by any other expression that is no more syntactically complex than the original one.

A lot of literature has attempted to refine Katzir's algorithm (Fox and Katzir, 2011; Trinh and Haida, 2015; Trinh, 2018; Breheny et al., 2018; Buccola, Križ, and Chemla, 2022; Schwarz and Wagner, 2024; Hirsch and Schwarz, 2025; Haslinger and Schmitt, 2025), usually while assuming that alternatives are sentences.

✓Negative direction

Katzir, 2007 makes correct predictions about negative direction.

- *Agnes attended* has *Beatrice attended* as an alternative, but not *Beatrice didn't attend*.
- *Some of the linguists* has *all of the linguists* as an alternative but not *some but not all of the linguists*.

This follows from the syntactic complexity of negation.

✗Positive direction

Katzir, 2007 makes wrong predictions about positive direction.

- *Not all of the linguists* has both *none of the linguists* and *some of the linguists* as alternatives, but only the former should be available.
- Negatively restricted sentences are predicted to have both "positive" and "negative" alternatives, which should cancel out (Romoli, 2013).

Negative answers should have only "negative" alternatives, but this result does not follow from a complexity condition.

✓ For the same reason, Katzir, 2007 gets extent right.

✗ *Not Beatrice* has both *Agnes* and *not Agnes* as alternatives, so no conclusions about Agnes are predicted.

"Negation-induced symmetry" prevents an analysis of positive direction, but enables an analysis of restricted extent.

Revising the structural theory

Several subsequent proposals have attempted to fix up Katzir's algorithm to handle positive direction.

- Haslinger and Schmitt, 2025 replace complexity with "syntactic similarity", so negative answers have only negative alternatives.
- Trinh and Haida, 2015 propose a more syntactically complex "atomicity" condition that leads to the same outcome.

Result: *not all of the linguists* has *none of the linguists* as an alternative, but not *some of the linguists*, correctly predicting the positive direction of exhaustivity.

✗Restricted extent

The problem with this strategy is that it sacrifices the account of restricted extent that was previously available.

- *Not Beatrice* then has only *not Agnes*, *not Charmian* etc as alternatives without their symmetric counterparts.

It should then be possible to understand B's answer in (10) as implicating that everyone else attended the panel, contrary to fact.

(10) **A**: Who attended the last panel?
B: Not Beatrice. $\not\rightsquigarrow$ *conclude nothing about anyone else*

Katzir's algorithm cannot simultaneously account for positive direction and restricted extent: any explanation for one will get the other one wrong.

→ The culprit for this problem is the underlying assumption that alternatives are sentences.

Extent = whether exhaustivity can cover the whole domain of discourse, or only individuals that an answer is "about".

- This isn't a syntactic distinction. Individuals are model-theoretic objects, not linguistic expressions.
- Aboutness is semantic.

✗Nondetachability

Aside from particular empirical problems, no syntactic theory of alternatives can provide a principled explanation for the fact that exhaustive inferences are NONDETACHABLE.

- Different ways of expressing the same meaning are by and large associated with the same exhaustive inferences.

(11) Who attended the last panel?
a. Agnes attended.
b. Agnes came.
c. Agnes showed up.
d. Agnes was there.
e. Agnes was there.
f. Agnes was present.
g. The panel was attended by Agnes.
$\rightsquigarrow$ *nobody else*

If alternatives are sentences, we have no explanation for nondetachability.

6.3 Verdict on formal alternatives

The majority of the literature on scales, alternatives, and symmetry breaking is based on a false premise. *Alternatives are not sentences*.

... ok so, what are they?

In what follows I will develop a direct semantic mapping from the literal meaning of an answer to its exhaustive interpretation, without using syntax as a middleman.

(12) **Claim**: *Alternatives are model-theoretic objects*.

Specifically, in terms of a focus/background structure, they are possible background extensions.

7 Circumscription

A small amount of semantic literature has made progress toward this goal (Szabolcsi, 1983; Groenendijk and Stokhof, 1984; von Stechow and Zimmermann, 1984; van Rooij and Schulz, 2004; Schulz and van Rooij, 2006; Spector, 2007) and my proposal will build on the ideas of these authors.

Schulz and van Rooij, 2006, building on Groenendijk and Stokhof, 1984, formalize exhaustivity as in (13).

(13) a. $\texttt{exh} = \lambda T \lambda P. T(P_w) \wedge \neg \exists v : T(P_v) \wedge v < w$
b. $v < w$ iff $P_v \subset P_w$

Exh operates on *focus/background* pairs: T is a focused term and P is a background term.

(14) **A**: Who attended the last panel? $\lambda w \lambda x.\textbf{attended}_w(x)$
B: Agnes. $\lambda P.P(a)$

The question provides an intensional background property term, and the answer provides an extensional generalized quantifier meaning that takes the background as its argument.

Exhaustivity quantifies over possible P-extensions.

- The first conjunct $T(P_w)$ says that the answer is true in the actual world w.
- The second conjunct says that there is no possible world v in which i) the answer is true, and ii) P's extension is smaller in v than it is in w.

Applied to the exchange in (14), here is what exh says:

- The set of attendees includes Agnes.
- The set of attendees is the *smallest set that includes Agnes*.

$\rightarrow P = \{a\}$.
So, Agnes attended, and nobody else did (or else P would include other individuals).

Schulz and van Rooij, 2006 call this approach CIRCUMSCRIPTION.

- It is a variant on the notion of "predicate circumscription" from the computational literature on nonmonotonic reasoning (McCarthy, 1980).
- It provides a direct mapping from literal to exhaustive content, given an information-structural division into focus and background.

The **symmetry problem** is resolved here by the fact that worlds are compared using set inclusion, and smaller sets are preferred to larger ones.

→ Directional facts are taken to be definitional of exhaustivity, instead of outsourced to an independent theory of alternatives.

Circumscription of negative and mixed answers

Circumscription makes correct predictions about positive answers, but it doesn't work for negative answers.

- The smallest set that doesn't contain Beatrice is the empty set, but *not Beatrice* doesn't implicate *nobody*.

(15) Who attended the last panel?
a. Exactly three linguists. $\rightsquigarrow$ *nobody else*
b. Most linguists but not a lot of computer scientists. $\rightsquigarrow$ *not all linguists, some computer scientists*

It works for some mixed answers but not others.

✓ The smallest set containing exactly three linguists doesn't contain anyone else.

- The smallest set containing most linguists but not a lot of computer scientists contains most linguists, ✓not all linguists, and ✗no computer scientists.

7.1 Verdict on circumscription

Circumscription provides the means to describe exhaustive inferences in strictly model-theoretic terms. This is a very good thing.

✗ It says nothing about direction (or extent).

Spector, 2007 proposes a variant on circumscription that does make correct predictions about some negative and mixed answers, but in a rather opaque way.

- Spector translates answers into propositional logic, and makes stipulations about which formulae count as alternatives to other formulae within the translation language.

This works, but unclear why.

I will build on the basic, semantically transparent exhaustivity recipe by Schulz and van Rooij, 2006 and make it sensitive to the polarity of an answer.

8 The plan

Here's exhaustivity again.

(16) $\texttt{exh} = \lambda T \lambda P.T(P_w) \wedge \neg\exists v : T(P_v) \wedge v < w$

I propose that the $<$ order over worlds is an indeterminate parameter. It can be resolved to several relations. One is the DEFAULT ORDER, which underlies existing circumscription accounts.

(17) **The default order.**
$v <_{\texttt{default}} w$ iff $P_v \subset P_w$

All positive answers and some mixed answers (*exactly three linguists*, *some but not all of the linguists*) use the default order. Negative answers and the remaining mixed answers use the HYBRID ORDER, which enforces restricted extent.
The plan:

1. I will provide a formal description of the hybrid order and show that it makes correct predictions for negative and mixed answers.
2. Then I will develop a theory of ORDER SELECTION: I will show how to predict which order an answer uses.
3. Sometimes there are exceptions to the restricted extent pattern. These call for a third order, the DUAL ORDER.

The orders play the role that alternatives play in theories that use alternatives, but are predictable on the basis of content.

9 The hybrid order

To define the hybrid order, we need a way to make the notion of EXTENT formally precise.

- Answers provide a generalized quantifier meaning against a backgrounded property term.
- One advantage to this approach is that it means we can import formal tools from the logic of generalized quantifiers into the analysis.

Proposal: an answer is "about" the objects in its RESTRICTION.

Quantificational restriction

A generalized quantifier LIVES ON all the sets that it is closed under intersection with (Barwise and Cooper, 1981).

(18) $\mathbf{L} = \lambda T \lambda S.\forall P : T(P) \leftrightarrow T(P \cap S)$

A generalized quantifier is RESTRICTED by the intersection of all its live-on sets (von Stechow and Zimmermann, 1984).

(19) $\mathbf{R} = \lambda T. \bigcap \mathbf{L}(T)$

(20) **Sample restrictions**

a.	Agnes	$\{a\}$
b.	Some of the linguists	**linguists**
c.	No linguists	**linguists**
d.	Not Beatrice	$\{b\}$
e.	Agnes or Beatrice	$\{a, b\}$
f.	Some cats but no hamsters	**cats** $\cup$ **hamsters**

Extent can now be given a precise formal characterization, as follows.

- Exhaustive interpretation of an answer T with UNRESTRICTED EXTENT regulates the question-property's intersection with the entire domain of discourse.
- Exhaustive interpretation of an answer T with RESTRICTED EXTENT regulates the question-property's intersection with $\mathbf{R}(T)$.

9.1 Restriction operators

The hybrid order should only regulate the question-property's intersection with an answer's restriction.

(21) Who attended the last panel?
  a. Not all of the linguists. ⇝ *some of them did*
  b. Most phonologists but not a lot of syntacticians. ⇝ *not every phonologist, some syntacticians*

For negative answers like *not all of the linguists* we want to maximize P's intersection with the set of linguists.

- Maximizing P's intersection with $\mathbf{R}(T)$ is (correctly) trivial for *not Beatrice* whose literal meaning completely fixes P's status with respect to $\mathbf{R}(T)$.

For mixed answers like *most phonologists but not a lot of syntacticians* we want to minimize P's intersection with the set of phonologists and maximize P's intersection with the set of syntacticians.

→ Different subsets of the restriction need to be treated differently.

Proposal: a GQ's restriction can be sorted into POSITIVE and NEGATIVE subsets.

Positive and negative restriction

(22) **Restriction operators.**
  a. $\mathbf{R}^+ = \lambda T \lambda x. \exists S : \neg T(S) \wedge T(S \cup \{x\})$
  b. $\mathbf{R}^- = \lambda T \lambda x. \exists S : \neg T(S) \wedge T(S - \{x\})$

T's POSITIVE RESTRICTION consists of objects that can be added to a set that T doesn't apply to, to create a set that T does apply to. T's NEGATIVE RESTRICTION consists of objects that can be added to a set that T doesn't apply to, to create a set that T does apply to.

(23) **Some useful facts that are provably true**
  a. $\mathbf{R}^+(T) \cup \mathbf{R}^-(T) = \mathbf{R}(T)$
  b. Iff T is increasing, $\mathbf{R}^-(T)$ is empty.
  c. Iff T is decreasing, $\mathbf{R}^+(T)$ is empty.

(24) **Sample restrictions**
  a. Agnes $\mathbf{R}^+(T) = \{a\}, \mathbf{R}^-(T) = \varnothing$
  b. not Beatrice $\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \{b\}$
  c. some of the cats $\mathbf{R}^+(T) = \textbf{cats}, \mathbf{R}^-(T) = \varnothing$
  d. no hamsters $\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \textbf{hamsters}$
  e. exactly 3 kittens $\mathbf{R}^+(T) = \mathbf{R}^-(T) = \textbf{kittens}$
  f. some but not all of the kittens $\mathbf{R}^+(T) = \mathbf{R}^-(T) = \textbf{kittens}$
  g. if Agnes, then Beatrice $\mathbf{R}^+(T) = \{b\}, \mathbf{R}^-(T) = \{a\}$
  h. most goslings but not many cygnets
  $\mathbf{R}^+(T) = \textbf{goslings}, \mathbf{R}^-(T) = \textbf{cygnets}$

As far as I know, the idea that quantificational restrictions can be sorted by polarity has not been explored in the GQ literature, but these tools might have broader applications if the proposal is on the right track.

9.2 Defining the hybrid order

The hybrid order is defined on the basis of the restriction operators.

(25) **The hybrid order.**
  a. $v <^+ w$ iff $P_v \cap \mathbf{R}^+(T) \subset P_w \cap \mathbf{R}^-(T)$
  b. $v <^- w$ iff $P_v \cap \mathbf{R}^-(T) \supset P_w \cap \mathbf{R}^-(T)$
  c. $v \leq_{\texttt{hybrid}} w$ iff $v <^+ w$ or $v <^- w$
  d. $v <_{\texttt{hybrid}} w$ iff $v \leq_{\texttt{hybrid}} w$ and $v \not\leq_{\texttt{hybrid}} w$

The hybrid order prefers a world v to a world w either if P's intersection with $\mathbf{R}^+(T)$ is smaller in v, or if P's intersection with $\mathbf{R}^-(T)$ is larger in v, *and* if the same criteria do not also favor w over v.

- There will be cases where $v <^+ w$ and $w <^- v$.
- The "antisymmetry clause" (25-d) ensures that such cases cancel out.

This is pretty abstract so let's look at some concrete cases.
Context: Agnes and Beatrice are linguists, Charmian isn't.

(26) **A**: Who attended the last panel?
  B: Not Beatrice.
  $\{abc, ab, bc, b\} <_{\texttt{hybrid}} \boxed{\{ac, a, c, \varnothing\}}$

The $\mathbf{R}^-$ set is $\{b\}$. $<^-$ prefers worlds where Beatrice attended. These are all ruled out by the literal meaning of the answer, so exh has no effect.

(27) **A**: Who attended the last panel?
  B: No linguists.
  $\{abc, ab\} <_{\texttt{hybrid}} \{a, b, ac, bc\} <_{\texttt{hybrid}} \boxed{\{c, \varnothing\}}$

Exh still has no effect here but ranks worlds differently: worlds where both linguists attended are preferred to worlds where one attended, which are preferred to worlds in which neither attended (and the answer is true).

(28) **A**: Who attended the last panel?
B: Not all of the linguists.

$\{abc, ab\} <_{\texttt{hybrid}} \boxed{\{a, b, ac, bc\}} <_{\texttt{hybrid}}$ ~~$\{c, \varnothing\}$~~

Exhaustivity does have an effect here. The order over worlds here is the same as in (27), but because two equivalence classes are compatible with the answer's literal meaning, `exh` selects the minimal ones. These are worlds in which one linguist did attend the panel.

- Predictions are similar for *not both Agnes and Beatrice*, *not a lot of linguists*, and so on.
- No inferences licensed about nonlinguists.

Now let's look at mixed answers.
Context: Agnes and Beatrice are linguists, Charmian and Doris are computer scientists.

(29) **A**: Who attended the last panel?
B: Agnes but not Beatrice.

$\{b, bc, bd, bcd\} <_{\texttt{hybrid}} \{\varnothing, c, d, cd\}, \{ab, abc, abd, abcd\} <_{\texttt{hybrid}} \boxed{\{a, ac, ad, acd\}}$

The $\mathbf{R}^+$ set is $\{a\}$ and the $\mathbf{R}^-$ set is $\{b\}$. There are four equivalence classes here. Exhaustivity has no effect here but it is useful to look at how the order works.

- The most minimal are worlds where Beatrice but not Agnes attended.
- Then there are worlds where neither attended (favored by $<^+$) and worlds where both attended (favored by $<^-$). These two cancel out.
- The least minimal are worlds where Agnes but not Beatrice attended.

Only the latter are consistent with the answer.

Now let's look at a case where `exh` does something.

(30) **A**: Who attended the last panel?
B: Some of the linguists but not all of the computer scientists.

$\boxed{\{ac, ad, bc, bd\}} <_{\texttt{hybrid}}$ ~~$\{abc, abd\}$~~, ~~$\{a, b\}$~~ $<_{\texttt{hybrid}}$ ~~$\{ab\}$~~

The $\mathbf{R}^+$ set contains the linguists and the $\mathbf{R}^-$ set contains the computer scientists. There are actually nine equivalence classes here (given our toy model) and I have only shown the ones consistent with the meaning of the answer.

- The hybrid order prefers worlds where not all of the linguists attended to worlds where they all did.
- It also prefers worlds where some of the computer scientists attended to worlds where none did.

→ Hybrid order simultaneously minimizes and maximizes over different subsets of an answer's restriction.

Conditionals

Here's how the analysis works for a conditional answer.

(31) **A**: Who attended the last panel?
B: If Agnes attended, then Beatrice attended.

$\{a\} <_{\texttt{hybrid}} \boxed{\{ab\}}, \boxed{\{\varnothing\}} <_{\texttt{hybrid}}$ ~~$\{b\}$~~

The $\mathbf{R}^+$ set is $\{b\}$ and the $\mathbf{R}^-$ set is $\{a\}$. I am ignoring Charmian and Doris.

- $<^+$ prefers worlds where Agnes didn't attend.
- $<^-$ prefers worlds where Beatrice did attend.
- They agree on excluding worlds where only Beatrice attended, but are "tied" on whether both or neither attended.

→ That's exactly the conditional perfection inference we wanted.[a]

[a] Most conditionals are more complicated than this but I take it as a good sign that the hybrid order predicts conditional perfection for very basic cases.

9.3 Interim summary

The HYBRID ORDER makes correct predictions about direction for negative and mixed answers.

✓ This is done by minimizing over the positive restriction and maximizing over the negative restriction.

Also makes correct predictions about extent.

✓ Objects outside the total restriction don't affect the hybrid order.

- Restricted extent follows from the need to sort objects by polarity using the restriction operators.

We did all this without formal alternatives!

10 Skew

The default order works for positive and some mixed answers. The hybrid order works for negative answers and the remaining mixed answers.

... what determines which order an answer uses?

The plan:

1. I will start by identifying a model-theoretic property that distinguishes these two classes of answers.

2. Then I will propose some predictive ORDER SELECTION CONSTRAINTS that derive the pattern.

Monotonicity

Usually the polarity of a quantifier is identified with its MONOTONICITY properties.

(32) **Monotonicity.**
 a. T is *increasing* iff $\forall A, B : T(A) \wedge A \subseteq B \rightarrow T(B)$.
 b. T is *decreasing* iff $\forall A, B : T(A) \wedge B \subseteq A \rightarrow T(B)$.

(33) **Sample quantifiers**
 a. *Increasing quantifiers*: Agnes, some of the cats, Agnes or Beatrice
 b. *Decreasing quantifiers*: not Beatrice, no cats, not all of the cats
 c. *Nonmonotonic quantifiers*: exactly three cats, some but not all of the kittens, most cats but not many hamsters, if Agnes then Beatrice

Does this work?

- Increasing answers use the default order and decreasing answers use the hybrid order.

- Nonmonotonic answers vary though.

DEFAULT: exactly three cats, some but not all of the kittens
HYBRID: most cats but not many hamsters, if Agnes then Beatrice

... what distinguishes the nonmonotonic quantifiers that use the default vs. the hybrid order?

Positive and negative skew

These two classes can be distinguished by the following model-theoretic property, which I propose to call SKEW.

Skew describes the relationship between a quantifier's total restriction and its positive and negative subsets.

(34) **Skew.**
 a. T is *positively skewed* iff $\mathbf{R}(T) = \mathbf{R}^+(T)$.
 b. T is *negatively skewed* iff $\mathbf{R}(T) = \mathbf{R}^-(T)$.

Because the restriction operators are new (to my knowledge), so is skew.

(35) **Some useful facts that are provably true**
 a. If T is increasing, T is positively skewed.
 b. If T is decreasing, T is negatively skewed.
 c. If T is positively and negatively skewed, either T is nonmonotonic or T is both increasing and decreasing.
 d. If T is unskewed, T is nonmonotonic.

Skew distinguishes between increasing and decreasing quantifiers in the expected way but allows us to draw additional distinctions between nonmonotonic quantifiers.

(36) exactly three cats — *positive and negative skew*
 $\mathbf{R}^+(T) = \mathbf{R}^-(T) = \textbf{cats}$

(37) some but not all of the kittens — *positive and negative skew*
 $\mathbf{R}^+(T) = \mathbf{R}^-(T) = \textbf{kittens}$

(38) most goslings but not many cygnets — *unskewed*
 $\mathbf{R}^+(T) = \textbf{goslings}, \mathbf{R}^-(T) = \textbf{cygnets}$

(39) if Agnes, then Beatrice — *unskewed*
 $\mathbf{R}^+(T) = \{b\}, \mathbf{R}^-(T) = \{a\}$

Generalization: positively skewed quantifiers use the default order, all others use the hybrid order.

- Some nonmonotonic quantifiers are skewed *in both directions* and these use the default order.

- Other nonmonotonic quantifiers are *unskewed* and these use the hybrid order.

Positive skew is necessary and sufficient to use the default order.

Answer	Skew	Extent
Agnes	+	unrestricted
Some of the linguists	+	unrestricted
Every linguist	+	unrestricted
Agnes and Beatrice	+	unrestricted
Most linguists	+	unrestricted
Agnes or Beatrice	+	unrestricted
Exactly 3 linguists	±	unrestricted
some but not all of the linguists	±	unrestricted
Not Beatrice	-	restricted
No linguists	-	restricted
Not all of the linguists	-	restricted
Not both Agnes and Beatrice	-	restricted
Not a lot of linguists	-	restricted
Neither Agnes nor Beatrice	-	restricted
Agnes but not Beatrice	n/a	restricted
Most cats but no hamsters	n/a	restricted
if Agnes then Beatrice	n/a	restricted

11 Selection constraints

The remaining task is to explain why skew matters. I propose two constraints whose interaction yields the order selection pattern.

(40) **Reduce markedness.**
Use the least marked order possible.

(41) **The markedness hierarchy (preliminary).**
`default` $\succ$ `hybrid`

The first is just a general preference for the default order.

- `Default` is less complex than `hybrid`, and it licenses stronger conclusions.
- `Hybrid` contains "more negation".

I propose that `default` is unmarked and `hybrid` is marked.

What would go wrong if `default` were used with answers that aren't positively skewed?

Information preservation

In a sense we already know the answer to this question.

- If `default` applied to *not Beatrice* the result would be *nobody*.
- If `default` applied to *most linguists but not a lot of hamsters* the result would be *most but not all linguists and nobody else*.

→ `Default` wipes out "negative information" contained in an answer.

The notion of "informational content" we have been working with is quantificational restriction. So, we need to look at how exhaustivity affects an answer's restriction.

Here's exhaustivity again.

(42) $\texttt{exh} = \lambda T \lambda P.T(P_w) \wedge \neg\exists v : T(P_v) \wedge v < w$

For any T, we can compare the restriction of T to the restriction of $\texttt{exh}(T)$ for both orders.[1]

Here's how default exhaustivity ($\texttt{exh}_D$) affects quantificational restriction:

Answer	$\mathbf{R}(T)$	$\mathbf{R}(\texttt{exh}_D(T))$	$\mathbf{R}^+(\texttt{exh}_D(T))$	$\mathbf{R}^-(\texttt{exh}_D(T))$
Skew: +				
Agnes	$\{a\}$	D_e	$\{a\}$	$D_e - \{a\}$
Some of the cats	**cats**	D_e	**cats**	D_e
Every cat	**cats**	D_e	**cats**	$D_e -$ **cats**
Agnes and Beatrice	$\{a,b\}$	D_e	$\{a,b\}$	$D_e - \{a,b\}$
Many cats	**cats**	D_e	**cats**	D_e
Agnes or Beatrice	$\{a,b\}$	D_e	$\{a,b\}$	D_e
Skew: ±				
Exactly 3 cats	**cats**	D_e	**cats**	D_e
Some but not all of the cats	**cats**	D_e	**cats**	D_e
Skew: −				
Not Beatrice	$\{b\}$	D_e	$\varnothing$	D_e
No cats	**cats**	D_e	$\varnothing$	D_e
Not all of the cats	**cats**	D_e	$\varnothing$	D_e
Not both Agnes and Beatrice	$\{a,b\}$	D_e	$\varnothing$	D_e
Not many cats	**cats**	D_e	$\varnothing$	D_e
Neither Agnes nor Beatrice	$\{a,b\}$	D_e	$\varnothing$	D_e
Skew: n/a				
Agnes but not Beatrice	$\{a,b\}$	D_e	$\{a\}$	D_e
Many cats but not many hamsters	**cats**$\cup$ **hamsters**	D_e	**cats**	D_e
If Agnes, then Beatrice	$\{a,b\}$	D_e	$\varnothing$	D_e

[1] Because the second argument of `exh` is intensional, we actually need slightly different, intensional restriction operators for $\texttt{exh}(T)$, but this is straightforward so I'm glossing over it.

Skew and information preservation

Generalization: an answer's total restriction is formally recoverable using the restriction operators whenever the answer is positively skewed, and not otherwise.

- Default exhaustivity typically preserves the positive restriction but expands the negative restriction.
- If the positive restriction is identical to the total restriction, then exhaustivity preserves the total restriction for this reason.
- Otherwise some or all of the restriction is lost.

`Default` would make for a bad heuristic applied to answers that are not positively skewed, because it would make some of their informational content redundant or irrelevant.

I won't show the equivalent table for `hybrid` but it always preserves the total restriction.

- This is because `hybrid` only regulates P's intersection with $\mathbf{R}(T)$, so hybrid exhaustive answers are always "about" the same objects as their literal counterparts.

I propose that this "information preservation" difference is what drives order selection.

→ Semantic content must be recoverable from pragmatic interpretation.

(43) **Preserve restriction.**
Ensure that:
a. $\mathbf{R}(\texttt{exh}(T)) = \mathbf{R}(T)$, or
b. $\mathbf{R}^+(\texttt{exh}(T)) = \mathbf{R}(T)$, or
c. $\mathbf{R}^-(\texttt{exh}(T)) = \mathbf{R}(T)$.

Preserve Restriction says that exhaustivity should not change the objects that an answer is about.

- The restriction must play a role in determining the partition between objects that have vs. lack the question-property.
- This is a reasonable principle to use in communication on the assumption that language users make utterances whose informational content is not arbitrary.

Deriving order selection

The two principles Reduce Markedness and Preserve Restriction yield the right order selection pattern.

- If T is positively skewed, both orders satisfy Preserve Restriction and so the unmarked `default` is chosen.
- Otherwise only `hybrid` satisfies Preserve Restriction and is therefore chosen.

This is a significant result because the resolution of the symmetry problem is often taken to be arbitrary.

Fox and Katzir, 2011: "Since the direction appears to be arbitrary from the perspective of context, it seems natural to make the choice in grammar."

→ *prematurely pessimistic!*

I have argued that the direction and extent properties of exhaustivity follow from how a semantic markedness hierarchy interacts with a general constraint on information preservation.

- To my knowledge, the analysis is the first formal proposal to derive the facts about scales, alternatives, and symmetry breaking from independent pragmatic principles while remaining explicit, predictive, and falsifiable.

12 The dual order

von Stechow and Zimmermann, 1984 claim that negative answers can have unrestricted extent.

- Usually this is not possible, but in certain contexts speakers are ok with it.

(44) **A**: Among the undergraduates, the graduate students, and the faculty, who was subjected to the Stanford prison experiment?
B: No faculty, and very few graduate students.⇝ *the undergraduates were subjected to it*

(45) **A**: Among the undergraduates, the graduate students, and the faculty, who came to the department barbecue?
B: No faculty, and very few graduate students.$\not\rightsquigarrow$ *conclude nothing about the undergraduates*

Apparently the horrific nature of the Stanford prison experiment is enough to motivate unrestricted extent here.

- If the department barbecue is considered a form of torture in the context of utterance, unrestricted extent is possible.

Suppose A walks into a bike shop and talks to the mechanic B.

(46) **A**: Which of these bikes have tubeless tires? I'd like to buy one.
B: Not these two. $\rightsquigarrow$ *the others do*

(47) **A**: I recently heard about tubeless tires.
B: Yeah, they're pretty trendy now.
A: Which of these bikes have tubeless tires?
B: Not these two. $\not\rightsquigarrow$ *conclude nothing about the others*

The relevant factor here is A's desire to buy a bike: given this, A can conclude from Bs' response in (46) that the other bikes are fair game.

→ The generalization with negative answers is that restricted extent is basic, but additional pragmatic factors can motivate unrestricted extent.

(48) **The dual order.**
$v <_{\texttt{dual}} w$ iff $P_v \supset P_w$

(49) **The markedness hierarchy (final).**
$\texttt{default} \succ \texttt{hybrid} \succ \texttt{dual}$

I propose that the DUAL ORDER is available, but extremely marked, so it takes extra motivation to use.

- Leaving open what counts as sufficient motivation or whether this could be formalized.

The markedness hierarchy

What motivates the markedness hierarchy?

- In particular, why is `dual` more marked than `hybrid`?

In a formal logic, the role of negation is to reverse the entailment order over truth values.

- *Conjecture*: the markedness of linguistic negation derives from its role as an order-reversing operator.
- The hybrid order is partially "negative" while the dual order is completely "negative".

So the markedness hierarchy follows from the general markedness of negation.

13 Summary

Here are the basic ingredients of the system I have proposed.

(50) **Exhaustivity.**
$\texttt{exh} = \lambda T \lambda P.T(P_w) \wedge \neg\exists v : T(P_v) \wedge v < w$

(51) **The default order.**
$v <_{\texttt{default}} w$ iff $P_v \subset P_w$

(52) **The hybrid order.**
a. $v <^+ w$ iff $P_v \cap \mathbf{R}^+(T) \subset P_w \cap \mathbf{R}^+(T)$
b. $v <^- w$ iff $P_v \cap \mathbf{R}^-(T) \supset P_w \cap \mathbf{R}^-(T)$
c. $v \leq_{\texttt{hybrid}} w$ iff $v <^+ w$ or $v <^- w$
d. $v <_{\texttt{hybrid}} w$ iff $v \leq_{\texttt{hybrid}} w$ and $w \not\leq_{\texttt{hybrid}} v$

(53) **The dual order.**
$v <_{\texttt{dual}} w$ iff $P_v \supset P_w$

(54) **Restriction operators.**
a. $\mathbf{R} = \lambda T. \bigcap \{S | \forall P : T(P) \leftrightarrow T(P \cap S)\}$
b. $\mathbf{R}^+ = \lambda T \lambda x. \exists S : \neg T(S) \wedge T(S \cup \{x\})$
c. $\mathbf{R}^- = \lambda T \lambda x. \exists S : \neg T(S) \wedge T(S - \{x\})$

(55) **Preserve restriction.**
Ensure that:
a. $\mathbf{R}(\texttt{exh}(T)) = \mathbf{R}(T)$, or
b. $\mathbf{R}^+(\texttt{exh}(T)) = \mathbf{R}(T)$, or
c. $\mathbf{R}^-(\texttt{exh}(T)) = \mathbf{R}(T)$.

(56) **Reduce markedness.**
Use the least marked order possible.

(57) **The markedness hierarchy.**
$\texttt{default} \succ \texttt{hybrid} \succ \texttt{dual}$

There are no lexical scales and no formal alternatives, and there is no appeal to the representational properties of alternative sentences.

- If exhaustivity is pragmatic, then I have shown how to account for nondetachability.
- If exhaustivity is semantic, then it works exactly like any other natural language quantification.

Exhaustive inferences can and should be systematically predicted on the basis of literal content, without a supplementary theory of alternatives.

14 Thanks!

I'd like to thank the department, my friends and colleagues, everyone I've discussed this work with over the past couple years, and my committee: Chris Kennedy, Ming Xiang, and Itamar Francez.

References

Barwise, Jon and Robin Cooper (1981). "Generalized quantifiers and natural language". In: *Linguistics and Philosophy* 4, pp. 159–219.

Breheny, Richard, Nathan Klinedinst, Jacopo Romoli, and Yasutada Sudo (2018). "The symmetry problem: current theories and prospects". In: *Natural Language Semantics* 26, pp. 85–110.

Buccola, Brian, Manuel Križ, and Emmanuel Chemla (2022). "Conceptual alternatives: competition in language and beyond". In: *Linguistics and Philosophy* 45, pp. 265–291.

Chierchia, Gennaro, Danny Fox, and Benjamin Spector (2012). "The grammatical view of scalar implicatures and the relationship between semantics and pragmatics". In: *Semantics: an international handbook of natural language meaning vol. 3*. de Gruyter.

Fox, Danny and Roni Katzir (2011). "On the characterization of alternatives". In: *Natural Language Semantics* 19, pp. 89–107.

Gazdar, Gerald (1979). *Pragmatics*. Academic Press.

Ginzburg, Jonathan (1996). "Dynamics and the semantics of dialogue". In: *Language, Logic, and Computation Vol. 1*. Stanford.

Grice, H.P. (1967). "Logic and conversation". In: *Studies in the Way of Words (1989)*. Harvard University Press.

Groenendijk, Jeroen and Martin Stokhof (1984). "Studies on the semantics of questions and the pragmatics of answers". PhD thesis. University of Amsterdam.

Haslinger, Nina and Viola Schmitt (2025). "Revisiting the role of structural complexity in symmetry breaking". In: *Proceedings of Sinn und Bedeutung 29*.

Hirsch, Aron and Bernhard Schwarz (2025). "Constraining alternatives". In: *Proceedings of Sinn und Bedeutung 29*.

Hirschberg, Julia (1985). "A theory of scalar implicature". PhD thesis. University of Pennsylvania.

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

Katzir, Roni (2007). "Structurally-defined alternatives". In: *Linguistics and Philosophy* 30, pp. 669–690.

Krifka, Manfred (1992). "A compositional semantics for multiple focus constructions". In: *Proceedings of SALT 1*.

Matsumoto, Yo (1995). "The conversational condition on Horn scales". In: *Linguistics and Philosophy* 18, pp. 21–60.

McCarthy, John (1980). "Circumscription: a form of nonmonotonic reasoning". In: *Artificial Intelligence* 13, pp. 27–39.

Roberts, Craige (2012). "Information structure in discourse: towards an integrated formal theory of pragmatics". In: *Semantics and Pragmatics* 5.6, pp. 1–69.

Romoli, Jacopo (2013). "A problem for the structural characterization of alternatives". In: *Snippets* 27.

Schulz, Katrin and Robert van Rooij (2006). "Pragmatic meaning and nonmonotonic reasoning: the case of exhaustive interpretation". In: *Linguistics and Philosophy* 29, pp. 205–250.

Schwarz, Bernhard and Michael Wagner (2024). "Symmetry resolution and blocking". In: *Proceedings of Sinn und Bedeutung 28*.

Spector, Benjamin (2007). "Scalar implicatures: exhaustivity and Gricean reasoning". In: *Questions in dynamic semantics*. Elsevier.

Szabolcsi, Anna (1983). "Focusing properties, or the trap of first order". In: *Theoretical Linguistics* 10.1-3, pp. 125–146.

Trinh, Tue (2018). "Keeping it simple". In: *Natural Language Semantics* 26, pp. 111–124.

Trinh, Tue and Andreas Haida (2015). "Constraining the derivation of alternatives". In: *Natural Language Semantics* 23, pp. 249–270.

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

von Stechow, Arnim and Thomas Ede Zimmermann (1984). "Term answers and contextual change". In: *Linguistics* 22, pp. 3–40.