Order selection: the logical foundations of exhaustivity

Lucas Fagen

The University of Chicago

Abstract

This thesis argues that exhaustive inferences in question/answer exchanges (*A: Who purred? B: Frederick* $\rightsquigarrow$ *not Gertrude*) are systematically predictable on the basis of semantic content, without a supplementary theory of formal alternatives. An explicit recipe for calculating exhaustive inferences is developed and tested against a wide range of answers.

Chapter 1 surveys the descriptive properties of exhaustivity and argues that exhaustive inferences should be modeled using PREDICATE CIRCUMSCRIPTION: the assumption that an answer provides the "best" answer to a question, calculated using an order over possible worlds determined by the question that was asked. However, existing circumscription accounts only cover positive answers to questions, and do not make correct predictions about how negative and mixed answers are interpreted.

Chapter 2 shows that the exhaustive interpretation of negative and mixed answers can be accounted for using a novel order over possible worlds that simultaneously integrates positive and negative information. The chapter also argues that circumscription makes successful predictions about how exhaustivity interacts with semantic type.

Finally, chapter 3 develops a theory of order selection, arguing that which order a particular answer uses is guided by the interaction of a general preference for markedness reduction and a constraint on information preservation that requires an exhaustive answer to share a restriction with its literal counterpart. A markedness hierarchy over the orders is developed and grounded in the more general markedness of negation.

Acknowledgements

Semantics and pragmatics are difficult endeavors and I am grateful to my committee, who supported and challenged me as I attempted to work out the ideas presented here. I would like to thank Chris Kennedy for his mentorship and intellectual influence, Itamar Francez for his principled refusal to accept inelegance, and Ming Xiang for letting a semanticist with no previous experimental background hang out in her lab and soak up psycholinguistic knowledge. I cannot believe my good fortune in getting to learn from these people.

I have also been fortunate to learn from many other current and former University of Chicago faculty, in particular Karlos Arregi, Monica Do, Allyson Ettinger, Malte Willer, and Erik Zyman, not to mention Michèle Lowrie.

Enthusiastic thanks to Eszter Ronai and Elitzur Bar-Asher Siegal for writing papers with me. Much as I enjoyed writing this thesis, the intimate and uncanny look inside another mind that collaboration allows is exciting on another level.

I owe many favors to Julian Grove and Paul Keyser for proofreading the appendix to this thesis (remaining errors are my fault) and for general mathematical edification.

The linguists who encouraged me to become one deserve special thanks. Matt Pearson got me started in linguistics, and the late, much missed Jiahui Huang pushed me to apply to graduate school when I was dragging my feet. Infinite thanks to Chris Barker, whose guidance since day one is a big reason why I'm here.

Many other scholars in the broader linguistics community discussed this and other projects with me, and I would like to especially thank Diti Bhadra, Brian Buccola, Cleo Condoravdi, Elizabeth Coppock, Donka Farkas, Chris Potts, and Aaron Steven White.

Thanks to all the University of Chicago students and alumni whose company I have enjoyed over the past six years, with a special shoutout to Marcia Gonzales Llanos, Mourad Heddaya, Sanghee Kim, Daniel Lam, Jacob Reed, Stephanie Reyes, Kutay Serova, Anna-Marie Sprenger, Mike Tabatowski, and Chih-chan Tien. Si Kai Lee and Xuetong Yuan do not belong in this category, but must be thanked in the same paragraph.

It's Ally Watson's fault I'm a linguist. It is striking to realize the extent to which our 2017 independent study on pragmatics prefigured my research interests throughout graduate school, and I hope my work continues to reflect her adventurous spirit.

Eternal gratitude as well to my friends and family outside the field. The debt I owe them is mostly unrelated to this work, but I do need to thank Lindsay for naming the individuals in my example sentences.

Contents

Introduction

This thesis is about how answers to questions are understood in conversation. Suppose there are three cats, Frederick, Gertrude, and Harriet. B's response to A's question in (1) can be heard as conveying that Frederick and Harriet did not purr, even though this was not explicitly stated.

(1) **A**: Who purred?
B: Gertrude.
⇝ *nobody else*

Similarly, B's response to A's question in (2) can be heard as conveying that Frederick and Gertrude didn't both purr, and that Harriet didn't purr, even though both possibilities are compatible with the literal inclusive meaning of linguistic disjunction.

(2) **A**: Who purred?
B: Frederick or Gertrude.
⇝ *not both*
⇝ *not Harriet*

This phenomenon, which has gone under many names in the literature—Quantity implicature, scalar implicature, exhaustivity, exhaustification, enrichment, strengthening, and more—is the subject of a rich and methodologically varied line of linguistic literature. I will use the descriptive, theory-neutral term EXHAUSTIVITY.

Exhaustivity has been extensively studied from pragmatic, syntactic, and psycholinguistic perspectives. Despite appearances, a semantic theory of exhaustivity has not been developed to the same extent.

Exhaustive implications are sensitive to information structure, background knowledge, the goals of the conversational participants, and other sources of contextual information. Yet a large class of exhaustive implications appear to be linguistically systematic and predictable on the basis of the sentence that was uttered and the context of utterance. Exhaustivity is thus a source of important insight into how semantic and pragmatic information is integrated in conversation. Because exhaustive implications appear to arise from the presumption that an utterance is the most informative utterance a speaker could have made, the study of exhaustivity additionally promises to reveal much about how informativity is implemented in natural language, and how language users reason about it.

One major question that has arisen in the study of exhaustivity is whether this phenomenon is semantic or pragmatic. Traditionally, exhaustive implications have been attributed to Grice's maxim of Quantity, which enforces a preference for informationally stronger utterances (Gazdar, 1979; Geurts, 2010; Grice, 1967; Hirschberg, 1985; Horn, 1972, 1984, 1989; Levinson, 2000). More recently, a productive line of literature has argued that exhaustive implications should be derived in compositional semantics (Chierchia, 2004, 2006; Chierchia, Fox, and Spector, 2012; Fox, 2007; Fox and Spector, 2018). At the same time, more explicit and testable pragmatic accounts have been developed using game theory (Bergen, Levy, and Goodman, 2016; Frank and Goodman, 2012; Franke, 2011; Parikh, 2001; Russell, 2012).

I do not take a stance on whether exhaustivity is semantic or pragmatic. The main question at issue in this thesis is whether the content of an exhaustive implication is systematically predictable from literal semantic content. According to existing theories, the answer is no. Both pragmatic and semantic views agree that exhaustive implications involve reference to alternative sentences, and that exhaustivity is a matter of negating these alternative sentences.

For example: to generate the implication from B's answer in (1) that Frederick didn't purr, such a theory would claim that the sentence *Frederick purred* is an alternative to the sentence *Gertrude purred*, and apply negation to the former sentence.

The task of predicting content is then outsourced to an independent theory of alternative sentences, which are usually characterized in syntactic terms: the alternatives to a sentence are all those sentences that stand in a particular syntactic relation to the original sentence. Theories whose goal is to predict content have accepted this assumption, and accordingly focused on elucidating the nature of this syntactic relation (Buccola, Križ, and Chemla, 2022; Fox and Katzir, 2011; Gazdar, 1979; Hirschberg, 1985; Katzir, 2007; Matsumoto, 1995).

Historically, the idea that exhaustivity is about negating alternative sentences arose from early attempts to formalize Gricean reasoning, in particular Gazdar, 1979. Because sentences are the linguistic objects that speakers use to make communicative utterances, this assumption seems unproblematic. However, it is not consistent with Grice, 1967, for whom Quantity implicatures were NONDETACHABLE.

An inference is nondetachable if any sentence with a particular meaning will license that inference. According to theories that use alternative sentences to derive exhaustive implications, they should be detachable, because synonymous sentences that vary in syntactic structure might be associated with different alternatives. This thesis will argue that exhaustive implications are largely nondetachable.

On the semantic side, one reason why theories that use alternative sentences appear plausible is because many logical languages do work this way. For instance, any first-order predicate logic in the tradition of Tarski, 1956 that handles quantification via assignment modification will involve metalinguistic reference to other logical formulae. However, natural language quantification is not usually analyzed in this metalinguistic way. Quantified sentences of natural language are usually given a compositional, model-theoretic interpretation, even if the logic we use to state truth-conditions is not itself compositional.

This thesis rejects the use of alternative sentences in the analysis of exhaustivity. Instead, I build on a minority semantic approach (Balogh, 2009; Groenendijk and Stokhof, 1984; Schulz and van Rooij, 2006; Sevi, 2005; Spector, 2007; Szabolcsi, 1981a,b, 1983; van Rooij and Schulz, 2004; von Stechow and Zimmermann, 1984). The thesis develops this approach in detail and draws out its implications for the contemporary literature on alternatives.

The primary claim of this thesis is that the content of exhaustive implications can be directly predicted on the basis of literal content and the context of utterance. Specifically, I argue for a CIRCUMSCRIPTION approach to exhaustivity—exhaustive implications arise when an answer to a question is taken to provide the "best" answer to that question—and I show that this proposal provides a solid foundation for principled analyses of many properties of exhaustivity that theories based on alternatives have struggled to explain.

The proposal is stated using predicate logic and the lambda calculus. These are the simplest and most general methodological tools at my disposal. The proposal is thus more abstract than probabilistic game-theoretic accounts, which explicitly model the decisions of the conversational participants. In the terminology of Franke and Jäger, 2016, this is an "algebraic" approach rather than an "interactive" approach. Although I believe that the ideas developed here could be implemented in a game-theoretic setting such as the "iterated best response" model of Franke, 2011, which is explicitly inspired by circumscription, game-theoretic analyses are typically focused on different questions and tend to take alternatives for granted (but see Bergen, Levy, and Goodman, 2016). From a game-theoretic perspective, I am concerned with how to set up the game, not how to play it.

The thesis is organized as follows. Chapter 1 summarizes the relevant empirical properties of exhaustive implications that any theory must explain—focusing in particular on what I call DIRECTION and EXTENT—and discusses two theoretical paradigms. The first one, FORMAL ALTERNATIVES, assumes that exhaustivity is about negating alternative sentences, and predicts the content of exhaustive implications by restricting the class of sentences that count as alternatives. The second one, CIRCUMSCRIPTION, defines a direct semantic mapping from the content of an answer to the content of its exhaustive interpretation. The chapter raises a number of problems for formal alternatives and argues that circumscription can resolve these problems. However, existing circumscription accounts only cover positive answers, and do not explain how the meaning of an answer—in particular, whether it has positive, negative, or mixed polarity—affects its direction and extent properties.

Chapter 2 develops an elaboration of circumscription that makes better predictions about these phenomena. The order over possible worlds that determines whether an answer is "best" is argued to be an indeterminate parameter that can be resolved to one of several relations. The DEFAULT ORDER is a simple, inclusion-induced order that minimizes the extension of a property made available by the question. This order implicitly underlies existing theories based on circumscription. To properly treat negative answers, however, we also need to recognize the more marked HYBRID ORDER, which simultaneously minimizes and maximizes different subsets of the question-property in a way determined by the semantics of the answer. The hybrid order is developed in detail and shown to make correct predictions about the direction and extent properties of negative and mixed answers. The chapter also argues that circumscription makes correct predictions about how exhaustivity interacts with semantic type.

Chapter 3 asks how answers choose their orders: why does exhaustive interpretation of positive answers use the default order, and exhaustive interpretation of negative answers use the hybrid order? The chapter develops a predictive theory of order selection based on two principles: REDUCE MARKEDNESS, which enforces a preference for a less marked order, and PRESERVE RESTRICTION, which requires the exhaustive interpretation of an answer to still be about the same objects (in a sense to be made precise) as its literal interpretation. The chapter also suggests that the DUAL ORDER, which is just the inverse of the default order, is sometimes used. The dual order is incorporated into the markedness hierarchy over available orders, which turns out to directly reflect the markedness of negation.

The conclusion summarizes the main contributions of the thesis and briefly discusses potential ramifications for the broader literature, in particular the question of whether exhaustivity is semantic or pragmatic. Finally, a short appendix proves a number of formal results pertaining to the logic of the analysis.

Chapter 1

Exhaustivity

This chapter summarizes the descriptive properties of exhaustive implications and discusses the two main characterizations of alternatives in the literature. The first and more mainstream approach, FORMAL ALTERNATIVES, is a syntactic approach. It starts from the idea that alternatives are sentences, and attempts to account for the exhaustivity facts by constraining the class of sentences that can count as alternatives to another sentence. The chapter discusses the core assumptions behind this approach and raises a number of problems for it, arguing that the characterization of alternatives as sentences is not viable.

The second approach, CIRCUMSCRIPTION, provides a semantic characterization of alternatives. Rather than defining alternatives as syntactic objects, this approach accounts for exhaustive interpretation by defining an order over possible worlds and taking the exhaustive interpretation of an utterance to consist of worlds in which a particular proposition provides the "best" answer to a question, on the basis of this order. The chapter argues that this approach can provide a solution to the problems formal alternatives face, and concludes by describing additional desiderata for a successful solution.

1.1 Characteristics of exhaustive implications

Exhaustive implications are usually viewed as inferences that result from the supposition that a particular utterance is the strongest or most informative utterance a speaker could have made in a particular context. From the fact that no stronger utterance was made, language users can enrich the meaning of the sentence that was uttered with additional information; specifically, certain stronger propositions can be taken to be false. All existing theories have in common the idea that exhaustivity arises from a process of maximizing informativity. Accordingly, all existing theories establish an informativity-based order over alternatives in some way, and use this order to derive exhaustivity.

The theories I discuss in this chapter differ primarily in the way they define this order and the objects they take to be ordered. In particular, they differ on whether these objects are model-theoretic objects (denotations), or expression/denotation pairs (sentences). Both approaches are designed to predict the exhaustive interpretation of a sentence on the basis of the sentence itself and the context in which it was uttered. The chapter starts by discussing ways in which exhaustivity interacts systematically with meaning and context.

A wide range of semantic and pragmatic phenomena have been argued by the contemporary literature to arise as consequences of exhaustive interpretation (free choice, homogeneity, and so on). In many cases, whether a particular phenomenon is really related to exhaustivity remains up for debate. In what follows, I stick to discussing cases whose status as exhaustive inferences is noncontroversial.

I follow the terminological convention in the literature of referring to exhaustive implications as IMPLICATURES, but wish to stay neutral on the question of whether exhaustivity is semantic or pragmatic.

1.1.1 Cancelability and nondetachability

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.

In this context, if B's answers in the following constructed dialogue exchanges are read with falling intonation, they can give rise to the exhaustive implications shown below.

(1) **A**: Who attended the last panel?
B: Agnes.
$\rightsquigarrow$ *nobody but Agnes*

(2) **A**: Who attended the last panel?
B: Most linguists.
$\rightsquigarrow$ *not all linguists*
$\rightsquigarrow$ *no computer scientists*

(3) **A**: What's the dress code for the panel?
B: Sneakers are allowed.
$\rightsquigarrow$ *sneakers aren't required*

Throughout this thesis, answers should be read with the falling intonation characteristic of exhaustive interpretation. The proposal I present in the following chapters will not explicitly cover modal implicatures like (3), but I include modal examples in this section to illustrate the range of the phenomenon.

Two fundamental properties of exhaustivity are that it is CANCELABLE and NONDETACHABLE. These features have traditionally been taken as evidence that exhaustive implications are conversational implicatures (Grice, 1967). Regardless of whether this view is correct, any theory needs to account for cancelability and nondetachability.

An inference is CANCELABLE if it can be withdrawn in the face of additional information. Exhaustivity is cancelable because an utterance that might otherwise be interpreted exhaustively can be followed with another utterance whose content is inconsistent with exhaustive interpretation of the first utterance, resulting in the the original exhaustive inference being withdrawn. The inferences shown in (1)-(3) are cancelable, as shown in (4)-(7).

(4) **A**: Who attended the last panel?
B: Agnes attended. Beatrice attended too.
$\not\rightsquigarrow$ *nobody but Agnes*

(5) **A**: Who attended the last panel?
B: Most linguists attended. In fact, every linguist attended.
$\not\rightsquigarrow$ *not all linguists*
$\rightsquigarrow$ *no computer scientists*

(6) **A**: Who attended the last panel?
B: Most linguists attended. Some computer scientists attended too.
$\rightsquigarrow$ *not all linguists*
$\not\rightsquigarrow$ *no computer scientists*

(7) **A**: What's the dress code for the panel?
B: Sneakers are allowed. In fact, they're required.
$\not\rightsquigarrow$ *sneakers aren't required*

B's followup sentences in (4)-(7) cancel the exhaustive implications that might otherwise arise. This property distinguishes exhaustive implications from noncancelable entailments, which cannot be withdrawn in the face of inconsistent information (*Agnes attended. #In fact, she didn't attend*). Cancellation is not absolute, as the contrast between (5) and (6) illustrates: in (5), B cancels the exhaustive implication that not every linguist attended, but it can still be concluded from B's answer that no computer scientist attended; conversely, in (6), B cancels the exhaustive implication that no computer scientist attended, but it can still be concluded that not every linguist attended. These examples show that exhaustivity does not categorically disappear when particular inferences are canceled, an observation that applies to the other examples too: although B's claim in (4) that Beatrice attended cancels the inference that nobody but Agnes attended, it can still be concluded that other individuals besides Agnes and Beatrice did not attend.

The content of an exhaustive inference can often be paraphrased by adding an overt exclusive modifier like *only* to an answer. Exclusive content from *only* is not cancelable: (4) forms a minimal pair with (8).

(8) **A**: Who attended the last panel?
B: Only Agnes attended. #Beatrice attended too.

I do not discuss the semantics of exclusives like *only* in this thesis, and the proposal I develop for exhaustivity is not intended to cover exclusives. Some exhaustive implications cannot be paraphrased with *only*. It is difficult to interpret *only most linguists attended* as denying that they all did, or *sneakers are only allowed* as denying that they're required.

There are several ways to account for cancelability. Most straightforwardly, if exhaustive implications are conversational implicatures then it suffices to say that exhaustive interpretation is only possible when consistent with the context of utterance (Gazdar, 1979). Compositional semantic theories of exhaustivity can access a version of this explanation by supposing that sentences are in general ambiguous between exhaustive and nonexhaustive readings, and that cancellation is a matter of ambiguity resolution in favor of the nonexhaustive reading (Chierchia, Fox, and Spector, 2012). It is also possible that exhaustive interpretation is always in force, but its output is subject to revision over the course of a conversation, with followup sentences inspiring revision of which alternatives are taken to be relevant for purposes of exhaustivity.

Explanations for cancelability typically lie in how alternatives are used, rather than how alternatives are defined. Both approaches to characterizing alternatives that I discuss in this chapter are therefore consistent with cancelability.

An inference is NONDETACHABLE if the linguistic form of an utterance does not affect the presence of the inference. Exhaustivity is nondetachable because distinct sentences with the same literal meaning generally allow the same exhaustive inferences. The facts in (9)-(11) illustrate nondetachability.

(9) **A**: Who attended the last panel?
B: {Agnes attended.
Agnes came.
Agnes showed up.
Agnes was there.
Agnes was present.
The panel was attended by Agnes.}
$\rightsquigarrow$ *nobody but Agnes*

(10) **A**: Who attended the last panel?
B: {Most linguists attended.
Most linguists came.
Most linguists showed up.
Most linguists were there.
Most linguists were present.
The panel was attended by most linguists.}
$\rightsquigarrow$ *not all linguists*
$\rightsquigarrow$*no computer scientists*

(11) **A**: What's the dress code for the panel?
B: {You're allowed to wear sneakers.
You can wear sneakers.
Sneakers are allowed.
Sneakers are permitted.
Sneakers are acceptable.
Wearing sneakers is not forbidden by the dress code.}
$\rightsquigarrow$ *you're not required to wear sneakers*

B's answers in the above exchanges differ in linguistic form, but not semantic content. In each case, varying the linguistic form does not affect the exhaustive interpretation of the answer: all the answers in (9) implicate that nobody but Agnes attended, for instance.

Similar nondetachability facts are shown in (12)-(14). If Agnes is the only student in the department working on sluicing, then B's answers in (12) are extensionally equivalent; similarly, if the only sneakers in the universe are Chuck Taylor All-Stars, then B's answers in (14) are extensionally equivalent. In each case, the same exhaustive implications arise.

(12) *Context: Agnes is the student working on sluicing.*
A: Who attended the last panel?
B: {Agnes attended.
The student working on sluicing attended.}
$\rightsquigarrow$ *nobody but Agnes*

(13) **A**: Who attended the last panel?
B: {Most linguists attended.
The majority of linguists attended.
The linguists mostly attended.
More than half of the linguists attended.
Over half of the linguists attended.}
$\rightsquigarrow$ *not all linguists*
$\rightsquigarrow$*no computer scientists*

(14) *Context: all sneakers are Chuck Taylors.*
A: What's the dress code for the panel?
B: {You're allowed to wear sneakers.
You're allowed to wear Chuck Taylors.}
$\rightsquigarrow$ *you're not required to wear sneakers*

Grice, 1967 takes nondetachability to be a characteristic property of conversational implicatures (except Manner implicatures). A nondetachable implicature is calculable solely on the basis of the meaning of a sentence and the context of utterance. A successful theory should explain why exhaustive implications are nondetachable. Here different theories of alternatives come apart. The two approaches I discuss in this chapter, in particular, make different predictions regarding nondetachability.

The semantic approach, CIRCUMSCRIPTION, takes only meaning and context into account when calculating the exhaustive interpretation of an utterance. This approach leads to the expectation that exhaustive interpretation should be nondetachable, because linguistic form plays no role in its calculation. The syntactic approach, FORMAL ALTERNATIVES, additionally takes linguistic form into account. On this approach, different sentences with the same meaning can in principle be associated with different alternatives. If this is correct, there is no reason to expect exhaustive implications to be nondetachable.

1.1.2 Direction

Exhaustive implications are often analyzed as the negation of stronger propositions than the sentence that was uttered. A successful theory needs to predict which stronger propositions are considered, and explain why these propositions count but not others.

The implication from B's answer in (2) (*Most linguists attended*) that not all linguists attended the last panel can be modeled as the negation of the proposition that all linguists attended, which is stronger than the answer that was given. If this proposition counts as an alternative to B's answer, exhaustive interpretation can add to the literal meaning of B's answer the negation of this stronger alternative. The interpretation of B's answer is then correctly predicted to be that most but not all linguists attended.

This story is not complete without an explanation of why the proposition that all linguists attended counts as an alternative. If any proposition counts as an alternative, this recipe for exhaustivity can predict implications that are never attested. The proposition that most but not all linguists attended is also stronger than the literal meaning of B's answer, and if this proposition counted as an alternative, we would predict that B's answer should implicate that all linguists attended, exactly the opposite of what actually occurs.

If both this proposition and the proposition that all linguists attended were alternatives, exhaustive interpretation would deliver the contradictory proposition, since these two propositions cannot both be excluded consistently with B's answer. Some theories define exhaustivity in such a way as to require consistency preservation (Fox, 2007; Sauerland, 2004; Schulz and van Rooij, 2006; Spector, 2007, 2016; van Rooij and Schulz, 2004); when two alternatives cannot be simultaneously excluded, they "cancel out". The proposition that most but not all linguists attended still cannot count as an alternative, since it would cancel out the alternative we need to generate the observed implication.

In the literature, the problem of characterizing alternatives is called the SYMMETRY PROBLEM. This problem was first observed by Kroch, 1972; the name "the symmetry problem" is attributed to von Fintel and Heim. A significant body of contemporary literature, including this thesis, aims to address the symmetry problem (Atlas and Levinson, 1981; Bergen, Levy, and Goodman, 2016; Block, 2008; Breheny et al., 2018; Buccola, Križ, and Chemla, 2022; Fox, 2007; Fox and Katzir, 2011; Gazdar, 1979; Haslinger and Schmitt, 2025; Hirsch and Schwarz, 2025; Katzir, 2007; Matsumoto, 1995; Schwarz and Wagner, 2024; Swanson, 2017; Trinh, 2018; Trinh and Haida, 2015; Westera, 2017a).

The symmetry problem is orthogonal to the question of whether exhaustivity is semantic or pragmatic. If exhaustive implications are the output of Gricean Quantity reasoning, we have to explain why this reasoning process concerns certain alternatives but not others. If exhaustive implications are compositionally generated, we likewise have to explain why the relevant compositional process can access certain alternatives but not others.

"Symmetrical" alternatives are propositions that COMPLEMENT each other in logical space. In possible world semantics, symmetry arises as a consequence of the fact that classical propositions are complemented: if W is the domain of possible worlds, any proposition $p \subset W$ will have a complement, $W - p$. More generally, for any proposition p that is associated with a stronger proposition q where $q \subset p$, there will also be a stronger proposition $p - q$. For this reason, if we want to model the exhaustive interpretation of p by excluding q, we have to explain why q is an alternative but $p - q$ is not also an alternative.

Complementation is a fundamental relation among sets. The logical foundations of the symmetry problem are thus not unique to classical logic or possible world semantics, but arise as soon as we adopt set theory. Analogous problems will arise for any theory that treats the denotations of sentences as sets of model-theoretic objects and implements entailment relations using a set inclusion criterion. If we are at all interested in using a set-theoretic metalanguage to describe the denotations of linguistic expressions, we need to address the symmetry problem.

The description of the symmetry problem I just gave was stated in a theory-dependent way. It presupposed that exhaustivity is a matter of excluding stronger alternative propositions. However, any theory needs to explain why a particular answer allows certain exhaustive implications but not others. In particular, a successful theory must predict whether the exhaustive interpretation of an answer is POSITIVE or NEGATIVE. Loosely following von Stechow and Zimmermann, 1984, who observe that different answers can have a different "direction of closure", I call this feature the DIRECTION of exhaustivity.

B's answer in (2) has a negative direction: in addition to the literal meaning that most linguists attended it is possible to infer that not all linguists and no computer scientists attended. Exhaustivity is negative here because it entails that some individuals do *not* have the property named by the question, that of having attended the last panel. Some answers, like (15), allow positive exhaustive inferences instead.

(15) **A**: Who attended the last panel?
B: Not a lot of linguists.
$\rightsquigarrow$ *some linguists*

The implication from B's answer that some linguists attended (sometimes called an INDIRECT IMPLICATURE, "indirect" because of the presence of negation; Atlas and Levinson, 1981; Horn, 1972; Sauerland, 2004) can also be accounted for in terms of maximizing informativity. The alternative proposition that no linguists attended is stronger, and excluding this proposition together with the literal meaning of the answer would generate the attested interpretation that some but not many linguists attended. In this case, however, the direction of exhaustivity is positive because it entails that some individuals *do* have the property named by the question. Negation reverses the direction.

The symmetry problem arises for (15) in exactly the same way it did for (2). If exhaustive interpretation of B's answer excludes the proposition that no linguists attended, we need to explain why this proposition is an alternative, but the proposition that some but not many linguists attended is not an alternative. Excluding this proposition would give us the inference that no linguists attended, which is the opposite of what we actually observe.

Similarly, we need to explain why exhaustivity is negative in (16) but positive in (17), negative in (18) but positive in (19), and negative in (3) but positive in (20).

(16) **A**: Who attended the last panel?
B: Some of the linguists.
$\rightsquigarrow$ *not all of the linguists*

(17) **A**: Who attended the last panel?
B: Not all of the linguists.
$\rightsquigarrow$ *some did*

(18) **A**: Who attended the last panel?
B: Agnes or Beatrice.
$\rightsquigarrow$ *not both*

(19) **A**: Who attended the last panel?
B: Not both Agnes and Beatrice.
$\rightsquigarrow$ *one of them did*

(3) **A**: What's the dress code for the panel?
B: Sneakers are allowed.
$\rightsquigarrow$ *sneakers aren't required*

(20) **A**: What's the dress code for the panel?
B: Sneakers aren't required.
$\rightsquigarrow$ *sneakers are allowed*

Informally, I will say that an answer to a question asking about a property of individuals is POSITIVE if it attributes the question-property to an individual or set of individuals; an answer is NEGATIVE if it attributes the complement of the question-property to an individual or set of individuals; a MIXED answer does both. The directional facts shown above suggest that the following generalization holds: positive answers have a negative direction of exhaustivity, and negative answers have a positive direction of exhaustivity. I am not aware of counterexamples to this generalization.

The direction of exhaustivity is not always uniform, however. Answers like (21) allow exhaustive implications that are both positive and negative.

(21) **A**: Who attended the last panel?
B: Most of the linguists, but not a lot of computer scientists.
$\rightsquigarrow$ *not all linguists, some computer scientists*

B's answer in (21) allows exhaustivity in both directions: the implication that not all linguists attended is negative while the implication that some computer scientists attended is positive. This example shows that exhaustivity can also apply to mixed answers. The symmetry problem must still be resolved in (21): we need to explain why we do not infer, for instance, that every linguist and no computer scientist attended, which is also consistent with the literal meaning of the answer.

One might wonder whether exhaustivity should apply uniformly to mixed answers like (21): perhaps exhaustivity should apply separately to each conjunct. If this were possible, then we could avoid the conclusion that exhaustivity has mixed direction. However, the problem with this strategy is that applying exhaustivity to the first conjunct *most of the linguists* would typically license the implication that no computer scientists attended, which contradicts the existential implication we want exhaustivity to deliver when applied to the second conjunct.

Another sort of mixed answer is shown in (22), which licenses a CONDITIONAL PERFECTION inference (Geis and Zwicky, 1971).

(22) **A**: Who attended the last panel?
B: If Agnes attended, then Beatrice attended.
$\rightsquigarrow$ *if Agnes didn't attend, then Beatrice didn't attend*

Conditional perfection (strengthening of *if* to *if and only if*) is often taken to be an instance of exhaustive interpretation (Horn, 2000; van der Auwera, 1997; von Fintel, 2001). The direction of exhaustivity in (22) is simultaneously positive and negative, but in a different way from (21). Instead of definitely inferring that some individuals have the property named by the question and others lack it, we infer what is in effect a disjunction of positive vs. negative exhaustivity: either both Agnes and Beatrice attended, or neither did. A successful theory should explain these directional facts too, which pose yet another instance of the symmetry problem: we need to explain why we infer conditional perfection instead of inferring that both Agnes and Beatrice attended, or that neither attended, or that only Beatrice attended.

The two characterizations of alternatives I discuss in this chapter make different predictions regarding the direction of exhaustivity. Theories based on formal alternatives attempt to break symmetry on the basis of whether sentences expressing alternative propositions stand in a particular syntactic relation to the sentence that was uttered. If this view is correct, then the directional facts are epiphenomenal. Because symmetry is broken syntactically, this approach predicts no principled connection between the meaning of an answer and the direction(s) its exhaustive interpretation will take.

Theories based on circumscription have largely not addressed the symmetry problem explicitly (although see Sevi, 2005). One goal of this thesis is to argue that circumscription makes available a more successful solution than formal alternatives can provide.

1.1.3 Extent

The meaning of an answer systematically affects what I call the EXTENT of exhaustivity. There is a clear contrast between (1) and (23).

(1) **A**: Who attended the last panel?
B: Agnes.
$\rightsquigarrow$ *nobody but Agnes*

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

von Stechow and Zimmermann, 1984 claim that, just as B's answer in (1) implicates that nobody but Agnes attended, B's answer in (23) implicates that *everyone* but Beatrice attended. Although this is sometimes possible, the speakers I consulted gave judgments that indicated this implication is infrequent and dispreferred. The more typical response to B's answer in (23) is that nothing can be concluded about other individuals (as also observed by Schulz and van Rooij, 2006; Spector, 2007; van Rooij and Schulz, 2004).

The claim that B's answer in (23) does not typically allow exhaustive implications about individuals other than Beatrice should not be mistaken for a claim that no other implicatures are possible. A might conclude from B's answer that B has a special interest in whether Beatrice attended. In a context where it is part of the common ground that Agnes attended if Beatrice didn't attend, then B's answer can convey that Agnes attended. This inference is a contextual entailment, and should not be attributed to exhaustive interpretation.

Independently of the directional contrast, (2) and (15) also differ in whether implications about computer scientists are attested.

(2) **A**: Who attended the last panel?
B: Most linguists.
$\rightsquigarrow$ *not all linguists*
$\rightsquigarrow$ *no computer scientists*

(15) **A**: Who attended the last panel?
B: Not a lot of linguists.
$\rightsquigarrow$ *some linguists*
$\not\rightsquigarrow$ *conclude nothing about computer scientists*

Although B's answer in (2) can implicate that no computer scientists attended, there is no analogous implication in (15). Given the directional contrast, we might expect (15) to implicate that every computer scientist attended. The speakers I have consulted agree that this is not the case. Perhaps B is ignorant about computer scientists, or doesn't care about them, or assumes A doesn't care about them; whatever B's pragmatic motivation, the usual exhaustive interpretation of this answer allows no clear conclusions about computer scientists and licenses only the conclusion that some linguists attended.

I use the term EXTENT to characterize the set of individuals whose status regarding the question-property can be determined via exhaustivity. The examples I just discussed suggest that positive and negative answers differ systematically in extent. Positive answers like (1) and (2) allow the extent of exhaustive interpretation to cover the entire domain of discourse. This is why (1) licenses inferences about individuals other than Agnes, and (2) licenses inferences about computer scientists in addition to linguists. I will say that these answers have UNRESTRICTED EXTENT.

Negative answers like (15) and (23) (usually) only allow the extent of exhaustive interpretation to cover individuals that, informally, the answer is "about". B's answer in (15) is about the linguists, and accordingly only allows exhaustive inferences about the linguists. B's answer in (23) is about Beatrice, and because Beatrice's status regarding the question-property is already determined by the meaning of the answer, no additional inferences are available. I will say that these answers have RESTRICTED EXTENT.

Mixed answers vary in extent. Some mixed answers, like (24), have restricted extent. B's answer conveys an exhaustive implication about the syntacticians and phonologists, but nothing is concluded about computer scientists or linguists who are not syntacticians or phonologists.

(24) **A**: Who attended the last panel?
B: Most syntacticians, but not a lot of phonologists.
$\rightsquigarrow$ *not all syntacticians, some phonologists*
$\not\rightsquigarrow$ *conclude nothing about other linguists or computer scientists*

Similarly, (22) allows a conditional perfection inference, but nothing is concluded about individuals other than Agnes or Beatrice.

(22) **A**: Who attended the last panel?
B: If Agnes attended, then Beatrice attended.
$\not\rightsquigarrow$ *conclude nothing about anyone else*

There are some mixed answers that do allow negative exhaustive implications with unrestricted extent, however. It is possible to conclude that no computer scientists attended from (25) and (26).

(25) **A**: Who attended the last panel?
B: Some but not all of the linguists.
$\rightsquigarrow$ *no computer scientists*

(26) **A**: Who attended the last panel?
B: Exactly three linguists.
$\rightsquigarrow$ *no computer scientists*

I am not aware of any mixed answers that allow positive exhaustive implications with unrestricted extent.

A successful theory should explain these contrasts in extent. Specifically, we need to explain why positive answers have unrestricted extent, why negative answers have restricted extent, and what determines the extent of a mixed answer.

A successful theory should also explain why it is sometimes possible for negative answers to have unrestricted extent, and why this is a marked option. I take restricted extent as the basic case for negative answers, and return to the marked cases in chapter 3.

The distinction between restricted and unrestricted extent has not been widely discussed in the literature. To my knowledge, only van Rooij and Schulz, 2004 and Spector, 2007 have suggested explanations for it. The challenge is to give formal substance to the informal notion of "aboutness". This thesis proposes an account based on the logic of generalized quantifiers.

Theories based on formal alternatives predict that the extent facts should reflect syntactic properties of sentences attributing the question-property to individuals that an answer is not about. They do not predict any connection between an answer's polarity and its extent, and must view such semantic generalizations as accidental.

1.1.4 Information structure

The extent of a sentential answer is not solely determined by its meaning, but systematically varies depending on the question that was asked.

(27) **A**: Who attended the last panel?
B: Beatrice didn't attend.
$\not\rightsquigarrow$ *conclude nothing about anyone else*

(28) **A**: Who didn't attend the last panel?
B: Beatrice didn't attend.
$\rightsquigarrow$ *everyone else did*

Although B's answer has restricted extent in (27), the same answer has unrestricted extent in (28). The source of restricted extent in (27) cannot be the presence of negation, but is rather the congruence relation between question and answer.

The exchanges in (29) and (30) display similar congruence effects. In (29), B's answer matches the polarity of A's question and has unrestricted extent; in (30), B's answer mismatches the polarity of A's question and has restricted extent.

(29) **A**: Who attended the last panel?
B: Agnes attended.
$\rightsquigarrow$ *nobody but Agnes*

(30) **A**: Who didn't attend the last panel?
B: Agnes attended.
$\not\rightsquigarrow$ *conclude nothing about anyone else*

The answers with unrestricted extent attribute the question-property to Agnes, whereas the answers with restricted extent attribute the question-property's complement to Agnes.

I suggested that the extent of an answer is predictable from its polarity: positive answers have unrestricted extent, and negative answers have restricted extent, with variation among mixed answers. The observation that extent is also sensitive to the question that was asked motivates a view of polarity as not intrinsic, but determined in relation to a question.

The behavior of fragment answers in this regard is telling. There is no contrast between B's answers in (1) and (31), which both have unrestricted extent: exhaustive interpretation of (1) conveys that nobody else attended the panel, and exhaustive interpretation of (31) conveys that everyone else attended the panel.

(1) **A**: Who attended the last panel?
B: Agnes.
$\rightsquigarrow$ *nobody but Agnes*

(31) **A**: Who didn't attend the last panel?
B: Agnes.
$\rightsquigarrow$ *everyone else attended*

Conversely, B's answers in (23) and (32) both have restricted extent.

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

(32) **A**: Who didn't attend the last panel?
B: Not Beatrice.
$\not\rightsquigarrow$ *conclude nothing about anyone else*

In particular, (32) is not ambiguous between a reading claiming that Beatrice attended and a reading claiming that Beatrice didn't attend. Only the former is available. The meaning of this answer must be determined by negating the application of the question-property to Beatrice. Since A's question asks about the property of not attending, B's answer means that

Beatrice *didn't* not attend; i.e., she attended. The complexity of the double negation explains why (32) sounds marked. Neither answer allows inferences about other individuals.

These observations can be accounted for if the meaning of fragment answers is uniformly determined by applying a generalized quantifier to the question-property to yield a propositional meaning.[1] The fragment *Agnes* applies the Montagovian individual Agnes (the generalized quantifier over properties that include Agnes) to the question-property, and the fragment *not Beatrice* applies the complement of the Montagovian individual Beatrice to the question-property. Accordingly, B's answers in (1) and (31) both attribute the question-property to Agnes, and B's answers in (23) and (32) both attribute the complement of the question-property to Beatrice. The extent properties of these answers can then be subsumed under the generalization that positive answers have unrestricted extent and negative answers have restricted extent.

If this analysis of fragment answers is right, neither the extent nor the direction of a fragment should vary across different questions. Exhaustive inferences with negative direction entail that some individuals lack the question-property. If the question contains negation, a negative exhaustive implication relative to this question will entail that some individuals have the negated property, but this should not be mistaken for a positive exhaustive implication, which would instead entail that some individuals have the question-property.

These predictions are borne out. B's answers in (33) and (34) have negative direction and unrestricted extent, just as they do in response to questions that do not contain negation. Exhaustive interpretation of (33) conveys that some linguists attended (not all linguists didn't attend) and every computer scientist attended (no computer scientists didn't attend), and exhaustive interpretation of (34) conveys that either Agnes or Beatrice didn't attend, one of them attended, and everyone else attended.

(33) **A**: Who didn't attend the last panel?
B: Most linguists.
$\rightsquigarrow$ *some linguists did*
$\rightsquigarrow$ *every computer scientist did*

(34) **A**: Who didn't attend the last panel?
B: Agnes or Beatrice.
$\rightsquigarrow$ *one of them did*
$\rightsquigarrow$ *everyone else did*

B's answers in (35) and (36) similarly have restricted extent, just as they do in response to questions that do not contain negation. Exhaustive interpretation of (35) conveys that some linguists didn't attend, but nothing is concluded about computer scientists; (36) conveys nothing about individuals other than Agnes and Beatrice. These fragment answers must be interpreted with double negation: B's answer in (35) can only mean that not a lot of linguists *didn't* attend, i.e., most linguists attended, and B's answer in (36) can only mean that neither Agnes nor Beatrice *didn't* attend, i.e., both attended.

(35) **A**: Who didn't attend the last panel?
B: Not a lot of linguists.
$\rightsquigarrow$ *some linguists didn't attend*
$\not\rightsquigarrow$ *conclude nothing about computer scientists*

[1] This could happen at an abstract level of syntactic representation (Merchant, 2004), or at the discourse level (Jacobson, 2016).

(36) **A**: Who didn't attend the last panel?
B: Neither Agnes nor Beatrice.
$\not\rightsquigarrow$ *conclude nothing about anyone else*

The mixed answers in (37) and (38) have different direction and extent properties. B's answer in (37) has mixed direction and restricted extent: exhaustive interpretation conveys that some syntacticians attended (not all syntacticians didn't attend) and some phonologists didn't attend, but conveys nothing about computer scientists or other linguists. B's answer in (38) has negative direction and unrestricted extent: everyone attended aside from exactly three linguists. This is parallel to how these answers are interpreted with questions that do not contain negation.

(37) **A**: Who didn't attend the last panel?
B: Most syntacticians, but not a lot of phonologists.
$\rightsquigarrow$ *some syntacticians attended*
$\rightsquigarrow$ *some phonologists didn't attend*
$\not\rightsquigarrow$ *conclude nothing about anyone else*

(38) **A**: Who didn't attend the last panel?
B: Exactly three linguists.
$\rightsquigarrow$ *everyone else attended*

Because fragment answers are uniformly interpreted by applying a generalized quantifier to the unique question-property, it is possible to give a more principled definition of answer polarity. A fragment is positive, negative, or mixed if its noun phrase denotes a MONOTONE INCREASING, DECREASING, or NONMONOTONIC quantifier (Barwise and Cooper, 1981).

(39) **Monotonicity.**
A generalized quantifier T is *increasing* iff $\forall A, B : T(A) \wedge A \subseteq B \rightarrow T(B)$.
A generalized quantifier T is *decreasing* iff $\forall A, B : T(A) \wedge B \subseteq A \rightarrow T(B)$.

The empirical claims I have made so far can be summarized as follows.

- Increasing quantifiers (*Agnes, some of the linguists, every linguist, Agnes and Beatrice, most linguists, Agnes or Beatrice*) give POSITIVE ANSWERS with negative direction and unrestricted extent.
- Decreasing quantifiers (*not Beatrice, no computer scientists, not every computer scientist, not both Agnes and Beatrice, not a lot of computer scientists, neither Agnes nor Beatrice*) give NEGATIVE ANSWERS with positive direction and restricted extent.
- Nonmonotonic quantifiers (*some but not all of the linguists, Agnes but not Beatrice, most linguists but not a lot of computer scientists, exactly three linguists, Agnes or else not Beatrice*) give MIXED answers whose direction and extent vary.

Exhaustivity does not distinguish between fragment and sentential answers: as far as I know, there are no fragments whose sentential variants receive a different exhaustive interpretation. Thus I propose to extend the same treatment to sentential answers.

I will assume that fragment and sentential answers both denote STRUCTURED PROPOSITIONS (Cresswell and von Stechow, 1982), pairs of quantifiers and properties. The property element corresponds to what I have been calling the question-property, the property whose extension the question is asking about. The quantifier element will yield an answer's propositional meaning when applied to the property element. In the case of fragments, the quantifier element will be the denotation of the overt noun phrase. In the case of sentential answers, the quantifier element will correspond to whatever remains of the answer's metalanguage representation after abstracting over the question-property.

Samples are given in (40).

(40) Who attended?

a. $[\![\text{Agnes attended}]\!] = \langle \lambda P_{et}.P(a), \lambda x_e.\textbf{attended}(x)\rangle$
b. $[\![\text{Beatrice didn't attend}]\!] = \langle \lambda P_{et}.\neg P(b), \lambda x_e.\textbf{attended}(x)\rangle$
c. $[\![\text{Agnes but not Beatrice attended}]\!] = \langle \lambda P_{et}.P(a) \wedge \neg P(b), \lambda x_e.\textbf{attended}(x)\rangle$
d. $[\![\text{if Agnes attended, Beatrice attended}]\!] = \langle \lambda P_{et}.P(a) \rightarrow P(b), \lambda x_e.\textbf{attended}(x)\rangle$

The literature on focus has used structured propositions to model the flow of information in discourse, with answers split into focused and backgrounded information (Jacobs, 1983; Krifka, 1992; von Stechow, 1991). In (40-a), the Montagovian individual Agnes provides the focus and the property of having attended provides the background. I do not want to claim that these structured representations reflect natural language syntax, and I will not discuss how answers are composed. A compositional analysis that delivers structured representations like those in (40) seems implausible: English *not* lacks NP-modifying uses outside fragments (Weir, 2020), and I know of no noun phrase in any language with the GQ meaning in (40-d). Moreover, Rooth, 1996 argues against using structured representations as a semantics for intonational focus.

The use of structured propositions in this thesis is intended to capture the distinction between old and new information in discourse. It is not intended as a claim about sentence-level grammar. I commit only to the hypothesis that old and new content can be distinguished at the level of information structure. I will assume that intonational focus tracks this distinction, without adopting the stronger hypothesis that the semantic role of intonational focus is to form pairs like those in (40) or assuming that this is the only thing that focus can do.

I have assumed that questions make properties formally accessible. One might wonder what sort of question semantics could support this assumption. Denotations for questions as sets of static propositions (Hamblin, 1973; Karttunen, 1977) or propositional concepts (Groenendijk and Stokhof, 1982) are not rich enough to determine a uniquely corresponding property (see Zimmermann, 1985, for discussion of this point). Some theories of questions claim they denote properties (Hausser, 1983; Hausser and Zaefferer, 1979; Jacobson, 2016; Xiang, 2021b). This position is compatible with my proposal, but not necessary. Li, 2021 shows that if questions denote sets of *dynamic* propositions, it is possible to reconstruct the question-property by taking the set of discourse referents introduced by propositions in the question denotation.

The distinction between structured and unstructured propositions recalls the distinction between classical propositional logic and TERM LOGIC (Horn, 1989; Łukasiewicz, 1951). von Stechow and Zimmermann, 1984, pg. 38 comment that without the formal accessibility of subject and predicate terms that structured representations allow, "a propositional theory of answers would have to define closure [i.e., direction] syntactically." Because one goal of this thesis is to avoid a syntactic characterization of alternatives, I side with von Stechow and Zimmermann, 1984 and Groenendijk and Stokhof, 1984 in favoring a term-logical analysis of exhaustivity.

The difference between examples like (27) and (28) then hinges on information structure. B's answers denote the same proposition, but negation is part of the focus in (27) and the background in (28).

(27) **A**: Who attended the last panel?
B: Beatrice didn't attend.
$\not\rightsquigarrow$ *conclude nothing about anyone else*

(28) **A**: Who didn't attend the last panel?
B: Beatrice didn't attend.
$\rightsquigarrow$ *everyone else did*

Against the background of A's question in (27), B's answer is structured as in (41-a); against the background of A's question in (28), B's answer is structured as in (41-b).

(41) a. $\langle \lambda P_{et}.\neg P(b), \lambda x_e.\textbf{attended}(x)\rangle$
b. $\langle \lambda P_{et}.P(b), \lambda x_e.\neg\textbf{attended}(x)\rangle$

The reason B's answer has restricted extent in (27) and unrestricted extent in (28) is because the focus in the former case is a decreasing quantifier, while the focus in the latter case is an increasing quantifier.

1.1.5 Granularity and suspension

Two important empirical properties of exhaustivity are that it is GRANULARITY-SENSITIVE and SUSPENDIBLE. These phenomena appear to challenge my earlier claim that exhaustive implications are nondetachable, and are therefore worth discussing in detail.

Horn, 1972 observed that the implications allowed by B's answer in exchanges like (42) depend on how specific an answer is expected in the context of utterance.

(42) **A**: Who attended the last panel?
B: Some of the linguists.
$\rightsquigarrow$ *not all*
$\rightsquigarrow$ *not most*
$\rightsquigarrow$ *not many*

Horn suggested that the regularity with which the inferences arise correlates with the strength of the excluded alternatives, a conjecture that the experimental literature has confirmed (van Tiel et al., 2016): B's answer conveys that all of the linguists didn't attend more robustly than it conveys that most or many of the linguists didn't attend.

The GRANULARITY of a set of alternatives refers to how many alternatives there are and how finely they partition logical space (Krifka, 2009). A set containing only the propositions that some of the linguists attended and every linguist attended is less granular than a set containing these in addition to the proposition that most linguists attended. Contextual variation in how fine-grained an exhaustive implication is can be explained if exhaustivity is stated relative to a set of alternatives whose granularity depends on the context of utterance. The correlation between the informative strength of an alternative and the robustness of an exhaustive implication excluding this alternative can be explained if stronger alternatives are distinguished at coarser granularity levels, whereas relatively weaker alternatives are only distinguished at finer granularity levels.

Sensitivity to granularity reflects the more general fact that exhaustive interpretation is affected by the interests of the conversational participants. Schulz and van Rooij, 2006 propose to determine an order over alternatives using general considerations of contextual relevance: what matters is the UTILITY of a proposition, defined using techniques from statistical decision theory. If exhaustive interpretation is defined relative to this order, a wide range of context-sensitive effects can be accounted for. If A is only interested in which linguists attended the panel (perhaps she is a department administrator keeping attendance records), she might conclude from B's response that no other linguists attended the panel but conclude nothing about computer scientists in this context.

(43) **A**: Who attended the last panel?
B: Agnes and Beatrice.
$\rightsquigarrow$ *no other linguists*

Propositions attributing the question-property to computer scientists are not relevant and thus ignored by exhaustivity. This particular example could involve ordinary contextual domain restriction but Schulz and van Rooij's relevance order is also useful for describing exhaustive implications that do not order alternatives by logical entailment. Suppose A and B are undergraduates who are primarily interested in attending the panels for their entertainment value. A wants to know whether the most engaging public speakers in the Linguistics and Computer Science departments have been showing up. If Agnes is the second-best public speaker, and Beatrice is the best, A might conclude from B's answer that Beatrice did not attend (i.e., nobody more entertaining than Agnes) but conclude nothing about the presence of less entertaining public speakers.

(44) **A**: Who attended the last panel?
B: Agnes.
$\rightsquigarrow$ *no better public speaker*

Hirschberg, 1985 claims that any partial order can support exhaustive interpretation, with entailment as a special case. Schulz and van Rooij's relevance order reduces to entailment in contexts when the conversational participants are interested in knowing the complete extension of the question-property. The analysis of exhaustivity I propose is developed within the same general setting and is therefore compatible with this solution.

An account of granularity that depends only on the prior interests of the conversational participants is not complete, however. Although B's answer in (42) (*Some of the linguists attended*) can implicate that not many linguists attended, it can get no more granular than this: it is impossible to interpret B's answer at a granularity level where alternatives are distinguished by exact quantity. The proposition that two linguists attended is stronger than B's answer, yet if exhaustive interpretation excluded this alternative, we would conclude from B's answer that exactly one linguist attended, which is impossible.

Theories that use formal alternatives can provide a straightforward account of this observation by stipulating that quantifiers count as alternatives to other quantifiers (*some, many, most, all*), but numerals do not count as alternatives to quantifiers. Theories based on circumscription do not have access to this account, and must instead claim that granularity levels are at least partially restricted by the utterance that was actually made. This claim has been developed most explicitly in the literature on numerals.

According to some theories of exhaustivity (Horn, 1972; Levinson, 2000), upper bound inferences on numerals like the one in (45) should be attributed to exhaustive interpretation. If this is correct, then it is a puzzle why the same inference is missing from (46), which ostensibly denotes the same proposition (Fox and Hackl, 2006; Krifka, 1999).

(45) **A**: Who attended the last panel?
B: Three linguists.
⇝ *four didn't*

(46) **A**: Who attended the last panel?
B: More than two linguists.
$\not\rightsquigarrow$ *four didn't*

Regardless of whether the upper bound in (45) is an exhaustive implication, which has been debated (Breheny, 2008; Carston, 1998; Geurts, 2006; Kennedy, 2015; Scharten, 1997), any theory of exhaustivity needs to explain why B's answer in (46) does not allow inferences about exact quantity. Theories based on formal alternatives can account for this observation by making stipulations about the alternatives to modified and unmodified numerals that render it impossible for implications excluding higher quantities to arise.

Cummins, Sauerland, and Solt, 2012 challenge this view, observing that answers like (46) still do allow upper bound inferences at coarser granularity levels. *More than two linguists attended* will never implicate that exactly three linguists attended but we might conclude that the number was less than five, or less than ten, or less than twenty, depending on granularity. Cummins, Sauerland, and Solt provide experimental evidence that participants systematically draw such inferences and argue that granularity depends on a speaker's choice of utterance. According to the pragmatic reasoning schema they propose, conveying that exactly three linguists attended via exhaustive interpretation of B's answer in (46) is suboptimal because it would have been easier for the speaker to convey this meaning using an unmodified numeral. Use of modified numerals like *more than two* and *at least three* thus serves as a signal that the speaker is assuming a coarser granularity level than one where alternatives are distinguished at the level of exact quantity. Exhaustivity is possible at this coarser granularity level, however.

This account could extend to (42): to convey a message about exact quantity, a speaker would not have used a quantificational determiner like *some* either. Linguistic form plays a role in this account, but it is an indirect role. The potential inferences in (42) and (46) are not missing because particular sentences are unavailable as alternatives. Rather, the form of a sentence is one of several cues regulating the granularity level that matters for exhaustive interpretation, which can itself be given a semantic characterization.

Exhaustive implications are also SUSPENDIBLE (Horn, 1972): it is possible to prevent an exhaustive implication from arising by indicating ignorance as to whether it holds. Suspension is different from cancellation, which involves the explicit rejection of exhaustive content. B's answer in (18) can convey that Agnes and Beatrice didn't both attend, but this implication is absent from (47) and (48). The puzzle is that all three answers have the same meaning. Suspender clauses thus appear to DETACH exhaustivity.

(18) **A**: Who attended the last panel?
B: Agnes or Beatrice attended.
⇝ *not both*

(47) **A**: Who attended the last panel?
B: Agnes or Beatrice, or both.
$\not\rightsquigarrow$ *not both*

(48) **A**: Who attended the last panel?
B: Agnes or Beatrice, if not both.
$\not\rightsquigarrow$ *not both*

Historically, suspension is one of the main reasons why formal alternatives were introduced. Gazdar, 1979 wanted an account of disjunctive and conditional suspender clauses like those in (47) and (48). Because B's answers in (18)-(48) denote the same proposition, Gazdar concluded that syntactic structure must play a role in implicature calculation and accordingly developed an analysis of exhaustivity as an operation on sentences.

However, formal alternatives are not the key factor in Gazdar's account of suspension. Rather, the suspender clauses in (47) and (48) are taken to trigger ignorance implicatures that take priority over exhaustive inferences in Gazdar's model of context update. This account is consistent with a semantic characterization of exhaustivity.

More recent theories of formal alternatives do not distinguish these sentences either. Sauerland, 2004 proposes that the formal alternatives to a disjunction include the corresponding conjunction as well as both disjuncts. This recipe predicts the alternatives in (49) for a sentence of the form $(A \vee B) \vee (A \wedge B)$.

(49)
a. $A \vee B$
b. $A, A \vee (A \wedge B)$
c. $B, B \vee (A \wedge B)$
d. $A \wedge B, (A \wedge B) \vee (A \wedge B), (A \vee B) \wedge (A \wedge B), A \wedge (A \wedge B), B \wedge (A \wedge B), (A \wedge B) \wedge (A \wedge B)$

These formulae are all equivalent to one of the following: $A \vee B$, A, B, $A \wedge B$, the same set of alternatives that Sauerland predicts for (18). The extra disjunctive clause in (47) does not actually add any new alternatives that could be leveraged to explain its role in suspension.

Theories that do not use formal alternatives can preserve nondetachability in the face of suspension by adopting a richer notion of content than classical static semantics makes available. Schulz and van Rooij, 2006 propose an analysis of suspension using dynamic semantics: the extra suspender clauses in (47) and (48) activate propositional discourse referents whose presence blocks the exclusion of these propositions. This strategy can also account for contrasts like those in (45) and (46) if the dynamic contributions of modified and unmodified numerals are analyzed appropriately. Fine, 2017 pursues a similar strategy using the notion of "exact verification" in truthmaker semantics. Because propositions that are true in the same possible worlds may still have different truthmakers, a theory of exhaustivity based on truthmakers can distinguish these sentences semantically.

Although this sort of hyperintensional solution is compatible with the proposal developed in this thesis, I would like to speculate about another possibility: perhaps a theory of granularity regulation can also provide the key to suspension. Support for this view comes from the fact that B's answers in (18), (47), and (48) all still allow the inference that nobody other than Agnes or Beatrice attended. Suspender clauses can thus prevent specific exhaustive inferences from arising, but do not completely suspend exhaustive interpretation. Similarly, although we cannot conclude from B's response in (50) that twenty linguists did not attend, we might still conclude that the number was less than thirty, or fifty. Schulz and van Rooij's proposal for dynamic exhaustivity would incorrectly block this inference.

(50) **A**: Who attended the last panel?
B: At least ten linguists, if not more than twenty.

If the reasoning schema suggested by Cummins, Sauerland, and Solt, 2012 were applied to these cases, it could potentially bring suspension under the larger pragmatic umbrella of granularity regulation. The idea would be that the explicit addition of a suspender clause serves as a signal that the speaker is assuming a granularity level that does not distinguish between the content of the suspender clause and the rest of the sentence. Many more details would need to be worked out to make this sketch of an account viable, including which constraints guide this reasoning, how they interact, and why the class of suspender clauses has the members it has (disjunctive and conditional clauses).

There is one other phenomenon that potentially challenges nondetachability. Marked or convoluted ways of expressing the same proposition do not typically give rise to the same implicatures, thanks to Grice's maxim of Manner (Grice, 1967; Rett, 2020). In particular, all positive answers can be paraphrased using double negation, yet the answers in (51) suggest different conclusions.

(51) **A**: Who attended the last panel?
  a. **B**: Beatrice attended.
  b. **B**: Beatrice didn't not attend.

A might conclude from the extra negation in (51-b) that B expected Beatrice not to attend, or that Beatrice was in the habit of not attending, or that Beatrice's failure to attend was somehow a relevant issue in the context, or that Beatrice's attendance was somehow marginal. These inferences are all missing from (51-a), and can be attributed to Grice's maxim of Manner, which enforces a preference against using marked linguistic forms like double negation. Conversely, it is less natural to conclude from (51-b) that nobody else attended aside from Beatrice—perhaps because the linguistic form of this answer serves as a signal that it is addressing a more specific issue than A's question.

Because negation is always marked, one might wonder whether similar reasoning could account for the phenomenon of restricted extent in general: perhaps negation serves as a signal to suspend exhaustivity. van Rooij and Schulz, 2004 suggest an account like this. However, we have seen that negative answers with restricted extent do still allow some exhaustive inferences, which are evidently not suspended in the same way.

By definition, Manner-driven implicatures are nondetachable. I suggest that in marked sentences like (51-b), Manner takes priority and the need to make sense of B's double negation outweighs the need to maximize information in this context. What drives the reasoning in (51-b) is the presence of double negation and the availability of a simpler form to express the same meaning, not the presence of negation per se. The contrast in (51) does not show that exhaustivity is detachable, but that competing detachable implicatures can motivate the suspension of exhaustivity.

The phenomena discussed in this section provide the main challenges to the claim that exhaustive implications are nondetachable. Any account that rejected nondetachability on the basis of suspension would then need to explain why exhaustive implications cannot be detached via substitution of synonymous lexical items, passivization, or in any other way. I will instead assume that suspender clauses form a principled exception to this generalization and set them aside, recognizing that the present approach will eventually need to account for them. Reaching a more developed understanding of granularity would be useful here and no doubt elsewhere.

1.1.6 Summary

This section described the main empirical phenomena that a theory of exhaustivity should explain, which include cancellation, nondetachability, direction, extent, sensitivity to information structure, sensitivity to granularity, and suspension. The phenomena that most directly speak to how alternatives should be characterized are nondetachability, direction, extent, and sensitivity to information structure. Explanations for cancellation and granularity typically lie in how alternatives are used rather than how they are characterized. Although suspender clauses appear to provide a restricted class of counterexamples to nondetachability, the right analysis of suspension will likely also leverage the use of alternatives, rather than their characterization.

The next two sections provide a detailed introduction to existing theories of alternatives, starting with formal alternatives and then turning to circumscription. The choice between a syntactic and a semantic characterization of alternatives has consequences for how we analyze the facts about direction, extent and nondetachability. I discuss these consequences and argue in favor of circumscription.

1.2 Formal alternatives

Theories of FORMAL ALTERNATIVES start from the premise that propositions can only count as alternatives to a sentence if they are expressible by sentences that stand in a particular syntactic relation to the original sentence. The theorist's task is then to provide an appropriate characterization of this syntactic relation so that it applies to sentences whose negations express possible exhaustive implications, and only these.

This section argues that the syntactic status of alternatives in these theories either does no work or is actively harmful. In particular, the specific syntactic criteria that existing theories have adopted directly prevent a successful account of direction and extent. These problems can be avoided if alternatives are not sentences.

The hypothesis that alternatives are sentences is very general. To conclusively falsify it, we would probably need to identify exhaustive inferences whose content cannot be paraphrased using any sentence of natural language. Given the productivity of language, I would be surprised to find any such cases. The following discussion argues only that existing proposals do not isolate the right class of sentences.

It is worth making a methodological point. To analyze an exhaustive implication, theories based on formal alternatives typically proceed as follows. First, the content of the implication is expressed as the negation of some proposition. This proposition is then paraphrased as a sentence of natural language that is presumed to be the unique alternative whose exclusion underlies the exhaustive implication in question. If a theory of formal alternatives can relate this sentential paraphrase to the sentence that was actually uttered, the theory's predictions are considered successful.

There is not generally a bijective correspondence between propositions and sentences, however, so this methodology leaves the theorist impressive latitude. For instance, the inference from (52) that not all of the linguists attended the panel is usually modeled as the exclusion of the alternative (52-a), which is synonymous with the other (52) sentences.

(52) **A**: Who attended the last panel?
B: Most of the linguists attended. $\rightsquigarrow$ *not all*
a. All of the linguists attended.

b. All linguists attended.
c. Every linguist attended.
d. Every single linguist attended.
e. Every one of the linguists attended.
f. Each linguist attended.
g. Each one of the linguists attended.
h. Everyone among the linguists attended.
i. Every person among the linguists attended.
j. Every linguist among the linguists attended.
k. Every linguist is a linguist who attended.
l. Every linguist was in attendance.
m. Every linguist was present.
n. The linguists all attended.
o. The linguists each attended.
p. The panel was attended by all of the linguists.
q. No linguist didn't attend.
r. None of the linguists didn't attend.
s. None of the linguists are linguists who didn't attend.
t. It's not the case that not all of the linguists attended.
u. No linguist was absent.
v. None of the linguists were absent.
w. For every linguist, that linguist attended.
x. For every person, if that person is a linguist, then that person attended.
y. All people are such that they attended if they are a linguist.
z. Everything is such that it attended if it is a linguist.

Which of these sentences is the "real" alternative? Is there only one? Do some of these sentences count, and not others? How would we know?

Because alternatives are a theoretical tool, not an observable phenomenon, there are no empirical tests that could answer these questions. Theories based on formal alternatives implicitly assume that there is no problem as long as at least one sentential paraphrase counts as an alternative. In what follows, I grant this assumption, for the sake of argument.

1.2.1 Scales

Formal alternatives were introduced to the literature in the guise of LEXICAL SCALES, sometimes called HORN SCALES after Horn, 1972. A lexical scale is a set of expressions that generate an entailment order, or some other informativity order, over otherwise identical sentences containing them. The set ⟨*allowed*, *required*⟩ forms a lexical scale because the sentence *You're required to wear sneakers* entails *You're allowed to wear sneakers*; the set ⟨*good*, *great*, *excellent*⟩ forms a lexical scale because the sentence *Ivy Compton-Burnett's novels are excellent* entails *Ivy Compton-Burnett's novels are great*, which entails *Ivy Compton-Burnett's novels are good*; and so on. Horn, 1972 introduced scales in order to describe exhaustive implications, but did not commit to a specific hypothesis regarding their grammatical status.

Gazdar, 1979 develops a formalization of scales that is explicit about the assumption that alternatives are sentences. Gazdar defines EXPRESSION ALTERNATIVES as a relation on sentences that differ only in the substitution of scalar expressions (*You're* {*allowed*/*required*} *to wear sneakers*). Exhaustive implications are then generated on the basis of this class of sentences: a speaker who utters a particular sentence will implicate that she knows that all of the sentence's alternatives containing stronger scalar expressions are false.

What is the grammatical status of scales in this system? Gazdar, pg. 56 comments that "It is both in the spirit of Grice's program and in the interests of economy to read these nonconventional inferences from the semantic representation. Presumably, at such a level a set of expressions such as {*perhaps, maybe, possibly*} will be represented by just one item for the reading they have in common." These remarks suggest that scales are not literally sets of expressions, but sets of EQUIVALENCE CLASSES of expressions, where each equivalence class includes synonymous expressions. As long as two expressions make the same contribution to "semantic representation," i.e., logical form, they share an equivalence class.

The theory of alternatives that emerges views alternatives as linguistic expressions, but determines scale membership via semantic criteria. Nondetachability is weakened as a consequence, but not lost. Although equivalent sentences with different logical forms are no longer guaranteed to license the same inferences, nondetachability still holds relative to logical form. Especially if the theorist is allowed some leeway in selecting appropriate logical forms for sentences, this is a weak position.

Gazdar's view of scale membership as determined by logical form has another consequence. The syntactic component of his theory does no work. There is never any instance of an exhaustive inference that is blocked because the requisite scalar expression does not exist in a given language, or leads to ungrammaticality when substituted in a particular position. Buccola, Križ, and Chemla, 2022 consider a number of cases where this situation could potentially arise, and argue in each case that the expected inference is still attested, concluding that expressibility is not necessary for alternative status (see also Charlow, 2019; Dionne and Coppock, 2022; Smith, 2020). For this reason, Gazdar's theory could be recast entirely in terms of logical form. As long as we can access a sentence's logical form and compare it to alternative formulae determined by the class of scales in the language, the appeal to "expression alternatives" could be eliminated without affecting empirical coverage.

Lexical scales as developed by Horn and Gazdar provide a descriptive solution to the symmetry problem. Symmetry is broken whenever a lexical scale exists that supports exhaustive inferences in one direction but not the other. This strategy is a way of labeling the facts without explaining them, as Gazdar, pg. 58 admits when he suggests that "scales are, in some sense, 'given to us"'.

Subsequent literature in the neo-Gricean tradition has attempted to give the class of scales a more principled characterization. Atlas and Levinson, 1981 claim that scalar items must share selectional restrictions, belong to the same "semantic field", and be lexicalized to the same degree, although these conditions on linguistic expressions are not clearly compatible with Gazdar's view of scale membership (which Atlas and Levinson endorse) as determined by logical form.

Hirschberg, 1985 argues that any expression denoting a value in a pragmatically salient partial order can appear on a scale. She surveys a wide range of naturally occurring answers in dialogue that license exhaustive implications over contextually available nonentailment orders, and shows that exhaustivity is unavailable when the defining properties of a partial order (transitivity and antisymmetry) do not hold. Hirschberg suggests that this observation follows from the basic nature of exhaustivity, since partial order is the minimal requirement for an informativity-maximizing operation to be definable. Matsumoto, 1995 provides an information-structural account of the conditions that can make a scale pragmatically salient: to serve as an alternative, a proposition must address a relevant contextual ISSUE. On this view, scales cannot be predicted in advance, but must be constructed in conversation. Entailment orders are at most a default.

However, as Hirschberg and Matsumoto observe, there are some partial orders that do not support exhaustive implications. The dual of entailment is one such order: exhaustivity cannot target weaker propositions. This restriction is easy to motivate (don't contradict the assertion) but other ordering restrictions do not follow on obvious pragmatic grounds. In the question/answer exchanges I discussed in the previous section, certain symmetry-violating entailment orders must be systematically unavailable if we want to make correct predictions about the direction of exhaustivity.

Matsumoto acknowledges that his pragmatic account does not derive the direction facts, which he calls the "problem of scalarity" (pg. 57), and follows Horn, 1989 in claiming that only monotone quantifiers can appear on scales. This stipulation correctly resolves the symmetry problem in a number of cases: *most* has *all* as an alternative but not *most but not all*; *allowed* has *required* as an alternative but not *optional* (or *allowed but not required*). Monotonicity cannot be the whole story, however, as we have seen that nonmonotonic answers also allow exhaustive implications.

Despite their pragmatic view of how scales are constructed, both Hirschberg, 1985 and Matsumoto, 1995 assume with Gazdar, 1979 that alternatives are sentences. This assumption does not play a significant role in either proposal, and in general the syntactic status of alternatives in neo-Gricean pragmatics is empirically idle.

1.2.2 The structural theory

In an influential paper, Katzir, 2007 develops a syntactic algorithm intended to replace lexical scales. Alternatives are generated on the basis of the original sentence using tree-pruning operations. The syntactic status of alternatives in this theory is not idle; it is the main driver of empirical predictions.

A productive line of subsequent work has revised and elaborated Katzir's proposal (Buccola, Križ, and Chemla, 2022; Fox and Katzir, 2011; Haslinger and Schmitt, 2025; Trinh, 2018; Trinh and Haida, 2015), which has also been challenged (Breheny et al., 2018; Feinmann, 2023, 2025; Hirsch and Schwarz, 2025; Romoli, 2013; Schwarz and Wagner, 2024; Swanson, 2010, 2017; Westera, 2017a). This work represents the most developed attempt to address the symmetry problem in the contemporary literature. I will call the hypothesis that formal alternatives are built using Katzir's algorithm the STRUCTURAL THEORY.

Instead of generating formal alternatives by replacing scalar expressions with their scale-mates, Katzir proposes that any constituent can be replaced by any other constituent that is equally or less complex. The syntactic algorithm is defined as follows.

(53) **Substitution source.**
If X is a tree, then X's *substitution source* $L(X)$ is the union of the lexicon of the language with the set of all subtrees of X.

(54) **Complexity**.
Iff we can transform a tree X into a tree Y by a finite series of deletions, contractions, and replacements of constituents in X by constituents of the same category taken from $L(X)$, then $Y \lesssim X$ ("Y is at most as complex as X").[2]

(55) **Structural alternatives**.
The set of *structural alternatives* for a tree X is $\{Y | Y \lesssim X\}$.

[2]Deletions, contractions, and replacements are implicitly defined as in graph theory (West, 2001).

Exhaustive interpretation of a sentence only considers its STRUCTURAL ALTERNATIVES, and thus amounts on this view to the supposition that there is no true sentence that is at most as complex as, and strictly stronger than, the one that was uttered.

The structural theory is more general than an approach based on lexical scales. Because any substitution is allowed modulo complexity, canonical lexical scales are all subsumed under this condition. The appeal to complexity requires a shift in the level of representation that matters for constructing alternatives, however. Expressions that make the same contribution to logical form are no longer guaranteed to be interchangeable. Two synonymous expressions might not be equally complex, with the result that only the less complex one counts as an alternative to some other expression. Idiosyncratic formal properties of particular linguistic expressions are then expected to affect alternative status, and should potentially make a difference in whether an exhaustive implication is available.

The structural theory does not make explicit predictions without auxiliary syntactic assumptions about constituency and labeling. To generate a sentence's alternatives, we first need to commit to a specific hypothesis about that sentence's constituent structure and assign category labels to each constituent before determining which replacement operations can be made consistently with (54). In the face of apparent counterexamples, it is often possible to preserve the structural theory while modifying the syntax in which it is grounded. Hence the structural theory is difficult to falsify, although I will make my best attempt.

The role of syntax in this theory is to address the symmetry problem. In (56) and (57), the structural hypothesis generates the (a) sentence but not the (b) sentence as an alternative to B's answer, because the (b) sentences are more complex. This is a good result because exhaustive interpretation of B's answer conveys that the (a) sentence is false, not the (b) sentence. In this way, the structural theory provides a unified account of symmetry breaking without stipulating separate scales for connectives, quantifiers, and so on.

(56) **A**: Who attended the last panel?
B: Agnes or Beatrice. ⇝ *not both*
a. Agnes and Beatrice.
b. Agnes or Beatrice but not both.

(57) **A**: Who attended the last panel?
B: Some of the linguists. ⇝ *not all*
a. All of the linguists.
b. Some but not all of the linguists.

Katzir, 2007 also claims that some exhaustive implications require formal alternatives that cannot be generated using scales. The evidence comes from examples like (58) and (59). I have modified Katzir's examples for the question/answer format.

(58) **A**: Who attended the panel?
B: Everyone who understands generative syntax but not deep learning.
⇝ *some people who understand both didn't attend*

(59) **A**: Who attended the panel?
B: If Agnes but not Beatrice attended, then Charmian attended.
⇝ *Charmian only attended if Agnes did and Beatrice didn't*

To model the inference in (58) that some people who understand syntax as well as deep learning didn't attend, Katzir's algorithm would use the formal alternative *Everyone who* *understands generative syntax*. Excluding the proposition that everyone who understands generative syntax attended the panel together with the meaning of B's answer accounts for the attested inference that it is possible for some syntacticians not to have attended, as long as they understand deep learning. To model the conditional perfection inference in (59) that Charmian only attended if Agnes did and Beatrice didn't, Katzir's algorithm would use the formal alternative *If Agnes attended, then Charmian attended*. Excluding this alternative accounts for the inference that Charmian's attendance is conditional not only on the presence of Agnes but the absence of Beatrice.

However, if the claims from Horn, 1989 and Matsumoto, 1995 are maintained that only monotonic quantifiers can appear on scales, then these alternative sentences cannot be generated because we would have to replace complex nonmonotonic coordinations like *generative syntax but not deep learning* and *Agnes but not Beatrice* with their simpler monotonic conjuncts *generative syntax* and *Agnes*.[3] This should not be possible if only substitutions of scalar expressions are allowed, and only monotonic quantifiers appear on scales. Katzir concludes that there is no monotonicity condition restricting the construction of formal alternatives. This means that monotonicity cannot be the relevant factor behind symmetry breaking in examples like (56) and (57). The structural theory can generate the formal alternatives needed in (58) and (59) (*Everyone who understands generative syntax*, *If Agnes attended, then Charmian attended*), while still breaking symmetry in (56) and (57).

It is important to understand what this argument establishes, and what it does not establish. Katzir, 2007 is sometimes inaccurately cited as arguing for a syntactic characterization of alternatives over a semantic one. However, both the structural theory and a recipe based on lexical scales provide a syntactic characterization of alternatives. They differ only in whether the class of available syntactic substitutions is restricted using syntactic or semantic criteria. Because Katzir presupposes that alternatives are sentences, it is no wonder that his investigation leads to a syntactic conclusion.

Fox and Katzir, 2011 propose a revision of Katzir's algorithm that allows only substitutions of focused constituents, rendering the algorithm sensitive to information structure.

(60) **A**: Who was present?
B: $[\text{Agnes}]_F$ was present.
$\rightsquigarrow$ *Beatrice was absent*

The implication in (60) that Beatrice wasn't present can be modeled as the exclusion of the formal alternative *Beatrice was present*, which can be constructed by replacing the focused subject in B's answer with *Beatrice*. Because substitutions of backgrounded constituents are disallowed, we do not have to worry about the symmetric alternative *Beatrice was absent*, which could otherwise have been constructed.

1.2.3 Problems with the structural theory

This section shows that the structural theory systematically fails to capture the direction and extent of exhaustivity. These empirical problems stem from the complexity metric used to break symmetry. The structural theory also does not provide a principled explanation

[3] In the case of (58), this is not a strong argument. Instead of replacing *generative syntax but not deep learning* with *generative syntax* inside the universal quantifier's restriction a proponent of scales could make the substitution at the level of the whole quantifier, replacing *Everyone who understands generative syntax but not deep learning* with *Everyone who understands generative syntax*. These are both increasing quantifiers, so there is no problem. This is not the solution I will advocate, however.

for nondetachability, a problem traceable to the underlying conception of alternatives as syntactic objects. Together, these problems indicate that not only is complexity the wrong metric, the identification of alternatives with sentences cannot be maintained.

1.2.3.1 Direction

The structural theory's solution to the symmetry problem hinges on the syntactic complexity of negation. The structural theory can explain why positive answers have a negative direction of exhaustivity: because positive answers do not include negation, the hypothetical negatively restricted, symmetric formal alternatives that would block negative exhaustivity cannot be generated. For the same reason, however, because deleting negation does not violate complexity, when an answer already includes negation it is possible to generate symmetric alternatives that do and do not include negation. The structural theory therefore cannot explain the directional properties of negative or mixed answers.

In (60), the inference that Beatrice didn't attend can be modeled by replacing *Agnes* with *Beatrice* and excluding this formal alternative. The symmetric alternative *Beatrice didn't attend* (or *not Beatrice*) cannot be generated without increasing complexity. The same result holds for all individuals in the domain of quantification, so B's answer is correctly predicted to have negative direction.[4]

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

[4]Actually, whether this result holds in the general case depends on what we say about the syntax of names, which can vary in complexity.

(61) **A**: Which *Sex and the City* cast members have you met?
B: Kim Cattrall.
⇝ *no other cast members*

B's answer in (61) implicates that she hasn't met Sarah Jessica Parker, whose middle name apparently does not affect complexity. To generate this alternative, the structural theory would have to stipulate that all names are equally complex, perhaps because all names are lexical items, a conclusion inconsistent with the syntactic literature on names (Acquaviva, 2019; Matushansky, 2008).

Titles pose a similar problem. For instance, the singer Fiona Apple has released five albums, whose titles are listed in (62).

(62) a. *Tidal*
b. *When the pawn hits the conflicts he thinks like a king What he knows throws the blows when he goes to the fight And he'll win the whole thing 'fore he enters the ring There's no body to batter when your mind is your might So when you go solo, you hold your own hand And remember that depth is the greatest of heights And if you know where you stand, then you know where to land And if you fall it won't matter, cuz you'll know that you're right*
c. *Extraordinary Machine*
d. *The Idler Wheel Is Wiser Than the Driver of the Screw and Whipping Cords Will Serve You More Than Ropes Will Ever Do*
e. *Fetch the Bolt Cutters*

(63) **A**: Which Fiona Apple albums do you own on vinyl?
B: *Fetch the Bolt Cutters*.
⇝ *none of the others*

Without further stipulation, the structural theory predicts that B's answer in (63) should implicate that she doesn't own *Tidal* or *Extraordinary Machine* on vinyl, while conveying nothing about the other two albums.

In (2), the implication that not all of the linguists attended can be modeled by replacing *most* with *all* and excluding this formal alternative (*All of the linguists*). The symmetric alternative *Most but not all of the linguists* cannot be generated without increasing complexity, and there is no lexicalized quantifier in English that could convey this meaning, so symmetry is avoided.

(2) **A**: Who attended the last panel?
B: Most of the linguists.
⇝ *not all linguists*

In (3), the implication that sneakers aren't required can be modeled by replacing *allowed* with *required* and excluding this formal alternative (*Sneakers are required*). The symmetric alternative *You're allowed but not required to wear sneakers* cannot be generated without increasing complexity.

(3) **A**: What's the dress code for the panel?
B: Sneakers are allowed.
⇝ *sneakers aren't required*

However, Swanson, 2010 observes that *allowed but not required* has in fact been lexicalized in English as *optional* (also: *elective*, *discretionary*, etc.). The structural theory allows *Sneakers are optional* as an alternative to B's answer in (3), which should block the attested inference. Swanson's observation shows that the complexity of negation does not always break symmetry even when a sentence does not already include negation. We also have to check if the meaning of an unwanted alternative can be expressed in any other way.

Romoli, 2013 observes that the structural theory does not explain why negative answers have a positive direction of exhaustivity (see also Breheny et al., 2018; Schwarz and Wagner, 2024; Trinh, 2018; Trinh and Haida, 2015). In (64), the inference that some linguists attended can be modeled by replacing *all* with *any* (or *not all* with *none*) and excluding this formal alternative (*None of the linguists attended*). However, Katzir's algorithm can also replace *all* with *some* and delete negation to derive the formal alternative *Some of the linguists attended*, which stands in a symmetric relationship with the alternative we want to derive. This latter sentence should not count as a formal alternative if we want to account for positive direction, yet the structural theory allows it. This kind of situation arises whenever an answer includes negation. I will call it NEGATION-INDUCED SYMMETRY.

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

The implication in (19) that Agnes or Beatrice attended can be modeled by deleting *both*, replacing *and* with *or*, and excluding this formal alternative (*Not Agnes or Beatrice*). Katzir's algorithm also derives the alternative *Agnes or Beatrice*, however, which should prevent the attested implication. This is another instance of negation-induced symmetry.

(19) **A**: Who attended the last panel?
B: Not both Agnes and Beatrice.
⇝ *one of them did*

The same problem is observed in (20): the implication that sneakers are allowed can be modeled by replacing *required* with *allowed* and excluding this formal alternative (*Sneakers aren't allowed*), but by additionally deleting negation Katzir's algorithm also derives the symmetric *Sneakers are allowed*, so again we see negation-induced symmetry.

(20) **A**: What's the dress code for the panel?
B: Sneakers aren't required.
⇝ *sneakers are allowed*

To account for positive direction, we want to exclude negatively restricted alternatives, but because these are more complex than their positive counterparts the structural theory has no way to isolate them.

The structural theory also faces negation-induced symmetry with mixed answers. The implication in (21) that not all linguists and some computer scientists attended can be modeled by replacing *most* with *all* and replacing *not a lot of* with *no* or *not any*, and excluding this formal alternative (*All of the linguists but not computer scientists*). Because Katzir's algorithm can also replace *not a lot of* with *some*, the symmetric alternative *All of the linguists but some computer scientists* (as well as *Most of the linguists but some computer scientists*) again prevents this implication from arising.

(21) **A**: Who attended the last panel?
B: Most of the linguists, but not a lot of computer scientists.
⇝ *not all linguists*
⇝ *some computer scientists*

Other mixed answers appear to require formal alternatives that are strictly more complex than the original sentence. Conditional answers are such a case. Conditional strengthening can be analyzed by excluding conditional alternatives: there are no sufficient conditions for the consequent other than the antecedent (van der Auwera, 1997). Key motivation for Katzir, 2007 came from conditional answers like (59). The inference that Charmian only attended if Agnes did and Beatrice didn't can be modeled by respectively deleting the first and the second conjunct in the antecedent to derive two conditional alternatives, *If Agnes attended then Charmian attended* and *If Beatrice didn't attend then Charmian attended* (I assume for the sake of argument the latter can be derived from the original sentence, which contains term negation instead of VP negation), and excluding these. Thus we can conclude that Agnes attending is insufficient, and Beatrice not attending is insufficient.

(59) **A**: Who attended the panel?
B: If Agnes but not Beatrice attended, then Charmian attended.
⇝ *Charmian only attended if Agnes did and Beatrice didn't*

However, conditional strengthening can license more inferences than these. The literal meaning of B's answer is consistent with a world in which there are two sufficient conditions for Charmian to attend: either she will attend if Agnes attends and Beatrice doesn't, or she will attend if Doris and Ethel attend but Felicity doesn't, or both. In some contexts, exhaustive interpretation of (59) is not consistent with this world. We presumably want to exclude the formal alternative *If Doris and Ethel but not Felicity attended, then Charmian attended*, but because the antecedent of this conditional includes more conjuncts than the antecedent of the original sentence it cannot be generated by Katzir's algorithm.

1.2.3.2 Extent

The directional properties of answers and the challenges they pose for the structural theory have been discussed extensively. Less attention has been paid to extent. The structural theory can explain why positive answers have unrestricted extent and negative answers have restricted extent by appealing to negation-induced symmetry. However, mixed answers vary in extent and the structural theory does not make the cut in the right place.

The structural theory explains why B's answer in (1) has unrestricted extent for the same reason that it has negative direction: *Beatrice* (*Charmian*, etc.) is equally complex and therefore counts as a formal alternative, while *not Beatrice* is more complex and ruled out as a formal alternative.

(1) **A**: Who attended the last panel?
B: Agnes.
$\rightsquigarrow$ *nobody but Agnes*

The structural theory can also explain why B's answer in (2) has unrestricted extent. The implication that no computer scientists attended can be modeled by replacing *linguists* with *computer scientists*, replacing *most* with *some*, and excluding this formal alternative (*Some of the computer scientists*).

(2) **A**: Who attended the last panel?
B: Most of the linguists.
$\rightsquigarrow$ *no computer scientists*

There are several caveats here. The nominal substitution required to generate this formal alternative is only allowed if *linguists* and *computer scientists* are equally complex, which is superficially not the case. Potentially, the structural theory could claim that *computer scientists* is idiomatic, noncompositional, and part of the lexicon, hence included in the substitution source.

Additionally, to successfully exclude *Some of the computer scientists*, we also have to make sure that *None of the computer scientists* is not a symmetric formal alternative. If we could replace *most* with *none*, the attested inference would not be predicted. One way to prevent this substitution would be to decompose *none* into *not* + *any* (Jacobs, 1980; Penka, 2011) and situate the generation of formal alternatives at the level of syntactic representation where this decomposition takes place. To avoid generating other negatively restricted alternatives, this move would have to apply to all negative quantifiers (*few* = *not many*, and so on).

Regardless of whether these are plausible moves, they highlight a weakness of the structural theory. It commits us to very specific syntactic assumptions that should ideally be grounded in syntactic evidence, not forced upon us by an analysis of orthogonal semantic and pragmatic phenomena.

The structural theory makes correct predictions about the restricted extent of negative answers. Katzir's algorithm generates both *Agnes* and *not Agnes* as alternatives to B's answer in (23), preventing exhaustive implications about Agnes from arising in either direction. The same result holds for other individuals in the domain of quantification.

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

Katzir's algorithm likewise generates *Some computer scientists* and *No computer scientists* as alternatives to B's answer in (15), preventing exhaustive implications about computer scientists from arising in either direction.

(15) **A**: Who attended the last panel?
B: Not a lot of linguists.
$\not\rightsquigarrow$ *conclude nothing about computer scientists*

Restricted extent is thus correctly predicted by the same feature of the structural theory, negation-induced symmetry, that prevented a successful account of positive direction.

Mixed answers vary in extent. The success of the structural theory in accounting for the extent properties of positive and negative answers in terms of negation-induced symmetry leads to the expectation that mixed answers might come in two classes, those that exhibit negation-induced symmetry and those that don't, with their extent properties attributed to this syntactic difference. However, the presence of negation does not uniformly lead to restricted extent.

The structural theory makes correct predictions about the restricted extent of B's answer in (24). Katzir's algorithm generates *Some computer scientists attended* as an alternative by deleting conjunction and the second conjunct, replacing *syntacticians* with *computer scientists* and replacing *most* with *some*. It also generates *No computer scientists attended* as an alternative by deleting conjunction and the first conjunct, replacing *phonologists* with *computer scientists*, and replacing *a lot of* with *any*. Thus negation-induced symmetry prevents exhaustive implications about computer scientists from arising in either direction. The same result holds for other properties of individuals (*sociolinguists*, etc.).

(24) **A**: Who attended the last panel?
B: Most syntacticians, but not a lot of phonologists.
$\not\rightsquigarrow$ *conclude nothing about other linguists or computer scientists*

The structural theory cannot explain why B's answer in (25) has unrestricted extent, however. The implication that no computer scientists attended can be generated by deleting conjunction and the second conjunct (*not all*), replacing *linguists* with *computer scientists*, and excluding this formal alternative (*Some of the computer scientists*). Katzir's algorithm also generates *None of the computer scientists* as an alternative by deleting conjunction and the first conjunct (*some*) and replacing *not all* with *none* (or *all* with *any*). This latter sentence should apparently not count as a formal alternative given the attested exhaustive interpretation of B's answer, but there is no way to rule it out using complexity.

(25) **A**: Who attended the last panel?
B: Some but not all of the linguists.
$\rightsquigarrow$ *no computer scientists*

Although both (24) and (25) include negation, they differ in extent, indicating that an answer's extent properties are independent of negation-induced symmetry.

1.2.3.3 Revising the structural theory

The structural theory does not provide a general solution to the symmetry problem. Although it correctly predicts the negative directional properties of positive answers as long as particular auxiliary syntactic assumptions are made, the problem of negation-induced symmetry prevents it from predicting the positive directional properties of negative answers or mixed answers. It does correctly predict the unrestricted extent of positive answers and the restricted extent of negative answers, but this solution does not generalize to mixed answers, whose extent properties are independent of negation.

Moreover, the explanation for restricted extent in terms of negation-induced symmetry is suspicious, since it is precisely this feature of the structural theory that prevents an account of positive direction. One might wonder whether the positive direction and restricted extent of negative answers can be simultaneously accounted for within the structural theory, or whether a solution for positive direction would sacrifice the account of restricted extent that is currently available.

More concretely, to predict the positive direction and restricted extent of an answer like *Not all of the linguists attended*, the structural theory must generate *None of the linguists attended* without generating *Some of the linguists attended*; at the same time, all alternatives regarding nonlinguists must have symmetric counterparts.

The proposal by Katzir, 2007 has been revised in subsequent literature, much of it intended to resolve the problems posed by the directional properties of answers. Trinh and Haida, 2015 propose a constraint on the derivation of formal alternatives that does successfully account for positive direction. Extent has not played a starring role in this discussion, and in this section I show that Trinh and Haida's proposal comes at the cost of restricted extent. The point is not to criticize this specific proposal, which is discussed in more detail by Breheny et al., 2018 and Trinh, 2018, but to make the more general point that the structural theory cannot account for positive direction and restricted extent simultaneously.

There are two crucial components to Trinh and Haida's proposal. The first is a simplification of Katzir's algorithm suggested by Fox and Katzir, 2011, according to which the only syntactic operation involved in the derivation of formal alternatives is replacement; deletions and contractions on trees are no longer allowed. Deletion is then implemented as replacement of a larger constituent with one of its subconstituents, or replacement of a terminal node with a null exponent.

The second component is a novel constraint on tree-pruning, ATOMICITY, which rules out further combinatorial operations on constituents that have been substituted for other constituents. The intuition is that constituents in the substitution source are treated as "syntactically atomic", so the alternative-building algorithm cannot see their internal structure.

To get a sense of how atomicity works, imagine a propositional language with three literals a, b, c, conjunction as the only connective, and binary branching syntax. The substitution source for the formula $a \wedge (b \wedge c)$ is the union of its subconstituents with the lexicon of the language: $a, b, c, b \wedge c, a \wedge (b \wedge c)$. Replacing each terminal node with the root of the tree gives us $a \wedge (b \wedge c) \wedge (a \wedge (b \wedge c) \wedge a \wedge (b \wedge c))$. If atomicity is adopted, this is the longest formal alternative we can generate. Otherwise, the possibility of substituting complex formulae into the new terminal nodes that are now available means that Katzir's algorithm can generate formal alternatives of arbitrary length. In this way, atomicity significantly reduces the number of formal alternatives for a given sentence.

If atomicity is adopted, it is possible to avoid negation-induced symmetry with negative answers. This allows a successful account of positive direction. Trinh and Haida assume the syntax in (65) for a sentence like *Not all linguists attended*, with negation in a high position projecting a Neg(ation)P(hrase) that embeds a V(erb)P(hrase), which includes the subject. It is possible to generate *Not any linguists attended* (or *No linguists attended*) by substituting *any* for *all*, as shown in (65). This is the formal alternative we want to exclude in order to model the exhaustive implication that some linguists attended.

(65) *all* → *any*

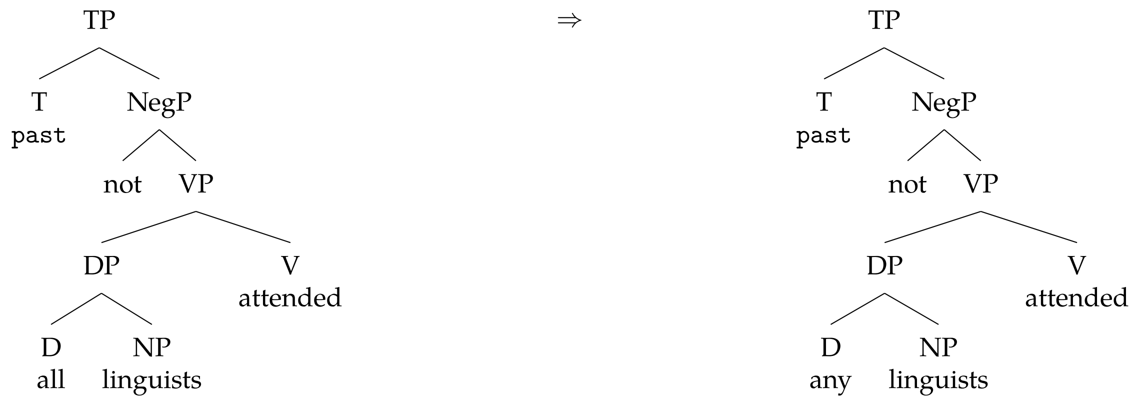

Atomicity prevents the generation of *Some linguists attended* as a formal alternative. Negation can be deleted by replacing NegP with VP, because VP is a subconstituent of the original tree and therefore included in the substitution source. However, atomicity renders the internal structure of VP inaccessible after this substitution step, so it is not subsequently possible to replace *all* with *some*.

(66) NegP → VP

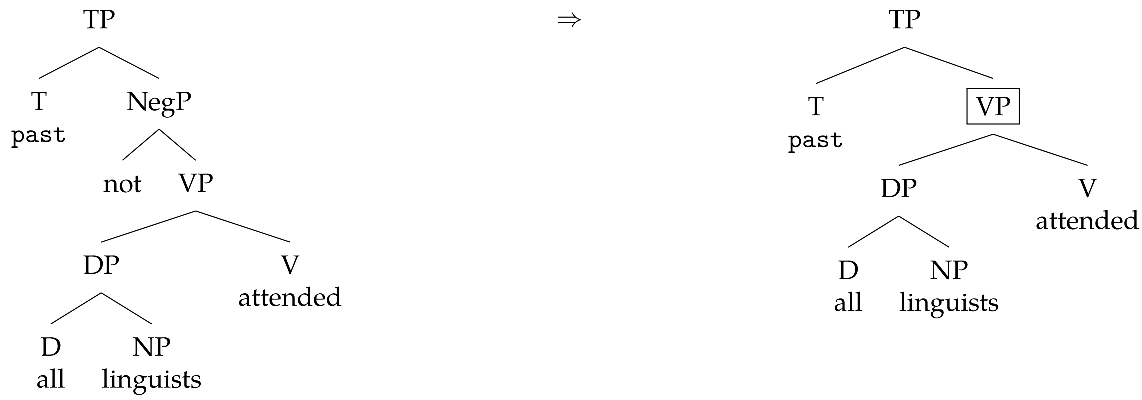

It is also not possible to substitute *some* for *all* before deleting negation. This deletion would involve replacing NegP with the derived VP *some linguists attended*, but because this derived constituent is not a member of the substitution source, the algorithm does not allow this substitution; a complex constituent is only included in the substitution source if it is a subconstituent of the original tree.

Because *No(t any) linguists attended* is a formal alternative and *Some linguists attended* is not, the former can be excluded and negation-induced symmetry is avoided. Augmented with atomicity, the structural theory is then able to predict the exhaustive implication that some linguists attended, overcoming the problem of positive direction.

Exactly the same feature of this proposal that successfully explains positive direction leads to wrong predictions about extent, however. A sentence like *Not all linguists attended* will have *Not all computer scientists attended* as a formal alternative without the symmetric *All computer scientists attended* necessary to block the exclusion of this alternative. The former can be generated by replacing *linguists* with *computer scientists*, as in (67).

(67) *linguists* → *computer scientists*

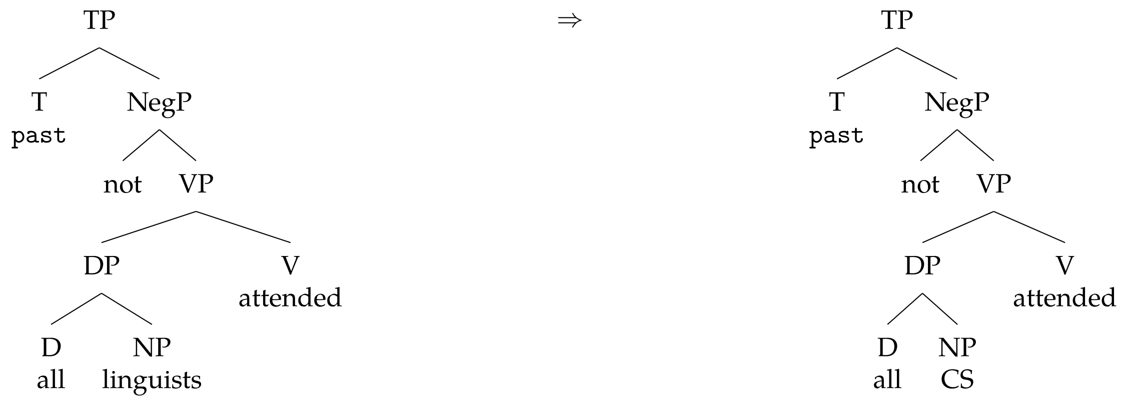

By contrast, *All computer scientists attended* cannot be generated as an alternative. We would have to delete negation and replace *linguists* with *computer scientists*, in either order. If negation is deleted first, as in (66), then atomicity prevents subsequent substitutions inside the VP. If we replace *linguists* with *computer scientists* first, then it is no longer possible to delete negation because the derived VP is not a member of the substitution source.

Trinh and Haida's proposal then predicts that B's answer in an exchange like *A: Who attended? B: Not all linguists attended* should have the formal alternative *Not all computer scientists attended* without a symmetric alternative to cancel it out. Excluding this alternative would lead to the inference that every computer scientist attended. This inference is not attested. For the same reason, B's answer in an exchange like *A: Who attended? B: Not Beatrice* is predicted to have *not Agnes*, *not Charmian*, and so on as formal alternatives but not their symmetric counterparts and should therefore implicate that everyone but Beatrice attended, which is impossible.

Without atomicity, the structural theory can explain restricted extent but not positive direction. With atomicity, the structural theory can explain positive direction but not restricted extent. The structural theory's primary strategy for addressing the symmetry problem is to regulate the distribution of negation, but negation must be regulated in contradictory ways to capture restricted extent: negative answers are either predicted to have unrestricted extent or no extent at all.

The problem is not specific to Trinh and Haida's proposal. Haslinger and Schmitt, 2025 propose that what matters for symmetry breaking is not overall complexity, but similarity to the sentence that was uttered, measured by the number of combinatorial operations that must apply to generate an alternative using Katzir's algorithm. Because negatively restricted alternatives are more similar to a sentence that already contains negation, this theory breaks symmetry in their favor and predicts the positive direction of negative answers, but makes wrong predictions about extent for the same reason.

The distinctions between alternatives we need to draw in order to get extent right are not distinctions a syntactic characterization of alternatives is capable of making. I suggested that extent can be characterized in terms of whether an answer allows exhaustivity to cover the whole domain of discourse, or only individuals that the answer is "about". The intuitive notion of aboutness can be made precise using the notion of quantificational restriction. The RESTRICTION of the quantifier *not all linguists* is the set containing the linguists (Barwise and Cooper, 1981; von Stechow and Zimmermann, 1984). Extent could then be analyzed as follows: answers with unrestricted extent allow exhaustivity to cover individuals outside their quantificational restriction, and answers with restricted extent only allow exhaustive inferences about individuals inside their quantificational restriction.

This perspective allows us to make sense of negative answers. Exhaustive interpretation of the answer *Not all linguists* allows us to conclude that some linguists attended because the linguists are in the answer's restriction and can be targeted by exhaustivity, but nothing is concluded about the computer scientists because they fall outside the restriction. An account of extent in these terms could provide a way to maintain positive direction and restricted extent simultaneously. I will develop such an account in this thesis.

The structural theory does not have access to this account. The relation between a quantifier and its restriction is not a relation that natural language syntax is sensitive to. It is a semantic relation.

1.2.3.4 Nondetachability

I do not want to claim that the structural theory is empirically unsalvageable. It is always possible that a creative technical solution could resolve the problems I have discussed. Regardless of whether the structural theory could eventually reach descriptive adequacy, however, no syntactic characterization of alternatives will ever achieve explanatory adequacy in light of the fact that exhaustive inferences are nondetachable.

According to the structural theory, synonymous expressions with different syntactic structures can be associated with different alternatives. Equivalent sentences that vary in structure are then expected to diverge in their exhaustive implications: for instance, B's responses in (68).

(68) **A**: What's the dress code for the panel?
B: {Sneakers are allowed./You may wear sneakers.}
$\rightsquigarrow$ *sneakers aren't required*

Katzir's algorithm generates both *Sneakers are required* and *Sneakers are forbidden* as alternatives to *Sneakers are allowed*. Since these alternatives are symmetric, no exhaustive implication is predicted. The structural theory likewise generates *You must wear sneakers* as alternatives to *You may wear sneakers*. However, there is no simplex modal auxiliary in English that means *forbidden*, so symmetry is avoided, successfully predicting the attested implication that sneakers aren't required.

The empirical contrast predicted between B's answers in (68) is not attested. Syntax does not affect exhaustivity in the way the structural theory leads us to expect. Both answers allow the conclusion that sneakers aren't required. Even if some syntactic constraint were devised that could avoid these problems, it would miss the generalization that B's answers in (68) receive the same exhaustive interpretation because they share the same meaning.

As long as alternatives are regarded as sentences, any instance of nondetachability must be regarded as incidental, arising only because object-language expressibility conspires to allow it in a particular case. This problem would persist even if the structural theory achieved perfect descriptive coverage.

There are predictable and systematic correlations between the meaning of a sentence and its exhaustive implications. The structural theory's failure to provide a principled explanation for this fact, even more than its particular empirical shortcomings, is why it must be rejected.

1.2.4 Summary

Theories of formal alternatives take the task of constraining alternatives to be syntactic in nature. These theories postulate that the facts about symmetry, polarity, and so on should correlate with formal properties of sentences that denote propositions exhaustivity can exclude, and attempt to develop a syntactic recipe that generates exactly the right set of excludable sentences. I argued that leading theories of formal alternatives do not make descriptively adequate predictions about the direction and extent of exhaustivity and cannot provide a principled explanation for nondetachability. Lexical scales (Gazdar, 1979) and structural alternatives (Katzir, 2007) both face these problems, which are direct consequences of the assumption that alternatives are sentences.

The following section discusses an alternate approach, CIRCUMSCRIPTION, that attempts to directly predict the exhaustive interpretation of an answer from its literal meaning.

1.3 Circumscription

A minority tradition has attempted to characterize exhaustivity semantically without reference to object-language expressibility (Balogh, 2009; Groenendijk and Stokhof, 1984; Schulz and van Rooij, 2006; Sevi, 2005; Spector, 2007; van Rooij and Schulz, 2004, 2007; von Stechow and Zimmermann, 1984). This approach has roots in NONMONOTONIC LOGIC and views exhaustive interpretation as an instance of "negation as failure" (Reiter, 1978), the defeasible assumption that whatever is not known to be true is false. As Schulz and van Rooij, 2006 observe, the basic exhaustivity schema underlying this approach can be seen as implementing a particular formalization of common sense reasoning that McCarthy, 1980, 1986 calls "predicate circumscription". For this reason, I call the approach CIRCUMSCRIPTION to distinguish it from theories that use formal alternatives.

This section introduces circumscription, focusing in particular on how it can address the direction and extent properties of answers. I will argue that Groenendijk and Stokhof's analysis of exhaustivity and its detailed elaboration by Schulz and van Rooij provide the kernel of a simple and explanatory semantic solution to the symmetry problem. However, this analysis alone is not sufficient to account for direction and extent. Spector, 2007 develops a related proposal that can account for direction and extent, but defines alternatives based on the syntax of a translation language and leaves the underlying semantic generalizations opaque. I will attempt to tease out the semantic core of this analysis and show how it can explain the exhaustivity facts. The task of chapters 2 and 3 will be to attain a greater degree of explanatory adequacy by deriving Spector's descriptively successful hypothesis about alternatives from first principles.

1.3.1 Minimal models

Groenendijk and Stokhof, 1984 give exhaustive interpretation the semantics in (69). It takes as input a TERM and a PROPERTY INTENSION. The output applies the input term to the input property and says that there is no subset of the input property that the input term applies to in the world of evaluation.[5] According to this proposal, exhaustive interpretation amounts

[5]A note on notation. In contrast to Groenendijk and Stokhof, who translate linguistic expressions into Ty2 (Gallin, 1975), I will assume that natural language is directly interpreted. I adopt predicate logic and the lambda calculus as my metalanguage. For simplicity I assume that focused content is extensional but backgrounded content is intensional; accordingly, I will treat quantificational term answers as type $\langle et \rangle t$ and back-

to the assumption that the question-property is a *minimal element* of the answer term.

(69) $\texttt{exh}_{GS}(\langle T, P\rangle) = T(P_w) \wedge \neg\exists Q : T(Q_w) \wedge Q_w \subset P_w$

If B's answer in (1) denotes a structured proposition consisting of the generalized quantifier over Agnes and the property of having attended the last panel, Groenendijk and Stokhof's exhaustivity schema interprets B's answer as in (70): Agnes attended the panel, and no proper subset of the set of attendees contains Agnes.

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

(70) $\texttt{exh}_{GS}(\langle \lambda P.P(a), \lambda w\lambda x.\textbf{attended}_w(x)\rangle)$
$= \textbf{attended}_w(a) \wedge \neg\exists Q : Q_w(a) \wedge Q_w \subset \lambda x.\textbf{attended}_w(x)$

This formula can only be true if no individuals other than Agnes attended the panel. If Beatrice also attended, then there would be a proper subset of the attendees that also contains Agnes, since $\{a\} \subset \{a, b\}$. The only set containing Agnes without any proper subset that also contains Agnes is the singleton containing Agnes. In this way, exhaustive interpretation of B's answer entails that the set of attendees in w is identical to $\{a\}$, which is the right result.

Because Groenendijk and Stokhof's proposal distinguishes between focused and backgrounded content, it accounts for the sensitivity of exhaustive interpretation to information structure. If B's answer in (28) denotes a structured proposition consisting of the generalized quantifier over Beatrice and the property of not having attended the last panel, Groenendijk and Stokhof's exhaustivity schema interprets B's answer as in (71): Beatrice didn't attend, and no proper subset of the set of individuals who didn't attend includes Beatrice.

(28) **A**: Who didn't attend the last panel? $\lambda w \lambda x.\neg\textbf{attended}_w(x)$
B: Beatrice. $\lambda P.P(b)$

(71) $\texttt{exh}_{GS}(\langle \lambda P.P(b), \lambda w\lambda x.\neg\textbf{attended}_w(x)\rangle)$
$= \neg\textbf{attended}_w(b) \wedge \neg\exists Q : Q_w(b) \wedge Q_w \subset \lambda x.\neg\textbf{attended}_w(x)$

This formula can only be true if no individuals other than Beatrice didn't attend; i.e., everyone other than Beatrice attended. The background property is still minimized in the right direction despite containing negation.

There are several differences between circumscription and formal alternatives. According to Groenendijk and Stokhof's proposal, alternatives are not sentences or even propositions, but possible background extensions. Exhaustive interpretation is not implemented as applying negation to alternative propositions, but shrinking the background property as much as possible given the meaning of the answer. It does involve establishing an informativity-based order and using this order to exclude alternatives. However, because this order is over property extensions and not propositions, the direction of exhaustivity is already built in and does not need to be specified by an independent theory of alternatives. This is how Groenendijk and Stokhof address the symmetry problem. The fact that exhaustivity excludes Beatrice from the extension of the question-property follows as a basic consequence of the inclusion-induced order over property extensions.

ground properties of individuals as type $s\langle et\rangle$. World arguments are represented as subscripts: P_w represents the extension of P at w. I freely mix set and function notation (sometimes egregiously), and gloss over the distinction between sets and characteristic functions.

Because Groenendijk and Stokhof do not appeal to object-language expressibility, they provide a principled account of nondetachability. Exhaustive interpretation is calculated solely on the basis of an answer's meaning and the context of utterance, so different answers with the same denotation are predicted to license the same exhaustive inferences.

The reason symmetry can be broken in a "term-logical" setting but not a propositional one comes down to whether informativity correlates with adding or subtracting items from a set. Possible world semantics measures informative strength using the subset relation. Because classical propositions already contain both the worlds that must be kept and the worlds that must be discarded, there is no way to determine which subset to exclude on the basis of an answer's propositional meaning. If exhaustivity operates at the property level, however, it is possible to establish an order by excluding objects from the background property extension whose presence is not already entailed by the meaning of the answer.

To be sure, the basic components of this recipe for exhaustivity could be preserved within a propositional setting. We could say that the alternatives to B's answer consist of propositions derived by applying the question-property to other individuals in the domain of quantification (Hamblin, 1973), and exclude as many of these as possible given the truth of the answer; Spector, 2007 and Balogh, 2009 pursue this strategy. Symmetry is then broken by the fact that alternatives are generated using only the question-property and not its complement. This solution to the symmetry problem is derivative of the fact that exhaustive interpretation of B's answer in (1) involves excluding objects from the extension of the question-property rather than including them, and preserves Groenendijk and Stokhof's insight that exhaustivity directly regulates the extension of a property.

Without access to the question-property, a propositional approach cannot define polarity semantically. Westera, 2017a,b proposes an approach like this. He assumes that alternatives are answers to an explicit or implicit question under discussion (Ginzburg, 1996; Roberts, 2012), implemented as a set of propositions, and that exhaustivity consists of excluding as many of these as possible. To rule out hypothetical QUDs like {Agnes attended, Beatrice didn't attend, Charmian attended...}, which would allow B's answer in (1) to implicate that Beatrice attended and Charmian didn't, Westera claims that such QUDs are dispreferred because the propositions they contain "vary along two dimensions" (Westera, 2017a, pg. 5) instead of one, but does not say what it means for a set of propositions to vary along a dimension.

Hirsch and Schwarz, 2025 propose a similar theory while appealing to object-language expressibility: "we will refer to propositions or properties as 'positive' or 'negative' based on whether their syntactic expression would contain negation" (pg. 712). In this way, despite rejecting the structural theory, Hirsch and Schwarz inherit its problems.

The solution to the symmetry problem implicit in Groenendijk and Stokhof's analysis relies on the fact that the domain of individuals is not complemented. If the question-property either had to contain some object or its complement, it would no longer be possible to establish an inclusion-induced order over property extensions: excluding that object would necessitate including the complement and vice versa, with the result that the desired property extension would no longer be a subset of the property extension we want to rule out. The only way to establish an inclusion-induced order in this scenario would be to restrict the domain of relevant objects to a set that was not complemented, and minimize the question-property's intersection with this set.

Questions whose abstracts have higher types lead to this scenario. There are many type $\langle et \rangle t$ quantifiers that must either contain some property or its complement, for instance. Suppose there are no "borderline" humans, then the generalized quantifier over Beatrice will either contain $\lambda x.\mathbf{human}(x)$ or $\lambda x.\neg\mathbf{human}(x)$. For this reason, symmetry is predicted to interact with semantic type. Whenever the question-property has a type higher than et, it should require contextual domain restriction to make available the necessary inclusion-induced order, and unless further conditions on this pragmatic process are assumed, symmetry should be breakable in both directions: it should be possible to include either some object or its complement in the set of relevant objects to be minimized. By contrast, this should be impossible with first-order properties. Complementation is a relation on sets, not individuals; there is no object $\bar{b}$ that must appear in any set that does not include b. Minimizing the extension of a first-order property can thus lead to the inference that Beatrice lacks the property, but cannot lead to the inference that Beatrice has the property.

Fox and Katzir, 2011 assume that this asymmetry between first-order and higher-order question abstracts is an undesirable result and reject circumscription for this reason (pg. 96, footnote 15). I will argue in chapter 2 that this feature of circumscription is not a bug, and that "contextual symmetry breaking" is possible precisely when the question-property has a higher type.

What matters for symmetry breaking is the type of the question-property, not the type of the focused term. The inference from *Some of the linguists attended* that not all of them did can be modeled by minimizing the set of attendees relative to the answer. There is no way on this account to generate the inference that all of the linguists attended, whereas theories that use focus alternatives need to worry about the symmetry between *All of the linguists* and *Some but not all of the linguists*. The switch from focus alternatives to background alternatives thus provides an explanation for why exhaustive interpretation of existentially quantified answers always negates rather than affirms universally quantified alternatives. It is because these alternatives attribute the question-property to a larger number of individuals.

In a series of papers, van Rooij and Schulz, 2004 et seq have developed a theory of exhaustivity that expands on Groenendijk and Stokhof's proposal in various directions. Their work draws out the conceptual connections between exhaustivity and "negation as failure" (Halpern and Moses, 1984; Reiter, 1978, 1980; van der Hoek, Jaspars, and Thijsse, 1999, 2000). In nonmonotonic logic, inference rules based on negation as failure allow the failure to derive some conclusion to defeasibly support the negation of that conclusion. Exhaustive interpretation can be seen as an instance of this reasoning pattern.

Groenendijk and Stokhof's exhaustivity schema is especially similar to the operation of PREDICATE CIRCUMSCRIPTION proposed by McCarthy, 1980, 1986. McCarthy's goal was to formalize common sense reasoning by adding a normality assumption: nothing unusual is the case unless explicitly stated otherwise. The normality assumption is implemented by minimizing the set of "unusual things" as much as possible, thus allowing the conclusion that "the objects that can be shown to have a certain property P by reasoning from certain facts A are all the objects that satisfy P" (McCarthy, 1980, pg. 27), which sounds a lot like exhaustive interpretation (see also van Benthem, 1989).

Inspired by this connection, Schulz and van Rooij, 2006 adopt a common implementation of circumscription in nonmonotonic logic, interpretation in MINIMAL MODELS. Their exhaustivity schema is based on the premise that the actual world is one of the "best" or most "minimal" worlds given the speaker's choice of answer, using an order over worlds that depends on the question-property.

(72) a. $\mathtt{exh}_{SvR}(\langle T, P\rangle) = T(P_w) \wedge \neg\exists v : T(P_v) \wedge v < w$
b. $v < w$ iff $P_v \subset P_w$

If the order over worlds is defined as in (72-b), then Schulz and van Rooij's proposal makes similar predictions to Groenendijk and Stokhof's, but there are several differences. Because Schulz and van Rooij's `exh` quantifies over worlds instead of property extensions it is sensitive to meaning postulates (restrictions on the class of possible worlds) and is hardwired to respect the entailments of the answer. Groenendijk and Stokhof's `exh` quantifies directly over property extensions without checking whether the most minimal one is actually realized in some world, and runs the risk of contradiction when applied to some answers.

For instance, on some theories of plurality, term conjunctions like *Agnes and Beatrice* are analyzed as denoting a sum individual, $a \oplus b$ (Link, 1983). This object cannot appear in the extension of a distributive predicate like *attend the panel* without its two parts a and b also appearing in the extension. If applied to the answer *Agnes and Beatrice*, Groenendijk and Stokhof's `exh` would require the extension of the background property to be $\{a \oplus b\}$, which is impossible, but Schulz and van Rooij's `exh` only picks out the most minimal background extensions that is possible in some world, namely $\{a \oplus b, a, b\}$.

Schulz and van Rooij's proposal defines the order over worlds separately from the basic exhaustivity schema. It can be refined by adopting a more sophisticated order, while retaining the core idea that exhaustivity uses this order to select the best worlds. In a dynamic semantic setting, exhaustivity can be refined using an order over possibilities (world-assignment pairs) that interacts with dynamic information; Schulz and van Rooij develop an analysis of suspender clauses along these lines that links suspension to the introduction of discourse referents. They also show that granularity, domain restriction, and other contextual influences on the output of exhaustivity can be accounted for by making the order sensitive to relevance and the goals of the conversational participants: v is more minimal than w iff the proposition that P's extension in v is its extension in the actual world is more relevant than the corresponding proposition for w, with relevance defined using formal techniques from statistical decision theory. The basic inclusion-induced order in (72-b) arises when the question-property's entire extension is at stake.

The power and generality of Schulz and van Rooij's analysis makes it straightforward to refine and I will propose a refinement in this thesis. One problem with the implementation in (72) as it stands is that it does not account for negative or mixed answers. Because all decreasing quantifiers contain the empty set, the most minimal property extension consistent with any decreasing answer will always be empty. This is a bad result insofar as B's answer in an exchange like *A: Who attended the panel? B: Not Beatrice* does not implicate that nobody attended. Similarly, a mixed answer like *Most linguists but not many computer scientists attended* is incorrectly predicted to implicate that no computer scientists attended. The problem is that because $\texttt{exh}_{SvR}$ does not take polarity into account, it predicts that all answers have negative direction.

van Rooij and Schulz, 2004 decompose the exhaustivity schema in (72-b) into two more basic components that regulate the speaker's knowledge. They propose that this epistemic analysis can provide an account of negative answers. The first function, $\texttt{eps}_1$, enforces the implication that the speaker does not believe that the question-property's extension is any larger than what is entailed by the answer. The second function, $\texttt{eps}_2$, enforces the implication that the speaker knows as much as possible about the question-property's extension, as is consistent with the answer that was given.

This proposal is intended as a formalization of the prevalent idea in the literature that exhaustive inferences arise from Grice's maxim of Quantity plus an additional principle of competence maximization (Geurts, 2010; Horn, 1989; Matsumoto, 1995; Sauerland, 2004; Soames, 1982). Applied to B's answer in an exchange like *Q: Who attended the panel? B:*

Agnes, $\texttt{eps}_1$ corresponds to Quantity and licenses the conclusion that B does not know that any other individuals attended. $\texttt{Eps}_2$ corresponds to the assumption that B is informed as to who attended. In contexts where this assumption is supported, A can conclude that B knows that Beatrice did *not* attend. In contexts where this assumption is not supported, only weaker inferences are predicted to arise. By assumption, $\texttt{eps}_1$ is always in force, while $\texttt{eps}_2$ is a feature of only some contexts. Decomposing circumscription into $\texttt{eps}_1$ and $\texttt{eps}_2$ makes no empirical difference for extensional answers, but makes better predictions about intensional answers like *Agnes and possibly Beatrice*.

van Rooij and Schulz suggest that the default way to answer a question is to use an increasing answer and that any deviation from this communicative strategy requires pragmatic justification. Assuming that a speaker who knew the entire extension of the question-property would have given an increasing answer, negative answers then serve as a signal to suspend the speaker competence principle: $\texttt{eps}_1$ applies but not $\texttt{eps}_2$. The prediction for an answer like *Not many computer scientists attended* is that B believes not many computer scientists attended, does not believe of anyone else that they attended, and has no other beliefs regarding panel attendance.

Although this proposal avoids the wrong prediction that negative answers have negative direction, it does not provide an account of positive direction. Exhaustive inferences with positive direction do not have a weaker epistemic status than those with negative direction, so we still need an explanation for why we can conclude from *Not many computer scientists attended* that some did. To get positive direction right, the order on worlds should not only be suspended, but reversed.

If negative answers served as a signal to suspend the speaker competence principle, we would still expect $\texttt{eps}_2$ to kick in when speaker ignorance is implausible. Mixed answers provide such a case. In (73), B does know that some individuals appear in the extension of the question-property, yet B's answer still does not implicate that no computer scientists attended, contrary to the predictions of $\texttt{eps}_2$.

(73) **A**: Who attended the panel?
B: Not many computer scientists. But most linguists did.
$\rightsquigarrow$ *not all linguists, some computer scientists*

Naturally occurring examples of this form are especially common in self-addressed questions. In (74) and (75) (from the Corpus of Contemporary American English; Davies, 2008), the speaker asks a question, provides a negative answer and then immediately provides a contrasting positive answer. This discourse move appears to be motivated by the speaker's desire to reject a salient contextual expectation. There is no sense that an ignorance implicature is being canceled.

(74) The number of cyclists on the road in my city has skyrocketed in the last couple of years, and what changed to make that happen? Not cycletracks or bike lanes, though the city is placing them now, after the fact. No, what changed was 1. Gas jumped to $3/gal 2. Significant mass transit cuts ensued so the busses are increasingly impractical 3. a popular fashion trend of urban hipsterism on fixed-gear bikes

(75) Its slogan—"New Mexico's Only Colorado Ski Area"—was part of the problem. Just where was it? Not on the west-facing slopes of Raton Pass, where remnants of a summer-only chairlift still can be seen. Instead, it was 12 miles northeast of Raton, N.M., but literally feet within the Colorado boundary. It ran from 1965-89.

These examples suggest that negation does not generally signal ignorance. Even if it did, though, van Rooij and Schulz's conjecture about negative answers would be insufficient to explain direction and extent.

1.3.2 Positive and negative formulae

Circumscription as developed by Groenendijk and Stokhof, 1984 and Schulz and van Rooij, 2006 makes correct predictions about positive answers, but not negative or mixed answers. von Stechow and Zimmermann, 1984 and Spector, 2007 have paid close attention to the effects of polarity on exhaustive interpretation, however, and the account I develop in this thesis shares important features with both proposals.

von Stechow and Zimmermann, 1984 recognize that exhaustivity has negative direction with positive answers and positive direction with negative answers, but do not notice the attested contrast in extent, claiming that answers like those in (27) implicate that everyone attended except Beatrice. The speakers I consulted, as well as those consulted by Schulz and van Rooij, 2006 and Spector, 2007 report to the contrary that negative answers typically have restricted extent, however.

(27) **A**: Who attended the last panel?
B: {Beatrice didn't attend./Not Beatrice.}

Like Groenendijk and Stokhof, von Stechow and Zimmermann adopt a structured approach to exhaustivity. Because propositions do not have a polarity, von Stechow and Zimmermann argue that explicit access to focused and backgrounded content is necessary in order to distinguish answers of different polarity. On their account, exhaustivity minimizes the question-property's extension when the answer is positive, maximizes the question-property's extension when the answer is negative, and delivers varied results with mixed answers. This analysis provides a successful existence proof that it is in principle possible to link the direction of exhaustivity to answer polarity, and they propose a formal characterization of polarity using the logic of generalized quantifiers. This is also the strategy I will take in chapter 2.

While pointing out that von Stechow and Zimmermann overlook extent, Spector, 2007 is inspired by their approach and develops an account of exhaustivity that hits all the empirical targets I have set. It correctly predicts the direction and extent properties of all varieties of answers. Spector translates answers into propositional logic and defines alternatives within the translation language in a way that distinguishes between positive, negative and mixed answers. The main limitation of Spector's descriptively adequate proposal is that these rules about alternatives are stipulations: they do not follow from any independent formal property of the system.

Spector states his proposal using a question-dependent logic. If Q is a question about the property P, then L_Q is the propositional language whose letters are based on the application of P to every individual in the domain of quantification. Interpretations are defined in the usual way (functions from formulae to truth-values). The connectives have classical semantics but only conjunction and disjunction can connect arbitrarily complex formulae; negation only applies to atomic proposition letters. The LITERALS of L_Q are the formulae that consist of either single proposition letters or their negations. Because L_Q does not allow for double negation, the following notational convention is useful: if L is a literal, then $-L$ is defined as $\neg p$ iff $L = p$ and p iff $L = \neg p$.

To model exhaustive interpretation, all formulae are associated with a set of alternative formulae. Alternatives are defined using the notion of FAVORING.

(76) **Favoring**

a. If w is an interpretation and L is a literal, w_{-L} is the unique interpretation that is exactly like w except that $w_{-L}(L) = w(\neg L)$.
b. If F is a formula and L is a literal, F *favors* L iff there is an interpretation w such that $w(F) = w(L) = 1 \neq 0 = w_{-L}(F)$.
c. If F is a formula and L is a literal, F *strongly favors* L iff F favors L and F does not favor $-L$.

A formula F favors a literal L iff changing L from true to false can change F from true to false. The formula $a \wedge b$ favors a and b: it is possible to find an interpretation w for which $w(a) = w(a \wedge b) = 1$ and $w_{-a}(a \wedge b) = 0$, just pick any w that assigns $a \wedge b$ to 1; similar reasoning holds of b. The formula $a \vee b$ favors a and b: it is possible to find an interpretation w for which $w(a) = w(a \vee b) = 1$ and $w_{-a}(a \vee b) = 0$, just pick any w that assigns a to 1 and b to 0; similar reasoning holds of b. Because neither formula favors the corresponding negative literals $\neg a$ and $\neg b$, they additionally STRONGLY FAVOR a and b. By contrast, $(a \vee b) \wedge (\neg a \vee \neg b)$ favors $a, b, \neg a$, and $\neg b$, but does not strongly favor any positive or negative literal.

Spector uses these definitions to ground the following theory of alternatives.

(77) **Alternatives**

a. A *positive* formula does not favor any negative literal.
b. A *quasi-positive* formula does not strongly favor any negative literal.
c. If F is quasi-positive, ALT(F) is the union of $\{F\}$ with the set of all positive formulae.
d. If F is not quasi-positive, ALT(F) is the closure under conjunction and disjunction of all the literals that F favors.

Spector proposes a theory of exhaustivity that is very similar to the two-step epistemic schema by van Rooij and Schulz, 2004, but defined relative to alternative formulae in L_Q. Quantity is modeled as the supposition that the speaker believes the answer and does not believe any of its stronger L_Q alternatives. Competence is modeled as the supposition that the speaker believes as many L_Q alternatives are false as possible given Quantity. Spector proves that this system is equivalent to van Rooij and Schulz-style circumscription when applied to positive answers. However, it also makes interesting and correct predictions with mixed and negative answers.

To get a sense of the system, here are some sample answers, their translations, and their alternatives, in a context where Q is *Who purred?* and the domain of quantification contains three individuals, Frederick, Gertrude, and Harriet. All three are cats; Frederick and Gertrude are kittens. L_Q is based on the atoms f, g, and h, where $w(f) = 1$ iff Frederick purred in w, and so on. I use $\longmapsto$ to indicate translation into L_Q.

(78) Frederick $\longmapsto f$
ALT = all positive formulae: $\{f \vee g \vee h, f \vee g, f \vee h, g \vee h, (f \vee g) \wedge (f \vee h), (f \vee g) \wedge (g \vee h), (f \vee h) \wedge (g \vee h), f, g, h, (f \wedge g) \vee (f \wedge h) \vee (g \wedge h), (f \wedge g) \vee (f \wedge h), (f \wedge g) \vee (g \wedge h), (f \wedge h) \vee (g \wedge h), f \wedge g, f \wedge h, g \wedge h, f \wedge g \wedge h\}$

Applied to a positive answer like *Frederick* in (78), exhaustivity targets all stronger and independent L_Q alternatives. Because all positive formulae are alternatives to f, this excludes g and h, correctly predicting the implication that Frederick purred and no other cats purred. Spector's hypothesis about alternatives thus successfully accounts for the negative direction and unrestricted extent of positive answers: other individuals are excluded from the extension of the question-property, and exhaustivity targets individuals that the answer is not about (*Frederick* is not about Gertrude or Harriet).

Applied to a negative answer like *not Gertrude* in (79), exhaustivity delivers nothing beyond the literal meaning of the answer. This is because only formulae constructed from the literals favored by a negative answer count as alternatives to that answer. Because the only literal $\neg g$ favors is itself, there are no alternatives to exclude and nothing can be inferred about other individuals. This is the correct result for *not Gertrude*.

(79) not Gertrude $\longmapsto \neg g$
$\textsc{alt} = \{\neg g\}$

Negative answers like *not all of the kittens* in (80) allow some exhaustive implications. The L_Q translation of this answer is the grand disjunction of all the negative literals based on kittens, in this case $\neg f$ and $\neg g$.

(80) not all of the kittens $\longmapsto \neg f \vee \neg g$
$\textsc{alt} = \{\neg f \vee \neg g, \neg f, \neg g, \neg f \wedge \neg g\}$

Although this is a negative answer, Spector's recipe for constructing alternatives closes this set of literals under conjunction and disjunction and thus allows additional distinct alternatives, namely the two disjuncts and the corresponding conjunction $\neg f \wedge \neg g$. Although neither disjunct can be excluded consistently with Quantity (otherwise there would be a stronger alternative that the speaker should have asserted, the other disjunct), exhaustivity can target the conjunction: it's not the case that Frederick didn't purr and Gertrude didn't purr. Thus we conclude that one of the kittens purred, without drawing inferences as to which one or concluding anything about Harriet.

In this way, Spector's hypothesis about alternatives makes correct predictions about the positive direction and restricted extent of negative answers. Because negative answers are based on negative literals, their alternatives are also based on negative literals. Excluding these alternatives has the effect of attributing the question-property to some individuals, through double negation. This is how Spector accounts for positive direction. Because only literals favored by the answer are used to construct the alternatives to negative answers, exhaustivity does not enforce implications regarding individuals that the answer is not about (*not all of the kittens* is about Frederick and Gertrude but not Harriet). This is how Spector accounts for restricted extent.

Spector's notion of "quasi-positivity" distinguishes between different classes of mixed answers. Answers that favor negative literals are predicted to still have unrestricted extent as long as the corresponding positive literals are also favored, whereas answers that "strongly favor" negative literals are predicted to have restricted extent. This is because quasi-positive answers (which do not "strongly favor" negative literals) have the entire class of positive formulae as alternatives, but answers that are not quasi-positive are only associated with alternatives based on the literals they favor.

The L_Q translation of *exactly one kitten* in (81) is the disjunction of all formulae that attribute the question-property to one kitten and exclude all other kittens, in this case $f \wedge \neg g$ and $\neg f \wedge g$. This formula counts as quasi-positive because it does not strongly favor any negative literals: every negative literal favored by the answer is "canceled out" by the corresponding positive literal. For this reason, all positive formulae count as alternatives.

(81) exactly one kitten $\longmapsto (f \wedge \neg g) \vee (\neg f \wedge g)$
ALT = the union of all positive formulae with $(f \wedge \neg g) \vee (\neg f \wedge g)$

Given Quantity, exhaustivity cannot exclude any alternatives that entail any of the positive literals favored by the answer (otherwise there would have been a stronger L_Q alternative that the speaker should have asserted), hence nothing is concluded about Frederick or Gertrude beyond the literal meaning of the answer. However, $(f \wedge \neg g) \vee (\neg f \wedge g) \wedge h$ is a stronger L_Q alternative that can be contingently excluded, predicting the implication that Harriet did not purr. More generally, Spector's hypothesis about alternatives correctly predicts that *exactly one kitten purred* can implicate that no nonkitten purred.

(82) if Gertrude purred, Frederick purred $\longmapsto f \vee \neg g$
ALT $= \{f \vee \neg g, f, \neg g, f \wedge \neg g\}$

The L_Q translation of *if Gertrude purred, Frederick purred* in (82) is the disjunction of f and $\neg g$ (L_Q does not include material implication as a connective). This formula favors the literals f and $\neg g$ and is not quasi-positive, because it does not also favor g. Spector's recipe for constructing alternatives closes these literals under conjunction and disjunction and thus allows the additional alternatives f, $\neg g$, and $f \wedge \neg g$. Exhaustivity cannot exclude either literal (or else the other one would have been a stronger alternative) but it can exclude the conjunction, predicting the implication that it's not the case that Frederick purred and Gertrude didn't. Together with the meaning of the answer, we can conclude that if Frederick purred, then Gertrude purred. This is the conditional perfection inference that we wanted exhaustivity to deliver. Meanwhile, because none of the alternatives favor h or $\neg h$, nothing is concluded about Harriet. The mixed direction and restricted extent properties of conditional answers are thus correctly predicted.

To summarize: whether a formula is associated with a larger or smaller set of alternatives depends on whether it is "quasi-positive". A quasi-positive formula either only favors positive literals, or favors some negative literals whose corresponding positive literals are also favored. Formulae that favor negative literals without also favoring the corresponding positive literals are classed as negative. This distinction allows a descriptively successful account of direction and extent. Positive and quasi-positive answers have negative direction and unrestricted extent because their alternatives include all positive formulae. Answers that are not quasi-positive have positive or mixed direction and restricted extent because their alternatives only include formulae based on the literals they favor.

Although these are the right results, it is not obvious why they should hold. Spector comments that "It remains to be seen whether the role played by polarity (namely, the distinction between positive and non-positive answers) can be derived in a more principled way" (pg. 241). That is the goal of this thesis. Spector has not pursued this question in his subsequent work, and neither have other scholars, perhaps because the contemporaneous publication of Katzir, 2007 shifted the literature's focus toward the structural theory.

1.3.3 Computational efficiency

I have so far argued for circumscription over formal alternatives on descriptive and explanatory grounds. There is a third sort of argument I want to make on grounds of computational efficiency. The proposals by Schulz and van Rooij, 2006 and Spector, 2007 can be implemented with relatively low time complexity, unlike the structural theory.

Adapting analytical techniques from computational complexity theory, Mascarenhas, 2021 proves that the structural theory is computationally intractable. Efficiency can be measured by comparing the size of an input source to the number of output alternatives an algorithm generates as a function of the input. In a propositional logic like Spector's L_Q, the number of atomic proposition letters in the language can be treated as the input n.

Mascarenhas shows that the Katzir, 2007 algorithm generates at least $(2n-1)^n * 2^{n-1}$ alternatives for a sentence with n atoms. In "big O notation" (Cormen et al., 1990), the order of Katzir's algorithm is $O(n^n)$, a growth rate even faster than ordinary exponential runtime because the base grows with n. This is unsatisfactory if exhaustivity is taken to be the result of online reasoning about alternatives, especially if alternatives are identified with mental models in cognitive psychology. As Mascarenhas discusses, research on the psychology of reasoning has found that human subjects can reason about at most five to seven mental models at once (Johnson-Laird, 2008), far fewer than the large numbers of formal alternatives generated by the structural theory. Regardless of whether we should think of alternatives in this way, exponential runtime is typically considered intractable in cognitive science (Frixione, 2001).

The Spector, 2007 algorithm is not tractable either, but has equivalent implementations that are. Positive answers in L_Q have the class of all positive formulae as their alternatives. This set can be generated by closing the positive literals under conjunction and disjunction. As Mascarenhas observes, if L_Q has n atomic proposition letters then the number of positive formulae will be equal to the nth Dedekind number, the number of monotone boolean functions on n variables (Kisielewicz, 1988; Quackenbush, 1986). Counting the Dedekind numbers is currently an open problem in lattice and order theory, but it is at least an exponential function of n.

However, Spector, 2007's treatment of positive answers is equivalent to the circumscription accounts of van Rooij and Schulz, 2004 and Schulz and van Rooij, 2006, who directly quantify over property extensions. For ease of comparison, Spector, 2016 shows that this approach can also be equivalently stated relative to sets of propositions that are not closed under Boolean operations (see also Schulz and van Rooij, 2006, pg. 219, footnote 19), given in (83).

(83) a. $\texttt{exh}(p, \textsc{alt}) = p_w \wedge \neg\exists v : p_v \wedge v < w$
b. $v < w$ iff $\{q_{\in \textsc{alt}} | q(v)\} \subset \{q_{\in \textsc{alt}} | q(w)\}$

The exhaustivity schema in (83) takes a proposition p and a set of propositions ALT as input. If ALT is the set of propositions that result from applying the question-property to every individual in the domain of quantification, the output is equivalent to the implementation we have already seen. It is not necessary to close ALT under conjunction and disjunction.

According to (83), a world v is more minimal than a world w iff v makes fewer alternatives true than w. Exhaustive interpretation of an answer will exclude worlds that the answer shares with other alternatives if possible, while keeping worlds that only appear in the denotation of the answer. For instance, exhaustive interpretation of the answer in an exchange like *Who purred? Frederick* excludes worlds in which Gertrude or Harriet also purred, since these worlds verify additional alternatives, whereas worlds in which only Frederick purred verify only the answer itself.

Circumscription therefore requires only as many alternatives as there are objects in the domain of quantification, which is also the number of atomic proposition letters in L_Q. Since this number was our input n, the order of the exhaustivity schema in (83) is $O(n)$. Circumscription can be implemented in linear time.

No comparable reduction in the cardinality of alternative sets could be made for the structural theory. The reason Spector's proposal can be made tractable is because the extra conjunctions and disjunctions that Spector, 2007 includes among alternative sets can be eliminated without affecting empirical coverage, as long as exhaustivity is characterized using circumscription. For the structural theory, however, alternatives are sentences and this sort of reduction would result in the loss of syntactic information, which would necessarily affect the structural theory's empirical predictions.

This argument from time complexity is presumably not decisive. The desiderata for theoretical linguistic proposals do not usually include computational efficiency, and intractable theories at the computational level of Marr, 1982 could in principle be coupled with tractable theories at the algorithmic level. Insofar as more efficient theories should be preferred on psychological grounds, however, these tractability results suggest that circumscription surpasses the structural theory in cognitive plausibility.

1.3.4 Summary

Circumscription accounts for exhaustive interpretation by quantifying over possible background property extensions. These "background alternatives" can either be excluded directly (Groenendijk and Stokhof, 1984; Schulz and van Rooij, 2006; van Rooij and Schulz, 2004; von Stechow and Zimmermann, 1984) or by applying the question-property to individuals in the domain of quantification and excluding the resulting propositions (Spector, 2007). I argued that circumscription can make correct predictions about direction and extent in virtue of having access to an answer's meaning and defining alternatives on this basis. Circumscription also leads to the expectation that exhaustive inferences should generally be nondetachable, a prediction I argued is correct. Finally, circumscription is computationally tractable. The structural theory cannot generally predict direction and extent, provides no explanation for nondetachability, and is computationally intractable. For these reasons, circumscription should be favored over formal alternatives.

1.4 Conclusion

The central message of this chapter is that alternatives should be characterized using the tools of model-theoretic semantics. The direction and extent properties of answers call for a semantic analysis of exhaustivity using circumscription. In this way, an explanatory solution to the symmetry problem is in sight.

Although existing circumscription accounts have achieved either descriptive or explanatory adequacy, currently no account achieves both. Spector, 2007 makes correct predictions about direction and extent, but the key features of his analysis exploit the syntax of a translation language in a rather cryptic way, leaving open the question as to why polarity affects alternatives in the first place.

Schulz and van Rooij, 2006, building on Groenendijk and Stokhof, 1984, provide a compelling answer to the explanatory question: modeling exhaustive interpretation as minimizing a property follows from the fact that questions are asked in order to gain information about property extensions, and a cooperative speaker should provide as much information about the question-property as possible. However, because this account invariably minimizes the question-property it only makes correct predictions for answers with negative direction, not positive or mixed direction, and it ignores extent.

The rest of this thesis attempts to match the descriptive coverage of Spector, 2007 while maintaining the explanatory power of Schulz and van Rooij, 2006. One way to account for negative answers within Schulz and van Rooij's setting would be to revise the order on worlds that matters for exhaustive interpretation by making it sensitive to answer polarity. This analysis would continue to model exhaustive interpretation by quantifying over background property extensions, but whether the background property is minimized, maximized, or manipulated in some more complex way would depend on whether the answer is positive, negative, or mixed. Chapter 2 develops such an analysis.

Chapter 2

Two orders

This chapter develops an analysis of how answers with restricted extent are interpreted. The order on possible worlds underlying exhaustive interpretation is argued to vary depending on the semantics of answers, which are sorted into two classes depending on their extent properties. Answers with unrestricted extent are interpreted using a simple inclusion-induced order, as in the proposals of Groenendijk and Stokhof, 1984 and Schulz and van Rooij, 2006. Answers with restricted extent are interpreted using a more complex order that simultaneously represents positive and negative information, and can resolve conflicting information in a principled way. Application of this order is shown to make correct predictions about the direction of negative and mixed answers. Finally, the chapter extends the analysis to questions whose abstracts denote sets of objects other than individuals, and shows that it makes correct predictions about how symmetry breaking interacts with semantic type.

2.1 Default and marked answers

2.1.1 Two classes of answers

Chapter 1 concluded with several open problems for the analysis of exhaustive interpretation. I argued that theories of exhaustivity based on formal alternatives do not make adequate predictions about the direction and extent properties of answers, which should instead be analyzed in logical semantic terms. Theories of exhaustivity based on minimizing worlds (Groenendijk and Stokhof, 1984; Schulz and van Rooij, 2006; Spector, 2007) offer a principled solution to these problems. Such theories analyze the interpretation of answers using the formal technique of CIRCUMSCRIPTION (McCarthy, 1980, 1986), developed in the computational literature on nonmonotonic reasoning. Circumscribing an answer amounts to minimizing the extension of the question-property as much as can be done consistently with the semantics of the answer (modulo granularity, domain restriction, and other sources of context-sensitivity). An analysis of exhaustivity as circumscription makes correct predictions about the direction of positive answers and provides a principled solution to the symmetry problem. However, it does not account for the direction of negative and mixed answers, and says nothing about extent.

I proposed to represent answers using structured propositions, pairs consisting of a property meaning in the background and a generalized quantifier in the foreground. The polarity of an answer depends on whether the foreground consists of an increasing, decreasing, or nonmonotone quantifier (1). Example pairs are given in (2) for positive (2-a), negative (2-b), and mixed answers (2-c)-(2-d).

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

(2) Who attended?
   a. $[\![$Agnes attended$]\!] = \langle \lambda P.P(a), \lambda w \lambda x.\textbf{attended}_w(x)\rangle$
   b. $[\![$Beatrice didn't attend$]\!] = \langle \lambda P.\neg P(b), \lambda w \lambda x.\textbf{attended}_w(x)\rangle$
   c. $[\![$Exactly two linguists attended$]\!] = \langle \lambda P.|\textbf{linguists} \cap P| = 2, \lambda w \lambda x.\textbf{attended}_w(x)\rangle$
   d. $[\![$Agnes but not Beatrice attended$]\!] = \langle \lambda P.P(a) \wedge \neg P(b), \lambda w \lambda x.\textbf{attended}_w(x)\rangle$

I made the following claims about answer polarity: exhaustive interpretation of positive (increasing) answers has negative direction and unrestricted extent, and exhaustive interpretation of negative (decreasing) answers has positive direction and restricted extent. Mixed (nonmonotone) answers vary: some (*Exactly two linguists*) have negative direction and unrestricted extent, and others (*Most linguists but not a lot of computer scientists*) have mixed direction and restricted extent. These claims are summarized in Table 2.1.

Answer	**Polarity**	**Direction**	**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	-	n/a	restricted
No linguists	-	n/a	restricted
Not all of the linguists	-	+	restricted
Not both Agnes and Beatrice	-	+	restricted
Not a lot of linguists	-	+	restricted
Neither Agnes nor Beatrice	-	n/a	restricted
Agnes but not Beatrice	±	n/a	restricted
Most linguists but not a lot of computer scientists	±	±	restricted
If Agnes, then Beatrice	±	±	restricted

Table 2.1: Typology of answers.

Direction and extent interact in the following way. Answers with restricted extent whose semantics completely determines whether the question-property holds of all individuals that they are "about" do not give rise to exhaustive inferences at all: *not Beatrice* is only about Beatrice, entails that Beatrice didn't attend, and does not license exhaustive inferences about other individuals. For this reason, it does not make sense to talk about the direction of these answers, hence the "n/a" label in Table 2.1. By contrast, answers with restricted extent that do not completely fix the status of all individuals they are "about" relative to the question-property do allow exhaustive inferences about precisely these individuals: *not all of the linguists* does not license conclusions about computer scientists, but can implicate that some linguists attended.

There is a pattern: the answers with unrestricted extent all have negative direction. By contrast, the answers with restricted extent all have positive, mixed, or no direction. It will be useful to group together the answers in the first two rows, and the answers in the second two rows. I will call answers with unrestricted extent DEFAULT answers, and answers with restricted extent MARKED answers.

To my knowledge, the only existing proposal that properly distinguishes between default and marked answers is that of Spector, 2007. Default answers correspond to Spector's "quasi-positive" answers. These are answers whose L_Q translations either only contain positive literals, or contain the corresponding positive literal for every negative literal they contain. Marked answers are answers whose L_Q translations contain at least one negative literal without the corresponding positive literal.

Quasi-positivity is a property of formulae, not denotations. Because my goal is to link the exhaustive interpretation of answers to their semantics, I will need to identify a model-theoretic property that tracks Spector's notion of quasi-positivity. This is the task of chapter 3. Before doing so, I will develop a recipe for exhaustivity that can adequately describe how marked answers are interpreted.

The goal of this chapter is to generalize the circumscription approach to cover marked answers. I retain Schulz and van Rooij's treatment of exhaustivity as selecting the "best" possible worlds relative to an answer. However, exhaustive interpretation of marked answers requires an order other than the inclusion relation over background property extensions. Although it is technically possible to define a single order that simultaneously covers all answers, my strategy will be to define two separate orders, one for default answers, and another for marked answers. Because we have seen that all answers with unrestricted extent have negative direction, the default order can simply be identified with the basic inclusion order from Schulz and van Rooij, 2006.

Developing the hybrid order, which regulates the interpretation of answers with restricted extent, is the focus of this chapter. In contrast to the simple default order, the need to track both positive and negative information will require the hybrid order to have a more complex algebraic structure. I will suggest that this is why the hybrid order is marked, and that the complexity of the order underlying exhaustive interpretation reflects the semantic markedness of an answer relative to a question.

2.1.2 The default order

Schulz and van Rooij, 2006 proposed the implementation of exhaustivity in (3), which takes a term and a property as input and selects the "best" or most minimal worlds in which the term applies to the property, i.e., worlds for which there are no "better" worlds in which the term applies to the property.

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

The (3-b) order ranks worlds by the size of background property extensions in those worlds: a world v is more minimal than w iff P's extension in v is a subset of P's extension in w. This order is thus isomorphic to the inclusion relation over possible extensions of P.

I propose that (3-b) is the DEFAULT ORDER, underlying exhaustive interpretation of default answers. More formally, this order has the algebraic structure of a LATTICE (Birkhoff, 1940; Davey and Priestley, 1990). A lattice is a special sort of ordered set for which every pair has a join ($\vee$) and a meet ($\wedge$). If the order is identified with inclusion, as in the default order, the join operation corresponds with union and meet with intersection.

(4) An ordered set L is called a **lattice** iff $\forall a, b \in L : a \vee b \in L$ and $a \wedge b \in L$.

If a lattice L has the complementation structure of a powerset, then it is a BOOLEAN LATTICE, and if the order on L is inclusion, then we specifically have a POWERSET ALGEBRA on L.

(5) A lattice L is called a **Boolean lattice** iff
  a. L is distributive: $\forall a, b, c \in L : a \vee (b \wedge c) = (a \wedge b) \vee (a \wedge c)$;
  b. L has minimal and maximal elements (0 and 1);
  c. L is complemented: $\forall a \in L : \exists! b \in L : a \wedge b = 0, a \wedge b = 1$.

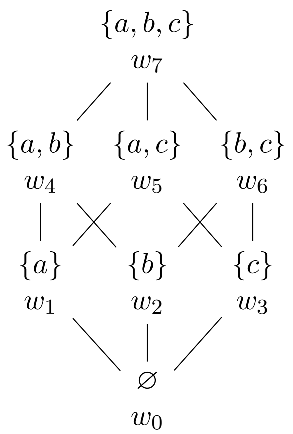

Figure 2.1: The default order.

The inclusion-induced order over possible extensions of a property has the structure of a Boolean lattice. If our domain of entities includes three individuals, Agnes, Beatrice, and Charmian, and there are no postulates restricting the possible extensions of some property P, then the default order over P-extensions is diagrammed in Figure 2.1, with $\subseteq$ proceeding from bottom to top.

As we have seen, the default order successfully predicts the exhaustive interpretation of positive answers. The literal meaning of B's answer in (6) selects worlds in which Agnes attended ($w1$, $w4$, $w5$, $w7$); `exh` then selects worlds from these in which the set of attendees is as small as possible (w_1). The conclusion that the actual world is w_1 (and $P = \{a\}$) correctly models the implication that only Agnes attended. In this way, minimizing worlds amounts to selecting the lowest worlds in the Hasse diagram.

(6) **A**: Who attended the last panel?
B: Agnes.
$\rightsquigarrow$ *nobody else*

The literal meaning of B's answer in (7) selects worlds in which either Agnes or Beatrice attended (w_1, w_2, w_4, w_5, w_7); `exh` then selects the most minimal of these (w_1, w_2). The conclusion that the actual world is either w_1 or w_2 (and either $P = \{a\}$ or $P = \{b\}$) correctly models the implication that they didn't both attend, and nobody else attended. Disjunctive answers show that the default order can have multiple minima.

(7) **A**: Who attended the last panel?
B: Agnes or Beatrice.
$\rightsquigarrow$ *not both, nobody else*

The default order also successfully predicts the interpretation of nonmonotonic answers like (8). Suppose that Agnes and Beatrice are linguists, but not Charmian: the literal meaning of B's answer in (8) selects worlds in which Agnes and Beatrice both attended (w_4, w_7), and `exh` selects the minimal one (w_4). The conclusion that the actual world is w_4 (and $P = \{a, b\}$) correctly models the implication that nobody attended other than exactly two linguists. Because our toy model only contains two linguists, w_4 is the only minimum. If there were more linguists in the model, there would be multiple minima, each with exactly two linguists in P, but Charmian would still be excluded.

(8) **A**: Who attended the last panel?
B: Exactly two linguists.
$\rightsquigarrow$ *nobody else*

All these results are familiar from chapter 1. It is worth paying close attention to the formal properties of the default order, however, to appreciate its relative simplicity compared to the hybrid order I develop in the remainder of this chapter. A Boolean lattice is a particularly elegant and attractive sort of order. It is the order induced by the relation of set inclusion, a fundamental relation in natural language semantics. Aside from the domain of properties, many other domains in language have been argued to instantiate this sort of structure. Link, 1983 finds lattice structure in the domain of individuals, and Bach, 1986 extends this treatment to the domain of events. Winter, 2001 claims that all semantic domains form Boolean lattices. Regardless of whether this is true, Boolean lattices often appear in our models, so it is unsurprising that this structure provides the default order for exhaustivity.

The default order cannot apply to negative answers, however. As we have seen, minimizing the question-property relative to an answer like *Not Beatrice* would predict the implication that *nobody* satisfies the question-property, incorrectly. Exhaustive interpretation of these answers must involve a different order. In the rest of this chapter, I develop one and demonstrate that it makes correct predictions about the exhaustive interpretation of marked answers.

2.2 Restriction operators

Marked answers, which comprise negative and mixed answers with restricted extent, vary in direction. Exhaustive interpretation of negative answers has positive direction when it has any effect at all. Exhaustive interpretation of (marked) mixed answers has mixed direction. These observations call for a theory of exhaustivity that links an answer's directional properties to its polarity.

So far, I have assumed that answer polarity corresponds to monotonicity: positive answers denote increasing quantifiers, negative answers denote decreasing quantifiers, and mixed answers denote nonmonotonic quantifiers. Because mixed/nonmonotonic answers vary in their extent properties, however, monotonicity cannot be the sole predictor. Moreover, to adequately describe the exhaustive interpretation of mixed/nonmonotonic answers with mixed direction, we still need to determine how direction is allocated across different subsets of individuals: to predict the implicatures of an answer like *Most linguists but not a lot of computer scientists attended the panel* (not all linguists attended, some computer scientists did), it should be predictable on the basis of the answer's meaning that "negative" inferences are drawn about the linguists and "positive" inferences are drawn about the computer scientists, rather than the reverse of this, or some other configuration.

This section develops a set of formal tools that can be used to predict the direction of exhaustivity for marked answers. The concept of "aboutness" that matters for extent is formally implemented using the notion of quantificational restriction: marked answers with restricted extent only license exhaustive implications about individuals inside their restriction. To account for positive, negative, and mixed direction, I develop two novel RESTRICTION OPERATORS that sort a quantifier's restriction into positive and negative subsets. The subsets yielded by the restriction operators are exactly the subsets of a quantifier's restriction that matter for exhaustive interpretation. Although the necessary tools are more fine-grained than sorting quantifiers by monotonicity, they nonetheless support the view that the logic of generalized quantifiers (Barwise and Cooper, 1981; Keenan and Stavi, 1986; Westerståhl, 1988) plays an important role in understanding exhaustivity.

2.2.1 Quantificational restriction

To analyze answers with restricted extent, the first task is to identify the set that restricts the extent of exhaustivity. Negative answers like *Not all of the linguists* and *Not many linguists* allow the implication that some linguists attended, but do not license implications about computer scientists. The extent of exhaustivity is thus restricted to the linguists. Similarly, *Not both Agnes and Beatrice* allows the implication that one of them attended, but does not license implications about other individuals. The extent of exhaustivity is thus restricted to the set containing only Agnes and Beatrice.

Mixed answers like *Most linguists but not a lot of computer scientists* can implicate that not every linguist attended, and that some computer scientists attended, but do not license implications about individuals who are not linguists or computer scientists. The extent of exhaustivity is thus restricted to the set containing the linguists and computer scientists. Conditional answers like *Beatrice attended if Agnes did*, which instantiate another species of mixed answer, can implicate that either both Agnes and Beatrice attended, or neither did; they do not license implications about other individuals. The extent of exhaustivity here is restricted to the set containing Agnes and Beatrice.

These answers and their corresponding sets stand in a familiar relation. It is the same relation that holds between *Every linguist* and the set of linguists, the generalized quantifier *Agnes or Beatrice* and the set containing Agnes and Beatrice, and the Montagovian individual *Agnes* and the singleton containing Agnes. It is the relation between a generalized quantifier and its RESTRICTION.

The intuitive notion of quantificational restriction should be clear. A generalized quantifier's restriction is a set that contains the objects whose presence in (typically) some other set must be checked in order to verify whether the quantifier applies to that set (whether the set is contained in a GQ term's denotation). For GQ terms composed by applying a quantificational determiner to a property-denoting nominal (*every linguist, no computer scientist*), the restriction is the denotation of the nominal argument. For Montagovian individuals like *Agnes*, the restriction is just the singleton containing the corresponding object, and for complex quantifiers like *Agnes and Beatrice* or *every linguist or no computer scientist*, the restriction typically corresponds to the union of the restrictions of the conjoined/disjoined quantifiers.

von Stechow and Zimmermann, 1984, exploring how the formal description of exhaustivity interacts with the logic of generalized quantifiers, show that quantificational restrictions can be characterized model-theoretically as follows. The definition uses the auxiliary notion of a LIVE-ON SET (Barwise and Cooper, 1981), as in (9).

(9) **Live-on sets.**
$\mathbf{L} = \lambda T_{\langle et \rangle t} \lambda S_{et}.\forall P : T(P) \leftrightarrow T(P \cap S)$

The **L** operator applies to a quantifier T and delivers (the characteristic function of) a set that contains the sets T "lives on". T lives on a set S iff both T and its complement are closed under intersection with S. If T applies to a set P, then T also applies to P's intersection with S. If T does not apply to P, then T does not apply to P's intersection with S.

Live-on sets play an important role in the formal statement of typological generalizations about quantificational determiners. Barwise and Cooper, 1981 and Keenan and Stavi, 1986 claim that all natural language determiners are CONSERVATIVE: the result of composing any determiner with a property-denoting nominal will be a quantifier that lives on the denotation of the nominal argument. Concretely: *Some cat purrs* is equivalent to *Some cat is a cat that purrs*. Typically, a quantifier will live on many other sets too, and every quantifier lives on the domain of discourse D (to see this, observe that every $P = P \cap D$, hence intersection with D can make no difference as to whether T applies). Concretely: *Some cat purrs* is equivalent to *Some cat is a thing that purrs*.

The restriction of a quantifier T can then be defined as T's unique smallest live-on set, or equivalently as the intersection of all the live-on sets (von Stechow and Zimmermann, 1984, pg. 28), as in (10).

(10) **Quantificational restrictions.**
$\mathbf{R} = \lambda T_{\langle et \rangle t}. \bigcap \mathbf{L}(T)$

The **R** operator takes a generalized quantifier T and delivers the intersection of every set that T lives on. Because every T lives on the intersection of its live-on sets, $\mathbf{R}(T)$ is always a member of $\mathbf{L}(T)$ and is hence also always the smallest member of $\mathbf{L}(T)$. I prove this claim in the appendix, but the basic idea is that, for any P that T applies to, T will also apply to P's intersection with any live-on set. By the logic of **L**, that set can subsequently be intersected with any other live-on set to create another set that T applies to; and so on.

To define **R**, it is not necessary to formally define a separate **L** operator: we could just as easily have defined **R** as in (11).

(11) $\mathbf{R} = \lambda T_{\langle et \rangle t}. \bigcap \{S | \forall P : T(P) \leftrightarrow T(P \cap S)\}$

However, having separate reference to **L** is expositionally useful to illustrate the connection between quantificational restriction and conservativity, and **L** plays a role in some of the proofs in the appendix. In particular, the fact that $\mathbf{R}(T)$ is itself a member of $\mathbf{L}(T)$ makes it possible to prove a variety of other results pertaining to the restriction operators.

The restriction of the generalized quantifier terms in (12) is $\{a, b\}$, and the restriction of the generalized quantifier terms in (13) is the set of linguists.

(12) a. Agnes or Beatrice
b. Agnes and Beatrice
c. Agnes but not Beatrice
d. neither Agnes nor Beatrice

(13) a. some linguists
b. most but not all linguists
c. not a lot of linguists
d. exactly 2 linguists

Here is a proof sketch that $\mathbf{R}(T) = \{a, b\}$ for $T = [\![\textit{Agnes or Beatrice}]\!]$. $\mathbf{R}(T)$ must contain Agnes, because otherwise T would apply to $\{a\}$ but not $\{a\}$'s intersection with $\mathbf{R}(T)$. Similar reasoning holds for $\{b\}$. Suppose counterfactually that $\mathbf{R}(T)$ included Charmian: this could be true in one of two situations. 1) There is some set P such that $T(P)$, $P(c)$, and $\neg T(P - \{c\})$, hence c must appear in $\mathbf{R}(T)$ so as to ensure that $T(P) = T(P \cap \mathbf{R}(T) = 1$. 2) There is some set P such that $T(P)$, $\neg P(c)$, and $\neg T(P \cup \{c\})$, hence c must again appear in $\mathbf{R}(T)$ so as to ensure that $T(P \cup \{c\}) = T(P \cup \{c\} \cap \mathbf{R}(T)) = 0$. However, adding or removing Charmian from some set cannot change whether this set contains Agnes or Beatrice, so neither situation is attested. Similar reasoning holds for every individual in the domain of discourse besides Agnes and Beatrice. This shows that the restriction of *Agnes or Beatrice* is precisely $\{a, b\}$. It can be shown along similar lines that the restrictions of the other quantifiers are what I claim they are.

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

- Exhaustive interpretation of an answer T with unrestricted extent regulates the background property's intersection with the entire domain of discourse.
- Exhaustive interpretation of an answer T with restricted extent regulates the background property's intersection with $\mathbf{R}(T)$.

This characterization of extent, which requires access to a generalized quantifier meaning, would not have been possible in a purely propositional setting. The structured account of question/answer exchanges is what allows us to import tools from the logic of generalized quantifiers into the analysis of exhaustivity. Because propositions do not have a polarity, the observation that polarity affects exhaustivity supports the view that focused and backgrounded content can be formally distinguished at the level of information structure. von Stechow and Zimmermann, 1984 reach similar conclusions.

2.2.2 Positive and negative restrictions

I have suggested that exhaustive interpretation should only license conclusions regarding the question-property's intersection with a marked answer's quantificational restriction. To get direction right, however, we need to draw further distinctions among the objects in the restriction. This is because the direction of exhaustivity varies according to the semantics of the answer.

(14) **A**: Who attended the last panel?
B: Most linguists but not a lot of computer scientists.
$\rightsquigarrow$ *not all linguists, some computer scientists*

In chapter 1 I argued that exhaustivity must apply uniformly to B's answer. It cannot be applied separately to each conjunct. Exhaustive interpretation of the first conjunct would typically license the conclusion that no computer scientists attended, yet we want exhaustive interpretation of the second conjunct to license the conclusion that some did. Appealing to cancellation does not work because the content of this latter inference is never asserted; it is an implicature driven by exhaustivity.

This clash is symptomatic of the larger problem that exhaustive interpretation of the first conjunct would typically have unrestricted extent. However, B's answer in (14) can only be interpreted with restricted extent. Instead of generating and selectively canceling exhaustive inferences, I will develop a recipe for exhaustivity that can uniformly account for mixed direction.

The restriction of B's answer in (14) is the set containing the linguists and the computer scientists. Exhaustive interpretation licenses the conclusion that not every linguist attended, but some computer scientists did. To deliver this implication, exhaustivity must interact differentially with different subsets of the restriction. We want to minimize the question-property's intersection with the set of linguists, and maximize its intersection with the set of computer scientists, as much as can be done consistently with the answer's literal meaning. This finding indicates that the set of linguists and the set of computer scientists must be formally recoverable from the meaning of B's answer in (14), as separate sets.

The idea I would like to pursue is that the objects in a quantifier's restriction can be classified as "positive" or "negative". The POSITIVE RESTRICTION of a nonmonotonic quantifier like *Most linguists but not a lot of computer scientists* should consist of the linguists, and the NEGATIVE RESTRICTION should consist of the computer scientists. To my knowledge, the idea of sorting a quantifier's restriction by polarity has not been explored in the GQ literature. To the extent that the proposal I develop is successful, it could potentially have some currency in other settings where generalized quantifiers are used.

A formal characterization of positive and negative restrictions should fulfill at least the desiderata in (15).

(15)
a. **Increasing** quantifiers should only have a positive restriction; their negative restriction should be empty.
b. **Decreasing** quantifiers should only have a negative restriction; their positive restriction should be empty.
c. **Nonmonotonic** quantifiers should have nonempty positive and negative restrictions.

Exhaustive interpretation of marked nonmonotonic answers like *Most linguists but not a lot of computer scientists* would then involve minimizing the question-property's intersection with the answer's positive restriction (the linguists), and maximizing the question-property's intersection with the answer's negative restriction (the computer scientists). Exhaustive interpretation of decreasing answers like *Not all of the computer scientists* would involve maximizing the question-property's intersection with the answer's negative restriction (the computer scientists).

Decreasing quantifiers have empty positive restrictions, so these reasoning schemata could be unified into a single schema: minimize the question-property's intersection with the positive restriction and maximize its intersection with the negative restriction. If one of these sets is empty, regulating its intersection with the question-property will be vacuous; otherwise, exhaustive implications are licensed. This schema will provide the basis of the HYBRID ORDER.

How should we define the restriction operators? The positive restriction should consist of objects whose presence in a set should in some sense correlate with that set's presence in the quantifier denotation, but the most obvious ways of making this idea formally precise do not deliver the desired result. The lambda term in (16-a) takes as input a generalized quantifier T and delivers a function that is true of an object x if every property in T contains x. This works for universal quantifiers (every property in the denotation of *every cat* contains every cat) but not existential quantifiers (some properties in the denotation of *some cat* do not contain every cat).

(16) a. $\lambda T_{\langle et \rangle t} \lambda x_e . \forall P : T(P) \rightarrow P(x)$
b. $\lambda T_{\langle et \rangle t} \lambda x_e . \forall P : P(x) \rightarrow T(P)$

The lambda term in (16-b) takes as input a generalized quantifier T and delivers a function that is true of an object x if every property that contains x appears in T. This works for existential quantifiers (any property that contains a cat appears in the denotation of *some cat*) but not universal quantifiers (not every property that contains a cat appears in the denotation of *every cat*). Yet increasing existential and universal quantifiers over cats should both be positively restricted by the set of cats.

The equivalent negative restriction operators fail for analogous reasons. The lambda term in (17-a) takes as input a generalized quantifier T and delivers a function that is true of an object x if no property in T contains x. This works for negated existentials (no property in the denotation of *no cat* contains a cat) but not negated universals (some properties in the denotation of *not every cat* contain cats).

(17) a. $\lambda T_{\langle et \rangle t} \lambda x_e . \forall P : T(P) \rightarrow \neg P(x)$
b. $\lambda T_{\langle et \rangle t} \lambda x_e . \forall P : \neg P(x) \rightarrow T(P)$

The lambda term in (17-b) takes as input a generalized quantifier T and delivers a function that is true of an object x if T contains every property that does not contain x. This works for negated universals (any property that does not contain some arbitrary cat appears in the denotation of *not every cat*) but not negated existentials (some properties that don't contain every cat fail to appear in the denotation of *no cat*). Yet negated existential and universal quantifiers over cats should both be negatively restricted by the set of cats.

The culprit for the inadequacy of the (16) and (17) definitions is apparently the presence of universal quantification. It is not generally the case that every property in a generalized quantifier's denotation needs to treat a particular object in the restriction in a certain way, or that the presence or absence of any particular object in the restriction is sufficient to determine whether the quantifier applies to a property.

The logic of quantificational restriction suggests a solution, however. Because every quantifier lives on its restriction, any T is closed under intersection with $\mathbf{R}(T)$: if P appears in T, so does P's intersection with $\mathbf{R}(T)$. An object x appears in $\mathbf{R}(T)$ if x's presence is necessary for intersective closure to hold: if x did not appear in $\mathbf{R}(T)$, there would be at least one P in T whose intersection with $\mathbf{R}(T)$ doesn't also appear in T, or conversely at least one P outside T whose intersection with $\mathbf{R}(T)$ does appear in T. That is, while universal quantification over the elements of T is generally too strong for the notions of positive and negative restriction we are after, existential quantification is just right. For every x in the restriction, there should be at least one property whose status regarding T is affected by x's presence or absence.

The definitions we want are those in (18). A generalized quantifier's POSITIVE RESTRICTION consists of objects that can be added to a non-T-set to create a T-set. The NEGATIVE RESTRICTION consists of objects that can be subtracted from a non-T-set to create a T-set.

(18) **Restriction operators.**
a. $\mathbf{R}^+ = \lambda T_{\langle et \rangle t} \lambda x_e . \exists S : \neg T(S) \wedge T(S \cup \{x\})$
b. $\mathbf{R}^- = \lambda T_{\langle et \rangle t} \lambda x_e . \exists S : \neg T(S) \wedge T(S - \{x\})$

There are several reasons why an object x might appear in a quantifier's restriction. It might be that some quantifier T contains some S that contains x, but subtracting x from S results in a set that T no longer contains. Then x must appear in the restriction or else intersective closure would be violated: T would apply to S but not the restriction's intersection with S. Objects that meet this condition are included in $\mathbf{R}^+(T)$, the positive restriction. It might also be that T contains some S that does not contain x, but adding x to S results in a set that T no longer contains. Then x must appear in the restriction or else intersective closure would again be violated: T would apply to the restriction's intersection with S but not S itself. Objects that meet this condition are included in $\mathbf{R}^-(T)$, the negative restriction.

The following useful facts about the restriction operators hold for all generalized quantifiers. These facts are proven in the appendix. Their validity shows that the restriction operators meet the desiderata in (15): positive and negative restriction interacts with monotonicity in the expected way.

(19) a. $\mathbf{R}^+(T) \cup \mathbf{R}^-(T) = \mathbf{R}(T)$ for all T.
b. Iff T is increasing, $\mathbf{R}^-(T)$ is empty.
c. Iff T is decreasing, $\mathbf{R}^+(T)$ is empty.

The fact that the union of any quantifier's positive and negative restriction always adds up to the total restriction is encouraging. It indicates that $\mathbf{R}^+$ and $\mathbf{R}^-$ indeed sort restrictions in the desired way, and that we are not missing any objects in the total restriction by confining our attention to the positive and negative subsets.

There is nothing special about the (18) definitions: we could equivalently adopt the definitions in (20). For any T, $\mathbf{R}^+(T)$ can also be thought of as a set of objects that can be subtracted from a T-set to create a non-T-set, and $\mathbf{R}^-(T)$ can also be thought of as a set of objects that can be added to a T-set to create a non-T-set.

(20) **Restriction operators (equivalent).**
a. $\mathbf{R}^+ = \lambda T_{\langle et \rangle t} \lambda x_e . \exists S : T(S) \wedge \neg T(S - \{x\})$
b. $\mathbf{R}^- = \lambda T_{\langle et \rangle t} \lambda x_e . \exists S : T(S) \wedge \neg T(S \cup \{x\})$

The former perspective emphasizes that adding (subtracting) an object to (from) a set is sufficient to qualify that set for T-membership, while the latter perspective emphasizes that an object's presence in (absence from) a set is necessary to maintain T-membership. These are different perspectives on the same two model-theoretic properties.

The (21) quantifiers are positively restricted by the set of linguists, and negatively restricted by nothing.

(21) $\mathbf{R}^+(T) = \textbf{linguists}, \mathbf{R}^-(T) = \varnothing$
a. some linguist $\lambda P.\textbf{linguists} \cap P$
b. every linguist $\lambda P.\textbf{linguists} \subseteq P$
c. most linguists $\lambda P.|\textbf{linguists} \cap P| > |\textbf{linguists} \cap \overline{P}|$

Every linguist can be added to a set that does not contain linguists and thereby create a set that *some linguist* applies to; every linguist can be added to a set that contains every linguist but one and thereby create a set that *every linguist* applies to; and every linguist can be added to a set that contains exactly half of the linguists and thereby create a set that *most linguists* applies to. However, it is impossible to create a set that contains some, every, or most linguists by subtracting objects from a set that does not already meet this condition, so nothing appears in the negative restriction.

The (22) quantifiers are positively restricted by the set containing Agnes and Beatrice, and negatively restricted by the empty set.

(22) $\mathbf{R}^+(T) = \{a, b\}, \mathbf{R}^-(T) = \varnothing$
  a. Agnes or Beatrice $\lambda P.P(a) \vee P(b)$
  b. Agnes and Beatrice $\lambda P.P(a) \wedge P(b)$

Agnes and Beatrice can both be added to a set that contains neither and thereby create a set that *Agnes or Beatrice* applies to; and Agnes and Beatrice can each be added to a set that only contains the other and thereby create a set that *Agnes and Beatrice* applies to. Because it is impossible to create a set that contains Agnes or/and Beatrice by only subtracting objects from a set that does not meet this condition, nothing appears in the negative restriction.

The (23) quantifiers are positively restricted by the empty set, and negatively restricted by the set of linguists.

(23) $\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \textbf{linguists}$
  a. no linguist $\lambda P.\textbf{linguists} \not\pitchfork P$
  b. not every linguist $\lambda P.\textbf{linguists} \not\subseteq P$
  c. not a lot of linguists $\lambda P.|\textbf{linguists} \cap P| < n$

Every linguist can be subtracted from a set that contains only that linguist and thereby create a set that *no linguist* applies to; every linguist can be subtracted from a set that contains every linguist and thereby create a set that *not every linguist* applies to; and every linguist can be subtracted from a set that contains one more linguist than the context-sensitive threshold for *a lot* (which I notate as n) and thereby create a set that does not contain a lot of linguists.[1] However, it is impossible to create a set that contains no, not every, or not a lot of linguists by only adding objects to a set that does not already meet these conditions, so nothing appears in the positive restriction.

The (24) quantifiers are positively restricted by the empty set, and negatively restricted by the set containing Agnes and Beatrice.

(24) $\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \{a, b\}$
  a. neither Agnes nor Beatrice $\lambda P.\neg P(a) \wedge \neg P(b)$
  b. not both Agnes and Beatrice $\lambda P.\neg P(a) \vee \neg P(b)$

Agnes and Beatrice can each be subtracted from a set that does not contain the other, thereby creating a set that satisfies *neither Agnes nor Beatrice*; and Agnes and Beatrice can each be subtracted from a set that contains both and thereby create a set that satisfies *not both Agnes and Beatrice*. It is impossible to create a set that lacks one or both by only adding objects to a set that does not already meet this condition, so nothing appears in the positive restriction.

In (25), *exactly 2 linguists* is both positively and negatively restricted by the set of linguists. Any set containing exactly one linguist can be turned into a set containing exactly two by adding a linguist, and any set containing exactly three linguists can be turned into a set containing exactly two by subtracting one.

(25) exactly 2 linguists $\lambda P.|\textbf{linguists} \cap P| = 2$
  $\mathbf{R}^+(T) = \mathbf{R}^-(T) = \textbf{linguists}$

[1] I assume that vague quantifiers contribute precise thresholds to truth-conditions that language users are unable to pinpoint for pragmatic reasons (Fara, 2000; Kennedy, 2007). This assumption is necessary for the restriction operators to correctly treat vague quantifiers.

Quantifiers like *exactly 2 linguists* show that $\mathbf{R}^+$ and $\mathbf{R}^-$ are not mutually exclusive.

In (26), *most linguists but not a lot of computer scientists* is positively restricted by the set of linguists and negatively restricted by the set of computer scientists. Every linguist can be added to a set containing half of the linguists and not a lot of computer scientists to create a set that *most linguists but not a lot of computer scientists* applies to. Every computer scientist can be subtracted from a set containing most linguists and one more linguist than the threshold for *a lot* to create a set that *most linguists but not a lot of computer scientists* applies to.

(26) most linguists but not a lot of computer scientists
$\lambda P.|\mathbf{linguists} \cap P| > |\mathbf{linguists} \cap \overline{P}| \wedge |\mathbf{computer\ scientists} \cap P| < n$
$\mathbf{R}^+(T) = \mathbf{linguists}, \mathbf{R}^-(T) = \mathbf{computer\ scientists}$

This example shows that $\mathbf{R}^+$ and $\mathbf{R}^-$ meet one of our main desiderata, which was to distinguish between the positive and negative restrictions of nonmonotone quantifiers like (26).

Finally, if natural language conditionals are analyzed using material implication, then (27) is positively restricted by the singleton containing Beatrice and negatively restricted by the singleton containing Agnes. (27) contains all sets except those that those that contain Agnes but not Beatrice. To transform such a set into one (27) applies to, we can either add Beatrice or subtract Agnes.

(27) if Agnes, then Beatrice $\lambda P.P(a) \rightarrow P(b)$
$\mathbf{R}^+(T) = \{b\}, \mathbf{R}^-(T) = \{a\}$

Conditionals lacking a predicate like the one in (27) are ungrammatical in English. The quantifier meaning in (27) is useful not as a denotation for actual linguistic expressions, but as a way to analyze novel/focused content at the level of information structure.

Most semantic literature on natural language conditionals assigns them a modal analysis that goes beyond what material implication can express. I do not address modal answers here. Since the goal of this thesis is to account for the direction and extent properties of answers, which are independent of modality, I have opted to adapt the "basic" circumscription approach developed by Schulz and van Rooij, 2006 instead of the more nuanced epistemic account of van Rooij and Schulz, 2004. The latter proposal successfully accounts for modal answers, however, and it would be straightforward to adapt it to the analysis I develop in this chapter. For now, I will show how a "basic" recipe for exhaustivity can predict conditional perfection for a subset of conditional answers.

Spector, 2007 defined a relation of "favoring" that holds between formulae and literals of propositional logic. The restriction operators recreate this notion within the logic of generalized quantifiers. Recall that a formula F of Spector's propositional logic L_Q *favors* a literal L if there is an interpretation w in which F and L are both true, and F is false in the interpretation that is exactly like w except that L is false. If we establish a correspondence between generalized quantifiers and propositions derived by applying a generalized quantifier to the question-property, then the objects in a quantifier's restriction will correspond to the literals favored by the corresponding formula. This is because the literals of L_Q denote propositions derived by applying the question-property to individuals in the domain of discourse.

Changing a positive literal of L_Q from true to false would mean that the corresponding object is subtracted from the question-property. If this has the consequence of changing a formula of L_Q from true to false, that would mean that the question-property is subtracted from the corresponding quantifier's denotation in virtue of missing the object. The positive literals favored by a formula thus stand in correspondence with the objects that appear in a quantifier's positive restriction.

Similarly, changing a negative literal of L_Q from true to false would mean that the corresponding object is added to the question-property. If this has the consequence of changing a formula of L_Q from true to false, this would mean that the question-property is subtracted from the corresponding quantifier's denotation in virtue of containing the object. The negative literals favored by a formula thus stand in correspondence with the objects that appear in a quantifier's negative restriction.

I do not provide an equivalence proof between Spector's system and mine. The requisite notion of "correspondence" is difficult or impossible to formalize. There is not generally a single unique quantifier that maps a particular property to a particular proposition. For this reason, there is no bijective mapping between the answers that are available in Spector's system and mine. Consider (28).

(28) a. $\lambda P.\mathbf{cats} \not\cap P$
b. $\lambda P.P \neq \mathbf{meowed} \rightarrow \mathbf{cats} \not\cap P$

The lambda term in (28-a) provides a reasonable denotation for the quantifier *no cats*. The lambda term in (28-b) behaves like *no cats* unless the input property is the property of meowing, in which case the output of (28-b) is trivially true. If the question is *Who purred?*, these quantifiers behave identically, yet are not equivalent in the general case. If some cats meowed, then (28-a) will map *meowed* to 0, but (28-b) will map it to 1.

To be sure, (28-b) is not a reasonable denotation for any answer that could be naturally expressed in language. The point is that, if A asks B *Who purred?* and B provides an answer denoting the proposition that no cats purred, there are no criteria for extracting a unique quantifier from this proposition. Still, without establishing formal equivalence, I take it as a positive result that the restriction operators can draw the same semantic distinctions as Spector's notion of favoring without recourse to a translation language.

2.2.3 Summary

Exhaustive interpretation of marked answers can be accounted for by restricting the extent of exhaustivity to the question-property's intersection with an answer's quantificational restriction. The restriction of a quantifier, **R**, can be defined formally as the intersection of the sets a quantifier lives on, which are in turn defined as all the sets that a quantifier's denotation is closed under intersection with (von Stechow and Zimmermann, 1984). I proposed that restrictions can be sorted by polarity and developed two novel restriction operators, $\mathbf{R}^+$ and $\mathbf{R}^-$, that sort a quantifier's restriction into (potentially disjoint, overlapping, or empty) positive and negative subsets. $\mathbf{R}^+$ and $\mathbf{R}^-$ were shown to make reasonable predictions for a wide range of quantificational expressions.

A successful account of direction requires explicit formal access to distinct subsets of a quantifier's restriction. The following section shows that the restriction operators deliver precisely the desired subsets. I use the restriction operators to define an order over background property extensions that can uniformly predict the positive, negative, and mixed directional properties of marked answers.

2.3 The hybrid order

The goal of this section is to define a complex order that captures the mixed directional properties of marked answers. I propose that exhaustive interpretation of marked answers minimizes the question-property's intersection with an answer's positive restriction, and maximizes the question-property's intersection with an answer's negative restriction. These operations apply simultaneously and are integrated into a single order over possible worlds, the HYBRID ORDER. If this order is plugged into Schulz and van Rooij's schema for exhaustivity, the resulting theory makes correct predictions about a wide range of negative and mixed answers. The difference between exhaustive interpretation of positive vs. negative/mixed answers then depends only on whether the default or hybrid order is chosen.

The hybrid order can be seen as an elaboration of circumscription. Although the default inclusion-induced order is unsuitable for marked answers, the hybrid order turns out to use the same formal techniques. The empirical success of the hybrid order in describing the directional properties of negative and mixed answers thus supports the general view that exhaustivity should be formally characterized in terms of circumscription.

2.3.1 Restriction-induced order

Recall that exhaustivity is implemented as in (29). Its output is true in only the minimal models in which the answer is true, where minimality is calculated using an order $<$.

(29) $\texttt{exh}(\langle T, P\rangle) = T(P_w) \wedge \neg\exists v : T(P_v) \wedge v < w$

I propose that $<$ is an indeterminate parameter. It can be resolved to either the DEFAULT ORDER or the HYBRID ORDER. The default order is the inclusion-induced order we have already seen (30).

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

The hybrid order is given in (31). It is defined using the restriction operators, which are repeated in (32).

(31) **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 $w \not\leq_{\texttt{hybrid}} v$

(32) **Restriction operators.**
a. $\mathbf{R}^+ = \lambda T_{\langle et\rangle t}\lambda x_e.\exists S : \neg T(S) \wedge T(S \cup \{x\})$
b. $\mathbf{R}^- = \lambda T_{\langle et\rangle t}\lambda x_e.\exists S : \neg T(S) \wedge T(S - \{x\})$

The hybrid order is calculated on the basis of two more basic orders, $<^+$ and $<^-$. These orders use the restriction operators $\mathbf{R}^+$ and $\mathbf{R}^-$ to rank possible worlds by the size of the question-property P's intersection with a term answer T's positive and negative restrictions. A world v is more $<^+$-minimal than w if P's intersection with $\mathbf{R}^+(T)$ in v is smaller. Similarly, v is more $<^-$-minimal than w if P's intersection with $\mathbf{R}^-(T)$ in v is larger. Overall, v is more $<_{\texttt{hybrid}}$-minimal than w if either $v <^+ w$ or $v <^- w$ as long as $<^+$ and $<^-$ do not conflict. In this way, the hybrid order simultaneously integrates positive and negative information into a single metric.

If $\mathbf{R}^+(T)$ is empty, as is the case for all decreasing answers, only $<^-$ matters. The empty set is never a strict subset of itself, so no world counts as more $<^+$-minimal than another world if a decreasing answer is given. However, mixed answers have nonempty $\mathbf{R}^+$ and $\mathbf{R}^-$ sets, so $<^+$ and $<^-$ have the potential to conflict.

To see why this is potentially a problem, consider an exchange like *A: Who attended? B: If Agnes attended, then Beatrice attended*. The $\mathbf{R}^+$ set for B's answer is the set containing Beatrice, and the $\mathbf{R}^-$ set is the set containing Agnes. The literal meaning of the answer excludes worlds in which Agnes attended and Beatrice didn't. Minimizing the question-property's intersection with the $\mathbf{R}^+$ set leads $<^+$ to prefer worlds where nobody attended to worlds where they both attended or only Beatrice attended. Maximizing the question-property's intersection with the $\mathbf{R}^-$ set leads $<^-$ to prefer worlds where they both attended to worlds where nobody attended or only Beatrice attended. There are no worlds that are simultaneously $<^+$-minimal and $<^-$-minimal. If both $<^+$ and $<^-$ are used, `exh` will exclude all possible worlds consistent with the answer.

This dilemma arises because $<^+$ prefers worlds where nobody attended to worlds where they both attended, and $<^-$ prefers worlds where they both attended to worlds where nobody attended. Neither of these (equivalence classes of) worlds should be excluded; we only want to exclude worlds where only Beatrice attended, which are not minimal according to either order. To get the right result, we need to integrate $<^+$ and $<^-$ into a relation that cannot simultaneously apply to a pair of distinct objects and its dual.

There is a name for this property: antisymmetry. The hybrid order defined in (31) is antisymmetric: conflicting pairs "cancel out". This property plays an important role in capturing the mixed direction of nonmonotonic answers. In the case of conditional answers, the only remaining pairs will be those that rank worlds where both attended or worlds where neither attended as more minimal than worlds where only Beatrice attended. Exhaustivity will only exclude the latter worlds, yielding the inference that Beatrice attended if and only if Agnes attended, the correct result for a conditional answer.

Compared to the default order, there are several formal properties of the hybrid order that plausibly contribute to its markedness. The default order amounts simply to the Boolean lattice over property extensions induced by inclusion. Although this order can only be used with default answers, the semantics of the answer itself do not play a role in calculating the order. We do not need access to model-theoretic properties of an answer, like its positive and negative restrictions and their intersection with P, to determine whether the default order prefers one world over another; only the overall extension of P matters. Moreover, all objects in the domain of discourse matter for the default order, which means that exhaustive interpretation can license stronger conclusions using the default order because objects outside an answer's quantificational restriction are excludable. The default order at least has the potential to produce a fully exhaustive answer.

The hybrid order is more complex, and it licenses weaker conclusions. It requires language users to track two orders over property extensions rather than one, as well as navigate potentially redundant or conflicting information. It cannot be calculated solely on the basis of the question-property P, but also interacts nontrivially with the semantics of the answer T. The default order will treat any two worlds v and w consistently as long as P is fixed, but the hybrid order might rank $v <_{\texttt{hybrid}} w$, $w <_{\texttt{hybrid}} v$, or neither depending on the choice of T. Finally, because of the domain restriction to $\mathbf{R}^+(T)$ and $\mathbf{R}^-(T)$, the hybrid order cannot license any conclusions about objects outside T's quantificational domain. Hence the hybrid order cannot even in principle produce an exhaustive answer except when it is contextually entailed that the question-property is included in an answer's restriction. These features all make it plausible that the default order is preferred.

The closest algebraic precedent to the hybrid order is probably the BILATTICE (Belnap, 1975; Fitting, 1991; Ginsberg, 1988, 1990), a structure used in the literature on nonmonotonic logic and symbolic computation to characterize reasoning in the face of conflicting information. A bilattice is a set equipped with two separate partial orders. These are usually construed as orders over the domain of truth values in multivalued logic, intended to simultaneously measure the amount of truth a formula has, and the amount of (potentially inconsistent) knowledge it is assigned. Unlike Boolean lattices, bilattices have not played a prominent role in natural language semantics, despite sporadic linguistic applications to attitude verbs (Muskens, 1989), presupposition and implicature (Schöter, 1996), and interrogative semantics (Nelken and Francez, 2002).

The hybrid order is not a true bilattice, for two reasons. First, the $<^+$ and $<^-$ orders are not guaranteed to be complete lattices. Second, a bilattice is supposed to come with a negation operation that reverses one of the partial orders (the "truth" order) while preserving the other one (the "knowledge" order). Such an operation could be defined for either $<^+$ or $<^-$, but as far as I can tell there is no reason to. Still, it is interesting that the notion of a multiply ordered set required for the analysis of marked answers was independently motivated in the same line of computational literature that developed circumscription.

Chapter 3 will investigate in detail which features of an answer determine whether the default or the hybrid order is used, and why. However, it should be clear that the use of the hybrid order is tied to the need to sort objects by polarity, which in turn is only necessary for negative and mixed answers. The hybrid order arises because conversational participants need a way to classify objects as positive or negative on the basis of a generalized quantifier meaning. The restriction operators provide such a basis for classification, but their use has restricted extent as a consequence.

The explanation for restricted extent shares some conceptual features with the structural theory's notion of "negation-induced symmetry". Recall that negation-induced symmetry occurred whenever two sentences, one the negation of the other, both count as formal alternatives to an answer. Because equally and less complex sentences count as formal alternatives according to the structural theory, the alternatives to a sentence containing negation will cancel out because of negation-induced symmetry. Both *Charmian* and *not Charmian* will count as alternatives to the answer *Agnes but not Beatrice*, preventing inferences from being drawn about Charmian.

The structural theory enforces this result by regulating the syntactic distribution of negation, and quantificational restriction is not a syntactic notion. The structural theory thus cannot generate exhaustive implications about objects inside an answer's restriction either, and does not explain why negative and mixed answers license these implications. In the present setting, this result follows from the fact that objects outside an answer's restriction cannot coherently be assigned a polarity, but objects inside the restriction can.

2.3.2 Negative answers

The hybrid order supports a principled and successful account of how negative answers are interpreted. Because all decreasing answers have empty $\mathbf{R}^+$ sets, to calculate the exhaustive interpretation of a decreasing answer it suffices to use the order induced by $<^-$, which maximizes the question-property's intersection with an answer's $\mathbf{R}^-$ set.

In what follows, I assume a domain of discourse with three individuals, Agnes, Beatrice, and Charmian. Agnes and Beatrice are linguists and Charmian is not.

The simplest cases are those like (33) where the meaning of the answer completely fixes the status of every object in its restriction relative to the question-property.[2]

(33) **A**: Who attended the last panel? $\lambda w \lambda x.\mathbf{attended}_w(x)$
B: Not Beatrice. $\lambda P.\neg P(b)$
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \{b\}$

The $\mathbf{R}^-$ set for *not Beatrice* is the set containing Beatrice. $<^-$ prefers worlds in which Beatrice attended to those in which she didn't, and otherwise enforces no additional structure on the domain of possible worlds. However, only worlds where she didn't attend are consistent with the answer. Exhaustivity has no effect in this case.

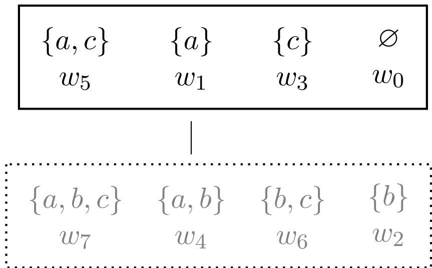

Figure 2.2: Not Beatrice.

Answers like (34) also fix the status of every individual in the restriction relative to the question-property, but induce a different order over worlds.

(34) **A**: Who attended the last panel? $\lambda w \lambda x.\mathbf{attended}_w(x)$
B: No linguists. $\lambda P.\mathbf{linguists} \not\cap P$
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \mathbf{linguists} = \{a, b\}$

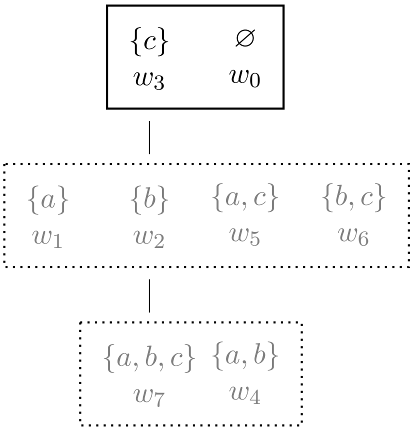

Figure 2.3: No linguists.

The $\mathbf{R}^-$ set for *no linguists* is the set containing Agnes and Beatrice. $<^-$ prefers worlds in which both Agnes and Beatrice attended to worlds in which just one of them attended, and prefers worlds in which just one of them attended to worlds in which neither of them attended. This is because there are multiple objects in the answer's restriction, allowing a more complex inclusion-induced order to arise. Although only worlds in which neither attended are consistent with the answer in this case, the example shows how the hybrid order over possible worlds is sensitive to the semantics of an answer.

In the context given, B's answers in (34) and (35) are extensionally equivalent. This will not always be the case across models, but exhaustivity likewise has no effect in (35) for the same reason.

(35) **A**: Who attended the last panel? $\lambda w \lambda x.\mathbf{attended}_w(x)$
B: Neither Agnes nor Beatrice. $\lambda P.\neg P(a) \wedge \neg P(b)$
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \{a, b\}$

[2] Here is how to read the diagrams in this section. Worlds that share a $\boxed{\text{box}}$ are unordered relative to each other. Lines between boxes indicate that every world in the lower box is more minimal than every world in the higher box. I sometimes draw directed lines labeled with $+$ or $-$ to indicate that every world in a box is more $<^+$- or $<^-$-minimal than every world in some other box; I only do this when $<^+$ or $<^-$ does not contribute to the overall order. I ~~cross~~ worlds excluded by exhaustivity and dim worlds inconsistent with the answer.

Although exhaustivity with the hybrid order has no effect in (33)-(35), it does have an effect when the answer does not fix the status of every individual in its restriction relative to the question-property, as in (36).

(36) **A**: Who attended the last panel? $\lambda w \lambda x.\textbf{attended}_w(x)$
B: Agnes and Beatrice didn't both attend. $\lambda P.\neg P(a) \vee \neg P(b)$
$\rightsquigarrow$ *one of them did*
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \{a, b\}$

The $\mathbf{R}^-$ set for *not both Agnes and Beatrice* is the set containing Agnes and Beatrice. The order given by $<^-$ is the same as that in Figure 2.3: $<^-$ prefers worlds in which they both attended to worlds in which just one of them attended, and prefers worlds in which just one of them attended to worlds in which neither attended. However, in this case there are more worlds to choose from: any world in which they didn't both attend is consistent with B's answer. Among these worlds, $<^-$ selects those in which one of them attended as minimal. Exhaustivity thus excludes worlds in which neither attended. In this way, the hybrid order accurately predicts the "indirect" implicature of B's answer that one of them attended the panel.

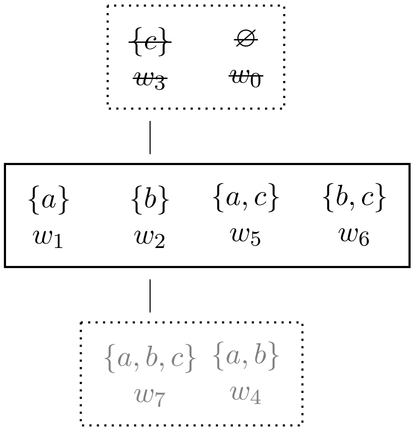

Figure 2.4: Not both Agnes and Beatrice.

Restricted extent is also captured. Charmian is outside the answer's quantificational restriction, and the presence of c in any of the possible extensions for P makes no difference as to how P is ranked vs. $P - \{c\}$.

The implicatures of B's answers in (37) and (38) can be predicted similarly, depending on the granularity of exhaustive interpretation and the threshold for *a lot*. In each case, $<^-$ will prefer worlds in which more linguists attended, and exhaustivity will exclude worlds in which fewer linguists attended from the denotation of the answer. The hybrid order thus predicts the implication in (37) that some linguists attended, and the implication in (38) that some or perhaps most linguists attended (depending on granularity).

(37) **A**: Who attended the last panel? $\lambda w \lambda x.\textbf{attended}_w(x)$
B: Not a lot of linguists. $\lambda P.|\textbf{linguists} \cap P| < n$
$\rightsquigarrow$ *some did*
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \textbf{linguists}$

(38) **A**: Who attended the last panel? $\lambda w \lambda x.\textbf{attended}_w(x)$
B: Not all of the linguists. $\lambda P.\textbf{linguists} \not\subseteq P$
$\rightsquigarrow$ *some did*
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \textbf{linguists}$

In the toy context given, these answers are not really distinguishable from (36), but exhaustivity will make more graded predictions in contexts with more linguists in the domain of discourse.

2.3.3 Mixed answers

So far, the hybrid order has reduced to $<^-$ when applied to negative answers. However, the utility of considering $<^+$ and $<^-$ at the same time reveals itself when mixed answers are taken into account.

In the case of mixed answers, it is convenient to assume a larger domain of discourse. Going forward I expand our toy model to include four individuals: Agnes and Beatrice are linguists, and Charmian and Doris are computer scientists.

I will warm up with another case where the meaning of the answer fixes the status of every object in its restriction relative to the question-property. Exhaustivity has no effect for (39), but it is expositionally useful to illustrate how the order works.

(39) **A**: Who attended the last panel? $\lambda w \lambda x.\mathbf{attended}_w(x)$
B: Agnes but not Beatrice. $\lambda P.P(a) \wedge \neg P(b)$
$\mathbf{R}^+(T) = \{a\}, \mathbf{R}^-(T) = \{b\}$

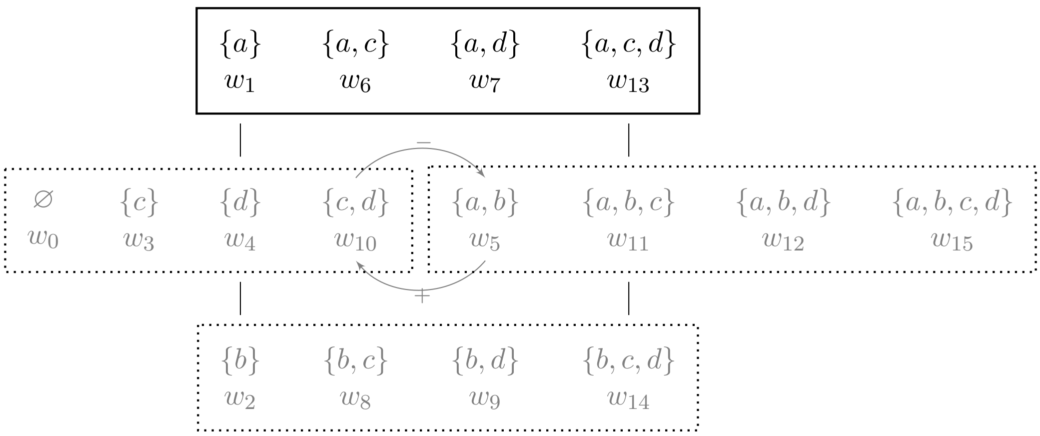

Figure 2.5: Agnes but not Beatrice.

The $\mathbf{R}^+$ set for *Agnes but not Beatrice* is the set containing Agnes, and the $\mathbf{R}^-$ set is the set containing Beatrice. The $<^+$ component of the hybrid order prefers worlds in which Agnes didn't attend to those in which she did, and the $<^-$ component prefers worlds in which Beatrice attended to those in which she didn't. Worlds in which they both attended and worlds in which neither attended are treated differently by $<^+$ and $<^-$: $<^+$ prefers worlds in which neither attended (since Agnes didn't attend in those worlds), and $<^-$ prefers worlds in which both attended (since Beatrice attended in those worlds). In Figure 2.5, the worlds in the left box in the second row are more $<^+$-minimal than those in the the right box, but the worlds in the right box are more $<^-$-minimal than those in the left box.

The hybrid order does not consider such cases in which $<^+$ and $<^-$ "disagree" and extracts only the consistent pairs, yielding the structure in Figure 2.5. Only worlds in which Agnes and Beatrice didn't attend are consistent with the answer, however, and the hybrid order does not distinguish between these. Charmian and Doris also do not affect the hybrid order, and so the exhaustive interpretation of B's answer in (39) is equivalent to its literal interpretation.

Exhaustive interpretation of B's answer in (40) does have an effect, however.

(40) **A**: Who attended the last panel? $\lambda w \lambda x.\mathbf{attended}_w(x)$
B: Some of the linguists but not all of the computer scientists.
$\lambda P.\exists x \in \{a, b\} : P(x) \wedge \neg \forall y \in \{c, d\} : P(y)$
$\mathbf{R}^+(T) = \{a, b\}, \mathbf{R}^-(T) = \{c, d\}$

The $\mathbf{R}^+$ set for *some of the linguists but not all of the computer scientists* is the set containing Agnes and Beatrice. The $\mathbf{R}^-$ set is the set containing Charmian and Doris. The hybrid order over possible worlds given this answer is shown in Figure 2.6. Worlds are preferred either in virtue of containing fewer members of the $\mathbf{R}^+$ set, or more members of the $\mathbf{R}^-$ set. For visual clarity I do not represent cases where $<^+$ and $<^-$ "cancel out" in Figure 2.6, but the reader may verify that every pair of unordered boxes is such a case.

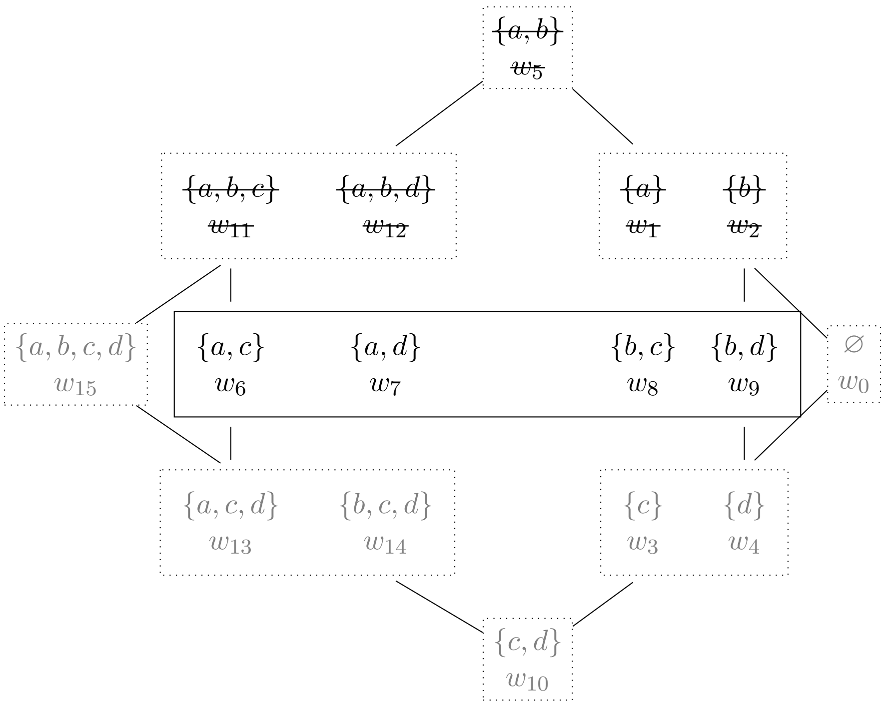

Figure 2.6: Some of the linguists but not all of the computer scientists.

As the boxes show, the worlds consistent with B's answer can be sorted into four groups. First we have worlds in which some but not all of the linguists attended and some but not all of the computer scientists attended (w_6, w_7, w_8, w_9); then we have worlds where all of the linguists attended and some but not all of the computer scientists attended (w_{11}, w_{12}); then we have worlds where some but not all of the linguists and no computer scientists attended (w_1, w_2); finally we have worlds where all of the linguists and no computer scientists attended (w_5). Worlds in the first group are the most minimal according to the hybrid order. The $<^+$ component of the hybrid order prefers these worlds to the group containing w_{11} and w_{12}, since fewer linguists ($\mathbf{R}^+$ members) attended in these worlds. The $<^-$ component of the hybrid order prefers these worlds to the group containing w_1 and w_2, since more computer scientists ($\mathbf{R}^-$ members) attended. Both $<^+$ and $<^-$ rule out w_5. Exhaustivity thus selects worlds where exactly one linguist and exactly one computer scientist attended, which is the desired prediction in this case.

This example shows that the hybrid order is capable of simultaneously minimizing the question-property's intersection with an answer's positive restriction and maximizing the its intersection with an answer's negative restriction. The requisite order over background property extensions is more complex than a simple inclusion relation, but can nevertheless be determined by combining multiple inclusion relations over different subsets of the background property.

Other mixed answers, like *Most linguists but not a lot of computer scientists*, can be treated in the same way. In the toy context given, this answer will not be truth-conditionally distinct from B's answer in (40), but in a model with more linguists and computer scientists the thresholds will be different. Applied to this answer, whose $\mathbf{R}^+$ set contains the linguists and whose $\mathbf{R}^-$ set contains the computer scientists, the hybrid order will similarly prefer worlds in which fewer linguists attended and more computer scientists attended, yielding the implication that most but not all of the linguists attended and not a lot but some computer scientists attended.

Finally, consider a conditional answer like (41).

(41) **A**: Who attended the last panel? $\lambda w \lambda x.\mathbf{attended}_w(x)$
B: If Agnes attended, then Beatrice attended. $\lambda P.P(a) \rightarrow P(b)$
$\mathbf{R}^+(T) = \{b\}, \mathbf{R}^-(T) = \{a\}$

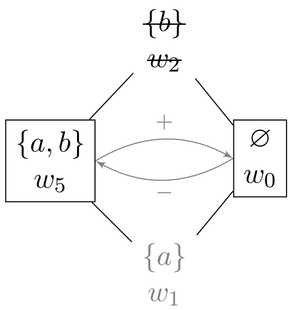

Figure 2.7: If Agnes, then Beatrice.

For simplicity of exposition I omit Charmian and Doris from the domain. Because of restricted extent, exhaustivity does not license any conclusions about them, so it is convenient to confine our attention to just those worlds where Agnes, Beatrice, both, or neither appear in the question-property's extension.

The $\mathbf{R}^+$ set for B's answer in (41) is the set containing Beatrice. The $\mathbf{R}^-$ set is the set containing Agnes. The hybrid order prefers worlds where Agnes attended to those where she didn't, as well as worlds where Beatrice didn't attend to those where she did. The literal meaning of the answer excludes worlds where only Agnes came. Among the remaining possible worlds the $<^+$ component prefers w_0, where nobody attended (because Beatrice, the single member of $\mathbf{R}^+$, did not attend), to the other two, and the $<^-$ component prefers w_5, where they both attended (because Agnes, the single member of $\mathbf{R}^-$, did attend), to the other two. The hybrid order prefers both these worlds to w_2, where only Beatrice attended, and accordingly exhaustivity excludes w_2 ($w_0 <^+ w_2, w_5 <^- w_2$).

However, because $<^+$ and $<^-$ conflict in their rankings of w_5 and w_0 ($w_0 <^+ w_5, w_5 <^- w_0$), neither pair is inherited by the overall hybrid order, which as a result ranks both of these worlds as minimal. Exhaustivity thus yields the implication that either Agnes and Beatrice both attended, or neither did, which is exactly the conditional perfection inference we were after.

There are some complications with conditional answers that the approach developed in this section cannot handle. I have deliberately confined my attention to conditional answers whose antecedent and consequent share the question-property as their main predicate. However, conditional perfection is a more general phenomenon (Horn, 2000) and a simple, circumscription-based account of exhaustivity does not cover all the cases. In particular, it cannot predict the conditional perfection inference in exchanges like (42).

(42) **A**: Who attended the last panel? $\lambda w \lambda x.\textbf{attended}_w(x)$
B: Doris attended if Charmian signed her permission slip.
$\lambda P.\textbf{signed}(c) \to P(d)$
$\mathbf{R}^+(T) = \{d\}, \mathbf{R}^-(T) = \varnothing$

The $\mathbf{R}^+$ set for B's answer in (42) is the set containing Doris. The $\mathbf{R}^-$ set is empty. Exhaustivity using the hybrid order minimizes the question-property's intersection with the answer's $\mathbf{R}^+$ set to the extent possible given the meaning of the answer. Since B's answer is consistent with Doris not attending, exhaustivity excludes all worlds in which Doris attended. The prediction is not adequate since in normal contexts B's answer cannot implicate that Doris didn't attend. The implication we want to predict is the usual conditional perfection inference that the antecedent and consequent share a truth value, i.e., Doris attended exactly if Charmian signed her permission slip. Because the antecedent contains a different predicate, unlike in (41) there is no way to achieve conditional perfection just by quantifying over possible background extensions of the question-property.

Schulz and van Rooij, 2006, aware of this problem, adopt an additional condition in their order over possible worlds: v cannot be more minimal than w unless all non-rigid terms occurring in an answer have the same extension at v and w. If this condition is adopted, worlds where Charmian signed Doris's permission slip cannot be excluded and conditional perfection is correctly predicted. Exhaustivity still excludes worlds where Charmian didn't sign the permission slip and Doris attended anyway, since these worlds are less minimal than worlds where Charmian didn't sign the permission slip and Doris didn't attend. The general idea is that exhaustivity should only minimize the extension of the question-property itself; it should not license additional conclusions about other properties.

Although Schulz and van Rooij's solution (using what I have called the default order) works for (42), it does not extend to (41), for which the hybrid order really is necessary. This condition also represents a weakening of nondetachability since the issue of which non-rigid terms appear in the answer is a syntactic matter, not a semantic one. One might attempt a semantic restatement of this condition in terms of modal similarity (Lewis, 1973): to count as more minimal than w, v must be exactly like w in all respects other than the extension of the question-property.

However, as van Rooij and Schulz, 2004 observe, such a condition would be too restrictive, since there are cases where exhaustive implications can serve as premises for further contextual entailments. Suppose it were known that Doris attended the panel if Charmian signed her permission slip. Exhaustive interpretation of the answer *Agnes and Beatrice* licenses the inference that Doris didn't attend, which in turn supports the conclusion that Charmian didn't sign her permission slip. If only maximally similar worlds could be compared using the default (or marked) order, this conclusion would be impossible to derive; background knowledge would prevent exhaustivity from excluding Doris. van Rooij and Schulz conclude that absolute modal similarity should not play a role in delimiting the order over possible worlds, which is why in Schulz and van Rooij, 2006 they restrict their similarity condition just to property terms that appear in an answer.

I do not extend the analysis of exhaustivity to conditional answers with different predicate terms in this thesis. Instead of adopting extra conditions on the order, the right solution is probably to leverage the modal nature of natural language conditionals (Kratzer, 1989; Lewis, 1973). van Rooij and Schulz's decomposition of circumscription into epistemic components intended to formalize the maxim of Quantity and the speaker competence assumption makes correct predictions about a wide range of answers containing modals and attitude verbs that a "basic" circumscription approach cannot cover (e.g., *Agnes and possibly Beatrice*). It should ideally be possible to extend this treatment to all varieties of conditionals given the right modal analysis of conditionals. The analysis I have developed in this chapter could be integrated into the van Rooij and Schulz, 2004 system by swapping out the default order for the hybrid order.

Despite these complications, I take it as a positive result that the present proposal correctly predicts conditional perfection when only the question-property is at stake.

2.3.4 Summary

Exhaustive interpretation of marked answers involves a form of reasoning that treats an answer's positive and negative restrictions differently. I proposed to formalize this reasoning using the HYBRID ORDER, which minimizes the question-property's intersection with an answer's positive restriction and maximizes the question-property's intersection with an answer's negative restriction. The hybrid order, which can be seen as the integration of two separate orders $<^+$ and $<^-$, makes correct predictions about the directional properties of a wide variety of negative and mixed answers.

Because the hybrid order requires reasoning using two partial orders instead of one, it is plausibly marked compared to the default order, and because the two orders it uses only regulate the question-property's intersection with different subsets of an answer's quantificational restriction, restricted extent is correctly predicted as well. Restricted extent is the hallmark of reasoning using the hybrid order.

The results discussed in this chapter so far all pertain to questions based on abstracts that denote first-order properties of individuals. The goal of the following section is to extend the account to a wider variety of questions, in particular questions whose abstracts have higher types.

2.4 Type and the symmetry problem

Circumscription addresses the symmetry problem by manipulating sets of objects drawn from an uncomplemented domain. Circumscribing a first-order property of individuals using the default order can only license exhaustive implications with negative direction. Circumscribing a first-order property of individuals using the hybrid order determines direction on the basis of an answer's positive and negative restriction. In both cases, the direction of exhaustivity is built into the order, so the symmetry problem does not arise.

This result is made possible by the fact that the domain of atomic individuals lacks Boolean structure. If the domain of individuals were complemented, so that any first-order property of individuals had to contain either some object a or $\overline{a}$, it would no longer be possible to exclude a or $\overline{a}$ from the extension of a first-order property using circumscription. There would be one minimal world in which the property contained a, and one minimal world in which the property contained $\overline{a}$. Whenever exhaustivity does lead to the exclusion of some object from a property's extension, then, that object cannot have a complement that is also excluded. The only way to exclude an object from a complemented domain would be to minimize the question-property's intersection with some contextually restricted subset of that domain, rather than minimizing the entire extension.

Questions with first-order abstracts do not lead to this situation, because the domain of individuals is not complemented. "Negative objects" do not exist in classical logic.[3] The domain of classical properties is complemented, however, so questions whose abstracts have higher types are predicted to interact differently with contextual information. It should be possible for circumscription to quantify either over a set that contains some property, or a set that contains its complement, without the other.

This section shows that these predictions are correct. Exhaustive interpretation of answers to questions with higher types does interact differently with context. I will argue the case in detail for questions whose abstracts denote second-order properties.

2.4.1 Second-order questions

Exhaustivity and the default order are repeated in (43) and (44).

(43) **Exhaustivity.**
$\texttt{exh}(\langle T, P\rangle) = T(P_w) \wedge \neg\exists v : T(P_v) \wedge v < w$

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

The symmetry problem arises whenever negation applies to the members of a contextually determined set of objects drawn from a complemented domain. This is why propositional accounts of exhaustivity cannot break symmetry semantically. As long as a proposition and its complement are both possible alternatives, some criteria must determine which is actually excluded, and because contextual domain restriction is usually treated as an unconstrained pragmatic process, ignoring the complement via domain restriction is insufficient in cases where symmetry is systematically broken in only one direction.

Circumscription provides a semantic solution to the symmetry problem by shifting the locus of quantification from the propositional level to the property level. Individuals can be excluded from a set without having to include some other individual. I have argued that this feature of circumscription enables a principled and empirically successful account of direction. Essentially, circumscription views the directional properties of answers as emerging definitionally from the nature of the exhaustive operation. In contrast, theories that use formal alternatives assume an underspecified exhaustive operation at the propositional level, which must then be supplemented with an independent theory of alternatives.

When the question-property has a higher type, however, symmetry has the potential to reappear. This is because at least some second-order properties of individuals must either contain some particular first-order property or its complement. The generalized quantifier over Beatrice must either contain the property of being human or the property of being nonhuman. As a result, if a question is asked whose abstract denotes the set of properties that Beatrice has (e.g., *What's Beatrice like*?), and an answer is given whose exhaustive interpretation licenses the conclusion that Beatrice is not human, the set of properties excluded from the abstract must contain the property of being human and cannot contain the property of being nonhuman, or else we would derive a contradictory conclusion.

[3] Bledin, 2025 develops a version of truthmaker semantics in which there are negative objects. However, negative objects do not actually appear in the extensions of predicate terms in Bledin's system; they are more like a compositional device, whose presence interacts with Bledin's composition rules in a way designed to make truthmaker conjunction compatible with NP coordination. If truthmaker semantics turns out to provide the right analysis of linguistic conjunction, then, it is still compatible with the solution to the symmetry problem advocated here. What matters is not whether negative objects are included in a model, but whether predicates can apply to them.

The relationship between type, exhaustivity, and contextual relevance is discussed extensively in the work of Szabolcsi, 1981a,b, 1983 on focus in Hungarian. In particular, the goal of Szabolcsi, 1983 is to understand the exhaustive interpretation of sentences with focused bare nouns, which she analyzes as property-denoting.

Focus in Hungarian involves movement to the front of a clause, and is associated with exhaustive semantics similar to the meaning of a cleft; Szabolcsi's semantics for exhaustivity involves minimizing a property extension in a way that is somewhat more complex, but similar in content to Groenendijk and Stokhof's implementation. The key sentences are those in (45)-(48).

(45) Biciklit$_F$ látott Mari.
Bikes.Acc saw Mari
'What Mari saw were bikes.'

(46) Biciklit$_F$ nem látott Mari.
Bikes.Acc not saw Mari
'What Mari didn't see were bikes.'

(47) Biciklit$_F$ látott mindenki.
Bikes.Acc saw everyone
'What everyone saw were bikes.'

(48) Biciklit$_F$ látott Mari kettőt.
Bikes.Acc saw Mari two.Acc
'What Mari saw two of were bikes.'
(Szabolcsi, 1983, pg. 130-131, ex. 13-15, pg. 135, ex. 25)

These sentences entail that everything that satisfies the scope property (analogous to the question-property in question-answer exchanges) also satisfies the focused property: (45) means that Mari didn't see any nonbikes, for instance. For this reason, these sentences raise the same problem of formally characterizing exclusive content that arises in the analysis of exhaustivity. The trick is to predict the attested meaning using a more general semantics for exhaustivity in terms of minimizing the extension of a potentially higher-order property. This is impossible unless the set of properties to be minimized is contextually restricted to a strict subset of the domain of properties.

In the case of (45), the set of properties that include things Mari saw cannot be minimized so as to contain only **bike**; it must also contain any superset of **bike** (**vehicle**, **thing**, etc). Since Szabolcsi's semantics for exhaustivity does not guarantee consistency preservation, (45) is predicted to either be false, or trivially true if the domain of discourse only contains bikes. The minimal models approach of Schulz and van Rooij, 2006 is hardwired to avoid this sort of result, but does not automatically generate the desired interpretation: for (45) to entail that Mari didn't see any cars, **car** must be excluded from the set of properties that contain things Mari saw, but the presence of the symmetric **noncar** in the domain of properties will prevent this.

Similarly, in the case of (46) the set of properties that don't contain objects Mari saw cannot only include **bike**. This set must also contain at least $\{x\}$ for any specific object x that she didn't see. If such singleton properties are considered, then (46) can only be true if Mari saw every nonbike in the domain of discourse. For similar reasons, if **thing** is considered then (47) can only be true if the domain of discourse only contains bikes. Szabolcsi's solution is to invoke domain restriction: exhaustivity does not shrink the scope property (question-property)'s extension absolutely but only minimizes its intersection with a contextually determined set of relevant properties.

Exhaustive interpretation relative to first-order properties also exhibits this sensitivity to context: if both Agnes and Beatrice attended the panel, exhaustive interpretation of the answer *Agnes* to the question *Who attended the panel?* can be taken as either true or false depending on whether Beatrice is relevant. However, there is an important difference. Agnes can appear in a model whose domain of individuals does not contain Beatrice, but any model with bikes also contains vehicles. If the entire set-theoretic universe is considered, first-order properties can still be minimized, but exhaustivity relative to higher-order properties is no longer possible. Higher-order properties therefore call for an even stronger variety of contextual restriction.

Examples like (48) show that this variety of domain restriction is independently necessary. Aside from the issue of avoiding triviality, exhaustive implications can be judged true or false depending on how the universe is partitioned. Szabolcsi considers a context in which Mari saw two bikes, a limousine, and a taxi. Is (48) true or false? Szabolcsi shows that both judgments are possible in this context, depending on which properties are relevant. If **car** is relevant then speakers will judge (48) as false. If **limousine** and **taxi** are relevant but not **car** then speakers will judge (48) as true. These observations indicate that the properties treated as relevant in a context should be allowed to vary freely.

Szabolcsi's discussion of this issue directly anticipates contemporary work on the symmetry problem. In the case of first-order properties:

> *Here we are not prima facie forced to acknowledge the role of the context in the formalism, however, because the two factors may as well be collapsed into one, by directly identifying the universe with the set of currently relevant individuals... The reason why from then on no problem arises lies in the fact that since there are no 'subindividuals' or 'superindividuals' no routine set theoretic operation may force us to take into consideration any individual which was not originally included among the relevant ones.*
> (pg. 139)

In the case of second-order properties, things are different:

> *...whether a property is relevant may vary from context to context but also has consequences for what happens within each such context. Namely, by singling out certain properties as relevant one actually kind of individuates them, that is, one disregards the fact that their extensions overlap with the extensions of other properties. In general, neither the union nor the intersection of two relevant properties needs to be a relevant property.* Relevance is not Boolean.
> *In set theory, however, properties are by no means individualistic in this sense and therefore in the course of testing the truth of a sentence there is no way to discriminate between properties that come to picture on their own right and properties that are just derived from the former by taking unions or intersections. Hence the relevant/irrelevant distinction for properties can no longer be made by manipulating the universe itself; consequently, the intended meaning of a CN-focus sentence* cannot even be approximated *unless the truth of the sentence is made relative to some context as well as some state of affairs.*
> (pg. 139-140)

This is exactly the perspective on symmetry and alternatives I have advocated throughout this thesis. Rooth, 1992 reaches similar conclusions in his discussion of *only*. Although Szabolcsi does not address the symmetry problem specifically, her discussion highlights the difference between complemented (Boolean) and uncomplemented domains implicit in the formal analysis of exhaustivity.

The orders that matter for exhaustivity should then be revised as follows, where D is a domain variable. The focus and background arguments T and P should now be viewed as unspecified for type. Szabolcsi only discusses inferences generated using the default order, but similar conclusions apply to the hybrid order.

(49) **Exhaustivity.**
$\mathtt{exh}(\langle T, P\rangle) = T(P_w) \wedge \neg\exists v : T(P_v) \wedge v < w$

(50) **The default order (general).**
$v <_{\mathtt{default}} w$ iff $P_v \cap D \subset P_w \cap D$

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

(52) **Restriction operators (general).**
a. $\mathbf{R}^{+} = \lambda T_{\langle\tau t\rangle t} \lambda x_{\tau}.\exists S : \neg T(S) \wedge T(S \cup \{x\})$
b. $\mathbf{R}^{-} = \lambda T_{\langle\tau t\rangle t} \lambda x_{\tau}.\exists S : \neg T(S) \wedge T(S - \{x\})$

The restriction operators, which were previously defined as applying only to type $\langle et\rangle t$ quantifiers, should also be restated in a type-general way as in (52).

The need to explicitly state exhaustivity relative to a contextually determined domain should not be surprising. I have argued that exhaustive interpretation is formally no different from any other natural language quantification, and all other varieties of quantification exhibit this need (von Fintel, 1994; Westerståhl, 1984).

2.4.2 Default answers to second-order questions

The approach in (50) and (51) predicts that questions with first-order vs. second-order abstracts allow different degrees of symmetry breaking. Regardless of which individuals are relevant, exhaustive interpretation of answers to questions with first-order abstracts should always have negative direction when using the default order; symmetry should be unbreakable by context. In contrast, for any particular property P, exhaustive interpretation of answers to questions with second-order abstracts should vary in whether it licenses the conclusion that P is excluded from or included in the abstract, depending on whether D contains P or $\overline{P}$. Context should be allowed to break symmetry, but only when the abstract has a higher type.

Fox and Katzir, 2011 have argued that this prediction is not borne out. They claim that symmetry is still systematically broken in only one direction even when exhaustivity applies to a set of objects drawn from a complemented domain, and that the alleged inability of context to break symmetry motivates an independent theory of formal alternatives.

Fox and Katzir's goal is to motivate the claim that domain restriction variables like D in (50) and (51) are subject to Boolean closure conditions. They argue that antecedent sets for contextual domain restrictions, at least for exhaustivity and focus-sensitive quantifiers, are necessarily closed under complementation and intersection operations. Fox and Katzir do not say how this constraint should be implemented, but it could potentially be viewed as a presupposition. Although they assume that exhaustivity is uniformly propositional, their constraint would pose a problem for the approach to exhaustivity I have advocated.

The key examples motivating Fox and Katzir's proposal have the form in (53) (I have revised their original examples to fit the question/answer format).

(53) **A**: Who attended last week's panel?
B: Some but not all of the linguists.
A: What about this week's panel?
B: Some of the linguists.
$\not\rightsquigarrow$ *all of the linguists attended*

The argument goes as follows. The surrounding linguistic context makes salient the proposition that some but not all of the linguists attended last week's panel. One might then naïvely expect it to be possible for the alternatives that matter for exhaustivity to include the proposition that some but not all of the linguists attended this week's panel, without including the symmetric proposition that all of the linguists attended this week's panel. In that case, exhaustive interpretation of B's answer to A's second question would implicate that "some but not all of the linguists" didn't attend (and hence all of the linguists attended). This implication is unattested, however.

Instead, Fox and Katzir claim that the set of contextually relevant propositions must be closed under Boolean operations: if the proposition that some but not all of the linguists attended is relevant, then so is the proposition that all of the linguists attended. Of these, exhaustivity only excludes the relevant propositions expressible by licit formal alternatives in the sense of Katzir, 2007. These include the proposition that all of the linguists attended, since this sentence can be expressed by a sentence no more complex than B's answer; they also include the proposition that some but not all of the linguists attended, since this sentence is salient in the surrounding linguistic context. These alternatives cannot be simultaneously excluded, so no exhaustive implications are predicted; in particular, (53) is not predicted to implicate that all of the linguists attended.

This prediction is not entirely satisfactory, since it is possible to hear B's answer as implicating that all of the linguists *didn't* attend the panel. However, one way to get around this problem in Katzir's framework is to suppose that formal alternatives based on contextually salient linguistic material can be optionally pruned, but equally or less complex alternatives are always active.

The argument depends on the assumption that alternatives are calculated via pointwise replacement of focused content (Rooth, 1985). If this is assumed, then it is indeed a puzzle why *all of the linguists* can serve as an alternative to *some of the linguists* but *some but not all of the linguists* cannot. The circumscription approach does not compose alternatives this way, but builds the directional facts directly into the order by quantifying over alternative background extensions. The default order uniformly minimizes the extension of the background property and could not possibly generate the missing inference in (53). What matters is the type of the background property, not the type of the focused expression.

The exchange in (54) appears to present a stronger argument for Fox and Katzir's view.

(54) **A**: What did the robbers do last week?
B: They stole the books, but not the jewelry.

A: What about this week?
B: They stole the books.
$\not\rightsquigarrow$ *they stole the jewelry*

Fox and Katzir use examples like (54) to support the same conclusions they drew from (53). If the proposition that the robbers stole the books but not the jewelry were used as an alternative to B's answer in (54), exhaustivity would exclude this proposition and license the conclusion that the robbers didn't steal "the books but not the jewelry" (hence they stole the jewelry), but this implication is unattested in the context given. Formal alternatives again come to the rescue: the verb phrases *stole the jewelry* and *stole the books but not the jewelry* are both licit formal alternatives, and since Fox and Katzir's contextual restriction constraint prohibits pruning one without the other, they cancel out.

If A's questions in (54) have abstracts over the domain of properties, and if the judgments in (54) are invariant across contexts, this is a challenging example. Without Fox and Katzir's closure constraint on domain restriction, it should be possible to resolve D to a set that only contains the property of stealing the jewelry; it should also be possible to resolve D to a set that only contains the property of stealing the books but not the jewelry. Minimizing the question-property's intersection with D should then license either conclusion (that the robbers did or didn't steal the jewelry) depending on which choice of D is adopted in context. The type of the background property cannot be counted on to break symmetry in this case.

However, the type of the background property in this example is not obvious. Fox and Katzir assume that because the focused content in B's answer is expressed by a tensed verb phrase, the question abstract must denote a property of individuals. That is one possibility, but A's questions in this example call for answers that denote stage-level properties and there is some debate on whether stage-level property terms are actually type et (Carlson, 1977; Kratzer, 1995). In an event semantic framework (Davidson, 1967; Parsons, 1990) it would be equally possible to analyze the question abstract as denoting a property of events or a property of relations between individuals and events. From this perspective a question like *What did the robbers do?* might be more accurately paraphrased along the lines of *What events have the robbers as agents?* rather than *What properties do the robbers have?*.

Whether symmetry is breakable semantically in an event semantic framework would then depend on what the domain of events is taken to be like. In particular, it would depend on whether "negative events" are admitted and what logical relationship these have to "positive" events. The proper treatment of negation in event semantics is an outstanding puzzle (Bernard and Champollion, 2024), and the empirical predictions here might depend on nontrivial theoretical choices yet to be made.

Since I am staying agnostic about type in (54), a more conclusive empirical testing ground would come from questions whose abstracts uncontroversially denote properties of properties of individuals. Questions whose short answers denote individual-level properties are a good candidate. I suggest that a question like *What's Agnes like?* has as its abstract the generalized quantifier over Agnes, a set of properties that contain Agnes. I will call these SECOND-ORDER QUESTIONS, in contrast to FIRST-ORDER QUESTIONS whose abstracts are over individuals. Exhaustive interpretation of any answer to a second-order question will have to resolve D to an uncomplemented set.

Symmetry breaking in either direction should be possible with such questions. This prediction is borne out: the exhaustive inferences in (55) and (56) are equally possible.

(55) *Context: B has a cat named Gertrude.*
A: I need a cat who's small and friendly for this scene in a film I'm directing. What's Gertrude like? $\lambda w \lambda P_{s\langle et\rangle}.P_w(g)$
B: She's small. $\lambda Q_{\langle s\langle et\rangle\rangle t}.Q(\lambda w \lambda x_e.\textbf{small}_w(x))$
$\leadsto$ *she's not friendly*

(56) **A**: I need a cat who's small and not friendly for this scene in a film I'm directing. What's Gertrude like? $\lambda w \lambda P_{s\langle et\rangle}.P_w(g)$
B: She's small. $\lambda Q_{\langle s\langle et\rangle\rangle t}.Q(\lambda w \lambda x_e.\textbf{small}_w(x))$
$\leadsto$ *she's friendly*

In (55) and (56), the background term denotes a set of first-order properties of individuals that contain Gertrude, and the focused term denotes a set of generalized quantifiers that contain the property of being small. The default order minimizes the extension of the background term: the generalized quantifier over Gertrude contains as few relevant properties as possible. B's answer entails that this quantifier contains the property of being small and exhaustivity dictates that it contains no other relevant properties. In (55), the only other relevant property is that of being friendly and so exhaustivity licenses the conclusion that Gertrude is not friendly. In (56), the only other relevant property is that of being not friendly and so exhaustivity licenses the conclusion that Gertrude is friendly.

There is no way for Fox and Katzir's theory to predict the equal felicity of the exhaustive inferences in (55) and (56). According to their closure condition on domain restriction the proposition that Gertrude is friendly cannot be relevant unless the proposition that Gertrude is not friendly is also relevant and vice versa. To explain why the former proposition but not the latter is excluded in (55), the structural theorist would have to claim that only the former is a formal alternative.

This claim is consistent with the Katzir, 2007 algorithm because the sentence *She's friendly* is as complex as B's answer (*She's small*), whereas the sentence *She's not friendly* is more complex. Yet precisely the opposite result is observed in (56), where only the more complex alternative is excluded. The structural theorist could claim that because *small and not friendly* is made salient in the linguistic context, the sentence *She's not friendly* counts as a formal alternative in (56) but not (55). This move is still not enough to generate the attested inference, since Fox and Katzir's closure condition requires the symmetric alternative (that she's friendly) to appear in the domain of quantification.

The equal felicity of (55) and (56) is expected on the approach to symmetry breaking I have advocated here, which says that context can break symmetry whenever the question-property has a higher type.

Since the work of Katzir, 2007 and Fox and Katzir, 2011, the literature has unearthed a number of cases where context does appear to break symmetry in directions unexpected for the structural theorist (Breheny et al., 2018; Haslinger and Schmitt, 2025; Schwarz and Wagner, 2024; Trinh, 2018; Trinh and Haida, 2015). In the face of this dilemma, the usual strategy has been to adopt auxiliary constraints on formal alternatives; the possibility that semantic type might play a role has not been considered. A typical case adapted from Trinh and Haida, 2015 is given in (57) and (58).

(57) **A**: I'm looking for a man who's tall and bald.
What's Godfrey like? $\lambda w \lambda P_{s\langle et\rangle}.P_w(g)$
B: He's tall. $\lambda Q_{\langle s\langle et\rangle\rangle t}.Q(\lambda w \lambda x_e.\textbf{tall}_w(x))$
$\leadsto$ *he's not bald*

(58) **A**: I'm looking for a man who's tall and not bald.
What's Godfrey like? $\lambda w \lambda P_{s\langle et\rangle}.P_w(g)$
B: He's tall. $\lambda Q_{\langle s\langle et\rangle\rangle t}.Q(\lambda w \lambda x_e.\textbf{tall}_w(x))$
$\rightsquigarrow$ *he's bald*

Trinh and Haida take examples like these, which parallel (55) and (56), to motivate their "atomicity" constraint on formal alternatives discussed in the previous chapter. On the circumscription approach, it would take special stipulation to rule out these cases. Their felicity follows from the fact that the question-property is a second-order property of individuals, not a first-order property.

2.4.3 Marked answers to second-order questions

Exhaustive interpretation of answers to questions with second-order abstracts displays exactly the interactions with context we would expect in a system where the default order, which minimizes the extension of the question-property, operates on a contextually determined subset of the question-property's domain. Similar conclusions can be shown for the hybrid order, which simultaneously minimizes and maximizes the question-property's intersection with an answer's $\mathbf{R}^+$ and $\mathbf{R}^-$ sets, respectively. This too should be done relative to a contextually determined subset of the overall domain. Exhaustive interpretation of marked answers to questions with second-order abstracts behaves exactly as in the case of first-order abstracts, modulo semantic type.

Before showing that this claim holds, some decisions must be made about how to represent negative answers to second-order questions. In particular, we need to decide whether B's answers in (59) and (60) have the same meaning.

(59) **A**: I'm looking for a man who's tall, bald, and Midwestern. What's Godfrey like?
B: Not Midwestern.

(60) **A**: I'm looking for a man who's tall, bald, and not Midwestern. What's Godfrey like?
B: Not Midwestern.

Although both answers should denote properties of generalized quantifiers that combine with the question abstract to deliver the proposition that Godfrey isn't Midwestern, there are several ways in which this could be done. Negation could in principle scope above or below the abstract (61).

(61) a. $\lambda Q_{\langle s\langle et\rangle\rangle t}.Q(\lambda w \lambda x_e.\neg\textbf{midwest}_w(x))$
b. $\lambda Q_{\langle s\langle et\rangle\rangle t}.\neg Q(\lambda w \lambda x_e.\textbf{midwest}_w(x))$

Because the abstract denotes a Montagovian individual and all individuals are either Midwestern or not, the choice between (61-a) and (61-b) will not affect the truth-conditions of the answer in this case. These lambda terms are not in general equivalent: if there are Midwestern and non-Midwestern linguists then *some linguists* will satisfy (61-a) but not (61-b), for instance. The equivalent for (61-a) did not exist in the case of first-order questions because there are no negative individuals, so negation had to scope over the abstract, but second-order questions allow this possibility.

To maintain consistency with the analysis of first-order question/answer exchanges I have proposed, the scope of negation should ensure that an answer's quantificational meaning is generated by a member of D; i.e., it should reflect whether the answer is pragmatically an affirmation or a denial. When the answer affirms that the abstract applies to some ("negative") member of D, negation should scope low. When the answer denies that the abstract applies to some member of D, negation should scope high. Accordingly I will analyze B's answer in (59) as denoting (61-a) and B's answer in (60) as denoting (61-b).

These decisions about scope have the following consequences. B's answer in (59) is a default answer, not really a negative answer; it is positively restricted by the set containing the property $\neg$**midwest** and negatively restricted by the empty set, as shown in (62).

(62) **A**: I'm looking for a man who's tall, bald, and not Midwestern.
What's Godfrey like? $\lambda w \lambda P_{s\langle et\rangle}.P_w(g)$
B: Not Midwestern. $\lambda Q_{\langle s\langle et\rangle\rangle t}.Q(\lambda w \lambda x_e.\neg\mathbf{midwest}_w(x))$
$\rightsquigarrow$ *he isn't tall or bald*
$\mathbf{R}^+(T) = \{\neg\mathbf{midwest}\}, \mathbf{R}^-(T) = \varnothing, D = \{\mathbf{tall}, \mathbf{bald}, \neg\mathbf{midwest}\}$

Exhaustivity uses the default order and licenses the conclusion that Godfrey has no other relevant properties, in this case **tall** and **bald**. This is the right result for (62).

By contrast, (60) is a true negative answer. This example shows that negative answers to second-order questions exhibit restricted extent, just like negative answers to first-order questions. In (63), the $\mathbf{R}^+$ set for *not Midwestern* (with high scope) is empty; the $\mathbf{R}^-$ set is the singleton containing the property **midwest**.

(63) **A**: I'm looking for a man who's tall, bald, and Midwestern.
What's Godfrey like? $\lambda w \lambda P_{s\langle et\rangle}.P_w(g)$
B: Not Midwestern. $\lambda Q_{\langle s\langle et\rangle\rangle t}.\neg Q(\lambda w \lambda x_e.\mathbf{midwest}_w(x))$
$\not\rightsquigarrow$ *conclude nothing about whether he's tall or bald*
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \{\mathbf{midwest}\}, D = \{\mathbf{tall}, \mathbf{bald}, \mathbf{midwest}\}$

Applied to this answer, the hybrid order predicts no further conclusions beyond the literal meaning. This prediction is correct as B's answer implicates nothing about whether Godfrey is tall or bald.

Exhaustive interpretation of negative answers to second-order questions also exhibits positive direction (64), as predicted by the hybrid order. Whether an exhaustive implication is "positive" or "negative" depends on the question-property and how D is resolved in the context of utterance; as before, implications that some member of D appears in the question-property's extension count as positive, and implications that some member of D does not appear in the question-property's extension count as negative.

(64) **A**: I'm looking for a man who's tall, bald, and Midwestern.
What's Godfrey like? $\lambda w \lambda P_{s\langle et\rangle}.P_w(g)$
B: He isn't all of those things.
$\lambda Q_{\langle s\langle et\rangle\rangle t}.\neg Q(\lambda w \lambda x_e.\mathbf{tall}_w(x)) \vee \neg Q(\lambda w \lambda x_e.\mathbf{bald}_w(x)) \vee \neg Q(\lambda w \lambda x_e.\mathbf{midwest}_w(x))$
$\rightsquigarrow$ *he's tall or bald or Midwestern*
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = D = \{\mathbf{tall}, \mathbf{bald}, \mathbf{midwest}\}$

B's answer in (64) denotes a property of generalized quantifiers that don't contain all of the properties **tall**, **bald**, and **midwest**. The $\mathbf{R}^+$ set for this answer is empty; the $\mathbf{R}^-$ set is identical to D. The hybrid order maximizes the question-property's intersection with $\mathbf{R}^-$ and yields the conclusion that Godfrey does have some of the relevant properties.

The hybrid order correctly predicts positive direction and restricted extent in (65).

(65) **A**: I'm looking for a man who's tall, bald, and Midwestern.
What's Godfrey like? $\lambda w \lambda P_{s\langle et\rangle}.P_w(g)$
B: He isn't both tall and Midwestern.
$\lambda Q_{\langle s\langle et\rangle\rangle t}.\neg Q(\lambda w \lambda x_e.\mathbf{tall}_w(x)) \vee \neg Q(\lambda w \lambda x_e.\mathbf{midwest}_w(x))$
$\rightsquigarrow$ *he's tall or Midwestern*
$\not\rightsquigarrow$ *conclude nothing about whether he's bald*
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \{\mathbf{tall}, \mathbf{midwest}\}, D = \{\mathbf{tall}, \mathbf{bald}, \mathbf{midwest}\}$

B's answer in (65) denotes a property of generalized quantifiers that don't contain both **tall** and **midwest**. The $\mathbf{R}^-$ set for this answer contains the properties **tall** and **midwest**. Applied to this answer, the hybrid order maximizes the question-property's intersection with $\mathbf{R}^-$ and yields the conclusion that Godfrey is either tall or bald. Since $\mathbf{R}^-$ does not contain **bald**, no conclusions are drawn about whether Godfrey is bald.

These examples show that the direction and extent properties of marked answers to second-order questions parallel those of marked answers to first-order questions.

However, there are further possible implications that arise with second-order questions whose analogues do not exist with first-order questions. Because the domain of individuals is not complemented, any particular exhaustive implication relative to a first-order question is necessarily always either positive or negative regardless of how D is resolved. Relative to the question *Who attended the panel?*, the implication that Beatrice didn't attend is always negative. In the case of second-order questions, however, different choices of D can affect the (superficial) polarity of exhaustive implications.

Changing *Midwestern* to *not Midwestern* in (66)-(68) predictably affects the content of exhaustive implications, but does not change the direction or extent properties of marked answers relative to the new domain. B's answer in (66) still has restricted extent, B's answer in (67) still has positive direction, and B's answer in (68) still has positive direction and restricted extent. These observations support the context-sensitive, type-theoretic approach to symmetry breaking I have advocated here.

(66) **A**: I'm looking for a man who's tall, bald, and not Midwestern.
What's Godfrey like? $\lambda w \lambda P_{s\langle et\rangle}.P_w(g)$
B: Midwestern. $\lambda Q_{\langle s\langle et\rangle\rangle t}.\neg Q(\lambda w \lambda x_e.\neg\mathbf{midwest}_w(x))$
$\not\rightsquigarrow$ *conclude nothing about whether he's tall or bald*
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \{\neg\mathbf{midwest}\}, D = \{\mathbf{tall}, \mathbf{bald}, \neg\mathbf{midwest}\}$

(67) **A**: I'm looking for a man who's tall, bald, and not Midwestern.
What's Godfrey like? $\lambda w \lambda P_{s\langle et\rangle}.P_w(g)$
B: He isn't all of those things.
$\lambda Q_{\langle s\langle et\rangle\rangle t}.\neg Q(\lambda w \lambda x_e.\mathbf{tall}_w(x)) \vee \neg Q(\lambda w \lambda x_e.\mathbf{bald}_w(x)) \vee \neg Q(\lambda w \lambda x_e.\neg\mathbf{midwest}_w(x))$
$\rightsquigarrow$ *he's tall or bald or not Midwestern*
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = D = \{\mathbf{tall}, \mathbf{bald}, \neg\mathbf{midwest}\}$

(68) **A**: I'm looking for a man who's tall, bald, and not Midwestern.
What's Godfrey like? $\lambda w \lambda P_{s\langle et\rangle}.P_w(g)$
B: He isn't both tall and not Midwestern.
$\lambda Q_{\langle s\langle et\rangle\rangle t}.\neg Q(\lambda w \lambda x_e.\mathbf{tall}_w(x)) \vee \neg Q(\lambda w \lambda x_e.\neg\mathbf{midwest}_w(x))$
$\rightsquigarrow$ *he's tall or not Midwestern*
$\not\rightsquigarrow$ *conclude nothing about whether he's bald*
$\mathbf{R}^+(T) = \varnothing, \mathbf{R}^-(T) = \{\mathbf{tall}, \neg\mathbf{midwest}\}, D = \{\mathbf{tall}, \mathbf{bald}, \neg\mathbf{midwest}\}$

One consequence of my assumption that the scope of negation is determined by the members of D is that the answer *Midwestern* in (66) is treated as involving double negation. This seems somewhat convoluted semantically, but reflects the pragmatic fact that this answer is interpreted as a denial, rejecting the application of the question abstract to a property that happens to be "negative". Because the property **midwest** is not in the domain of discourse, B's answer is preferentially interpreted relative to the property ¬**midwest** that is in the domain of discourse.

These examples show that exhaustive interpretation of marked answers to second-order questions exhibits exactly the same variety of context-sensitivity that default answers to second-order questions do, driven by the greater expressive possibilities that arise with the switch to higher types.

This variety of context-sensitivity is only made possible by the fact that abstracts with higher types denote functions over complemented domains. Whenever the question abstract has a lower type, symmetry should be unbreakable. Example (53), adapted from Fox and Katzir, 2011, already showed that this prediction is correct. A similar example is given in (69).

(69) **A**: I hope Agnes but not Beatrice attended the AI panel. Agnes always has interesting things to say, but I really don't care for Beatrice's Panglossian techno-hype.
B: I was in the audience.
A: Oh really? Who attended? $\lambda w \lambda x_e.\textbf{attended}_w(x)$
B: Agnes. $\lambda P_{et}.P(a)$
$\rightsquigarrow$ *Beatrice didn't attend*
$\not\rightsquigarrow$ *Beatrice attended*

There is no way to hear B's answer in (69) as implicating that Beatrice attended. This inference remains impossible despite the context, which makes the potential alternative *Agnes but not Beatrice* highly salient and pragmatically desirable; nevertheless, exhaustivity cannot exclude the proposition that "Agnes but not Beatrice" attended. Exhaustivity can only license the conclusion that Beatrice didn't attend. This result stands in contrast to examples like (55) and (56), which did allow exhaustive implications in the "wrong" direction once a sufficiently explicit and supportive context was given.

The theory of exhaustivity, alternatives, and symmetry breaking I have advocated directly captures this contrast: it follows from the fact that the question abstract in (55) and (56) denotes a property of properties of individuals, whereas the question abstract in (69) denotes a property of individuals.

The question of whether context can break symmetry has been extensively debated in the literature (Bergen, Levy, and Goodman, 2016; Breheny et al., 2018; Buccola, Dautriche, and Chemla, 2018; Buccola, Križ, and Chemla, 2022; Feinmann, 2023, 2025; Fox and Katzir, 2011; Haslinger and Schmitt, 2025; Hirsch and Schwarz, 2025; Katzir, 2007; Schwarz and Wagner, 2024; Trinh, 2018; Trinh and Haida, 2015). This question is often taken to have an absolute, definitive answer, yet examples supporting both conclusions have proliferated. For the structural theorist, neither answer is satisfactory. If the answer is yes, then formal alternatives do no explanatory work. If the answer is no, then exactly the right set of additional constraints on formal alternatives must be devised to yield flexibility in only the cases where it is observed and nowhere else.

I have argued that an empirically successful and explanatory answer to this question can be given if we recognize that exhaustivity necessarily operates on focus/background structures. The key factor determining whether context can break symmetry is the type of the background term. Higher-order background properties over complemented domains allow for contextual symmetry breaking; first-order background properties over uncomplemented domains do not. This difference does not have to be stipulated, but follows from independent model-theoretic properties of the objects we are quantifying over, together with the nature of exhaustivity.

There are many varieties of *wh*-questions whose abstracts plausibly denote functions over different sorts of objects besides individuals and properties of individuals. In particular, answers to degree questions and modal questions allow exhaustive inferences too, and I have suggested that questions targeting stage-level properties may require an analysis in terms of events. A complete solution to the symmetry problem will need to address such questions, as well as multiple *wh*-questions. The theory I have developed in this chapter is really more like an existence proof that a semantic solution is possible.

Although I leave this investigation to future work, the logic of the analysis should make clear what it would take to extend the theory to additional cases and what the predictions would be. Questions whose abstracts denote functions over uncomplemented domains should allow exhaustive implications in only one direction, and questions whose abstracts denote functions over complemented domains should allow contextual symmetry breaking in either direction. In any new empirical setting, the proposal could be straightforwardly confirmed or falsified by checking whether this is so.

2.4.4 Summary

The analysis of answers to questions with higher types raises complications that do not arise with first-order questions. Because the domain of properties is complemented, it is often the case that either some property or its complement must appear in a GQ term's denotation. Some domain restriction is therefore necessary if exhaustivity is to exclude a property from the denotation of a second-order abstract. I argued that domain restriction is also sufficient to explain the directional properties of answers to second-order questions.

This perspective on symmetry and alternatives, which has been most clearly articulated in the work of Szabolcsi, 1983, assigns equal roles to meaning and context. Default and marked answers to second-order questions are interpreted using the same orders as default and marked answers to first-order questions, but unlike answers to first-order questions they allow context to break symmetry between a property and its complement. Formal alternatives and Boolean closure conditions over quantificational domains (Fox and Katzir, 2011) are not necessary to achieve this result. The pattern follows automatically from the semantic type of the background property.

I take this claim to be the null hypothesis about how symmetry interacts with type. It is always pleasing when the null hypothesis is correct.

2.5 Conclusion

This chapter developed a comprehensive semantic analysis of how exhaustivity interacts with the polarity of answers, building on previous circumscription accounts (Groenendijk and Stokhof, 1984; Schulz and van Rooij, 2006; Spector, 2007; von Stechow and Zimmermann, 1984). Answers were sorted into default and marked classes based on their extent properties. Default answers are interpreted using the default order, which minimizes the question-property's overall extension. Marked answers are interpreted using the hybrid order, which minimizes the question-property's intersection with an answer's positive restriction and maximizes its intersection with an answer's negative restriction.

The hybrid order, and the notion of positive and negative quantificational restriction it uses, were developed in detail. Circumscribing the question-property relative to the hybrid order makes correct predictions about the directional properties of a wide range of marked answers. I also showed that the default and hybrid orders are sufficient to capture the directional properties of answers to second-order questions once domain restriction is taken into account, following in the footsteps of Szabolcsi, 1983.

The system I have presented precisely matches the empirical coverage of Spector, 2007. This is a good result, since Spector's proposal correctly captures the direction and extent properties of positive, negative, and mixed answers. Unlike Spector, I do not translate answers into propositional logic and associate their translations with different alternative logical formulae; instead I have provided a direct model-theoretic characterization of exhaustive interpretation in terms of quantification over background property extensions. Negative and mixed answers differ from positive/default answers in using a different order over background extensions.

Several questions remain. First, I have not indicated what properties of an answer make it default or marked; so far this categorization is a stipulation. Second, an explanatory account should go beyond identifying some model-theoretic property of generalized quantifier terms that happens to correctly distinguish between these two classes of answers. It should ground the distinction in independent pragmatic principles. These are the goals of chapter 3.

Chapter 3

Order selection

This chapter develops a theory of order selection. Although the default order and the hybrid order make correct predictions about exhaustive interpretation for a wide range of answers, to avail ourselves of this result we need a systematic way to determine which order a particular answer will use. Answers that use the default vs. the hybrid order are distinguished in terms of a model-theoretic property I call SKEW: a POSITIVELY SKEWED quantifier's total restriction and positive restriction are identical. This property holds of all answers whose exhaustive interpretation uses the default order, and only these. The role of skew in order selection is derived from how the pragmatic preference for a less marked order interacts with a general constraint on information preservation that requires an answer's restriction to be formally recoverable from its exhaustive interpretation. The markedness reduction principle selects the least marked order compatible with the information preservation requirement, yielding the observed order selection pattern. Finally, the chapter examines the rare, but attested instances where exhaustive interpretation of decreasing answers does have unrestricted extent, as observed by von Stechow and Zimmermann, 1984, and shows how the overall framework can be modified to account for these cases.

3.1 Skew

Chapter 2 proposed two orders over possible worlds, the DEFAULT order and the HYBRID order, that were shown to make correct predictions about the direction and extent of exhaustive interpretation. I sorted answers into two classes, default and marked answers. Default answers use the default order, a simple inclusion-induced lattice over background property extensions that enforces negative direction and unrestricted extent. Marked answers use the hybrid order, which establishes a more complex algebraic structure by integrating two sources of ordering information: a subset-induced order $<^+$ over the background property's intersection with an answer's positive restriction, and a superset-induced order $<^-$ over the background property's intersection with an answer's negative restriction.

The hybrid order uniformly enforces restricted extent, while the directional properties of the hybrid order vary depending on the answer: exhaustivity has positive direction with decreasing quantifiers and mixed direction with nonmonotonic quantifiers. If the hybrid order were used with an increasing quantifier, exhaustivity would have negative direction and restricted extent. However, increasing quantifiers are default answers, so they use the default order.

What is the basis of the default vs. marked distinction? Familiar ways of formalizing quantifier polarity, like monotonicity, do not make the cut in the right place: although all increasing quantifiers are default and all decreasing quantifiers are marked, nonmonotonic quantifiers vary in this regard. Some nonmonotonic quantifiers (*exactly 3 linguists*) have negative direction and unrestricted extent, while others (*most linguists but not a lot of computer scientists*) have mixed direction and restricted extent. The default vs. marked distinction must then interact with monotonicity without reducing to it. Making this distinction precise calls for some model-theoretic property that classes nonmonotonic quantifiers of the first kind with increasing quantifiers, and nonmonotonic quantifiers of the second kind with decreasing quantifiers.

This section develops a novel formalization of quantifier polarity that I call SKEW. Intuitively, skew describes the relationship between a quantifier's total restriction and its positive and negative subsets. A skewed quantifier is one whose restriction is identical to its positive or negative restriction. All increasing quantifiers are positively skewed, but some nonmonotonic quantifiers are too. I claim that default answers denote positively skewed quantifiers.

To my knowledge, the concept of skew has not been developed before. Because this formal property is based on the restriction operators from chapter 2, which were new, previous GQ literature has not classified quantifiers in this way. Skew is a very general notion, however, and if the proposal is on the right track it could potentially have further linguistic applications beyond the analysis of exhaustivity.

3.1.1 Quantifier polarity

Chapter 2 developed two novel positive and negative RESTRICTION OPERATORS $\mathbf{R}^+$ and $\mathbf{R}^-$, which are listed in (1) along with the **R** operator that delivers a quantifier's total restriction (von Stechow and Zimmermann, 1984).

(1) **Restriction operators.**

a. $\mathbf{R} = \lambda T_{\langle\tau t\rangle t}.\bigcap\{S|\forall P : T(P) \leftrightarrow T(P \cap S)\}$
b. $\mathbf{R}^+ = \lambda T_{\langle\tau t\rangle t}\lambda x_\tau.\exists S : \neg T(S) \wedge T(S \cup \{x\})$
c. $\mathbf{R}^- = \lambda T_{\langle\tau t\rangle t}\lambda x_\tau.\exists S : \neg T(S) \wedge T(S - \{x\})$

R takes a generalized quantifier T and delivers the intersection of the sets that T lives on. $\mathbf{R}^+$ takes a generalized quantifier T and delivers its POSITIVE RESTRICTION, a set of objects that can be added to a set that T does not apply to and create a set that T does apply to. $\mathbf{R}^-$ takes a generalized quantifier T and delivers its NEGATIVE RESTRICTION, a set of objects that can be subtracted from a set that T does not apply to and create a set that T does apply to. In the appendix, I prove that the union of any quantifier's positive and negative restriction is always equal to its total restriction, that all increasing quantifiers have empty negative restrictions, and that all decreasing quantifiers have empty positive restrictions.

The restriction operators can be used to define SKEW, a formal property of generalized quantifiers whose total restriction is identical to their positive or negative restriction.

(2) **Skew.**
A generalized quantifier T is *positively skewed* iff $\mathbf{R}(T) = \mathbf{R}^+(T)$.
A generalized quantifier T is *negatively skewed* iff $\mathbf{R}(T) = \mathbf{R}^-(T)$.

Skew is a simple notion. A positively skewed quantifier is one whose total restriction is identical to its positive restriction. A negatively skewed quantifier is one whose total restriction is identical to its negative restriction. Because the union of a quantifier's positive and negative restrictions always adds up to its total restriction, and it is possible for a quantifier's positive or negative restriction to be empty, many quantifiers are skewed in one direction or the other.

The useful facts in (3) regarding skew and monotonicity are proven in the appendix.

(3)
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 systematically interacts with monotonicity, but is distinct from it. Recall that increasing quantifiers have empty negative restrictions and decreasing quantifiers have empty positive restrictions. This fact guarantees that increasing quantifiers are positively skewed and decreasing quantifiers are negatively skewed.

Positive and negative skew are not mutually exclusive: some quantifiers are skewed in both directions. Trivially, this class includes quantifiers that are both increasing and decreasing in virtue of (3-a) and (3-b), but also includes nonmonotonic quantifiers with nonempty, identical positive and negative restrictions, like *exactly 3 linguists* and *between 5 and 10 computer scientists*. The remaining nonmonotonic quantifiers are those with nonempty, distinct positive and negative restrictions, like *most linguists but not a lot of computer scientists*. These are the only UNSKEWED quantifiers.

Skew thus yields a four-way typology of generalized quantifiers, which can be positively skewed, negatively skewed, both, or neither. This new formalization of quantifier polarity rather recalls the monotonicity-induced typology of quantifiers, since quantifiers can similarly be increasing, decreasing, both increasing and decreasing, or nonmonotonic. The main difference is that simultaneously increasing and decreasing quantifiers (*zero or more linguists*, *some linguists and no linguists*) are trivial or vacuous, whereas positively and negatively skewed quantifiers can be nontrivial; skew crosscuts the nonmonotonic quantifiers in a substantive way. It is this feature of skew that captures the distinction between default and marked answers.

The increasing quantifiers in (4) are positively skewed: their total restrictions and positive restrictions both contain the linguists, and their negative restrictions are empty.

(4) $\mathbf{R}(T) = \mathbf{R}^+(T) = \mathbf{linguists}, \mathbf{R}^-(T) = \varnothing$ *positive skew*
a. some of the linguists $\lambda P.\mathbf{linguists} \cap P$
b. every linguist $\lambda P.\mathbf{linguists} \subseteq P$
c. most linguists $\lambda P.|\mathbf{linguists} \cap P| > |\mathbf{linguists} \cap \overline{P}|$

Similarly, the increasing quantifiers in (5) are positively skewed: their total restrictions and positive restrictions both contain Agnes and Beatrice, and their negative restrictions are empty.

(5) $\mathbf{R}(T) = \mathbf{R}^+(T) = \{a, b\}, \mathbf{R}^-(T) = \varnothing$ *positive skew*
a. Agnes or Beatrice $\lambda P.P(a) \vee P(b)$
b. Agnes and Beatrice $\lambda P.P(a) \wedge P(b)$

The decreasing quantifiers in (6) are negatively skewed: their total restrictions and negative restrictions both contain the linguists, and their positive restrictions are empty.

(6) $\mathbf{R}(T) = \mathbf{R}^-(T) = \textbf{linguists}, \mathbf{R}^+(T) = \varnothing$ *negative skew*

a. no linguist $\lambda P.\textbf{linguists} \not\cap P$

b. not every linguist $\lambda P.\textbf{linguists} \not\subseteq P$

c. not a lot of linguists $\lambda P.|\textbf{linguists} \cap P| < n$

Similarly, the decreasing quantifiers in (7) are negatively skewed: their total restrictions and negative restrictions both contain Agnes and Beatrice, and their positive restrictions are empty.

(7) $\mathbf{R}(T) = \mathbf{R}^-(T) = \{a, b\}, \mathbf{R}^+(T) = \varnothing$ *negative skew*

a. neither Agnes nor Beatrice $\lambda P.\neg P(a) \wedge \neg P(b)$

b. not both Agnes and Beatrice $\lambda P.\neg P(a) \vee \neg P(b)$

The nonmonotonic quantifiers in (8) are both positively and negatively skewed. Their positive restrictions, negative restrictions, and total restrictions all contain the linguists.

(8) $\mathbf{R}(T) = \mathbf{R}^+(T) = \mathbf{R}^-(T) = \textbf{linguists}$ *positive and negative skew*

a. exactly 3 linguists $\lambda P.|\textbf{linguists} \cap P| = 3$

b. between 5 and 10 linguists $\lambda P.5 \leq |\textbf{linguists} \cap P| \leq 10$

c. some but not all of the linguists $\lambda P.\textbf{linguists} \cap P \wedge \textbf{linguists} \not\subseteq P$

However, the nonmonotonic *Agnes but not Beatrice* is an unskewed quantifier: its positive and negative restrictions are nonempty and disjoint. This quantifier is positively restricted by the set containing Agnes, negatively restricted by the set containing Beatrice, and restricted overall by the set containing Agnes and Beatrice.

(9) $\{a\} = \mathbf{R}^+(T) \neq \mathbf{R}(T) \neq \mathbf{R}^-(T) = \{b\}$ *unskewed*

Agnes but not Beatrice $\lambda P.P(a) \wedge \neg P(b)$

The nonmonotonic *most linguists but not a lot of computer scientists* is similarly unskewed. This quantifier is positively restricted by the set containing the linguists, negatively restricted by the set containing the computer scientists, and restricted overall by the set containing the linguists and the computer scientists.

(10) $\textbf{linguists} = \mathbf{R}^+(T) \neq \mathbf{R}(T) \neq \mathbf{R}^-(T) = \textbf{computer scientists}$ *unskewed*

most linguists but not a lot of computer scientists

$\lambda P.|\textbf{linguists} \cap P| > |\textbf{linguists} \cap \overline{P}| \wedge |\textbf{computer scientists} \cap P| < n$

Conditional answers are also unskewed: an answer like *if Agnes attended, Beatrice attended* is positively restricted by the set containing Beatrice, negatively restricted by the set containing Agnes, and restricted overall by the set containing Agnes and Beatrice.

(11) $\{b\} = \mathbf{R}^+(T) \neq \mathbf{R}(T) \neq \mathbf{R}^-(T) = \{a\}$ *unskewed*

if Agnes, then Beatrice $\lambda P.P(a) \rightarrow P(b)$

These results concerning quantifier restrictions are mostly familiar from chapter 2, where the restriction operators were introduced. However, they are worth reviewing in order to see how the restriction operators, which are independently necessary to define the hybrid order, also make it possible to define skew, which I claim grounds the distinction between default and marked answers.

3.1.2 Default answers and positive skew

The notion of skew provides a formal basis for predicting which order an answer will use. The goal was to find a model-theoretic property that classes increasing quantifiers together with nonmonotonic quantifiers like *exactly 3 linguists*, and classes decreasing quantifiers together with nonmonotonic quantifiers like *most linguists but not a lot of computer scientists*.

Skew distinguishes nonmonotonic quantifiers in exactly the right way. Nonmonotonic quantifiers are either both positively and negatively skewed, or unskewed. This is the distinction we need to predict whether a nonmonotonic quantifier is default or marked. The class of default answers consists of increasing quantifiers and nonmonotonic quantifiers that are skewed in both directions. The class of marked answers consists of decreasing quantifiers and unskewed nonmonotonic quantifiers. I propose the generalization in (12): positive skew is both necessary and sufficient for use of the default order.

(12) **Order selection.**
If T is positively skewed, exhaustive interpretation of T uses the default order; otherwise exhaustive interpretation of T uses the hybrid order.

Table 3.1 categorizes a range of different answers by skew, direction, and extent. It is the same as Table 2.1 in Chapter 2, with the monotonicity column replaced by skew. Unlike monotonicity, skew correlates directly with extent. The pattern reflects the generalization in (12): positively skewed answers have negative direction and unrestricted extent; all other answers have positive or mixed direction and restricted extent.

Answer	Skew	Direction	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	-	n/a	restricted
No linguists	-	n/a	restricted
Not all of the linguists	-	+	restricted
Not both Agnes and Beatrice	-	+	restricted
Not a lot of linguists	-	+	restricted
Neither Agnes nor Beatrice	-	n/a	restricted
Agnes but not Beatrice	n/a	n/a	restricted
Most linguists but not a lot of computer scientists	n/a	±	restricted
If Agnes, then Beatrice	n/a	±	restricted

Table 3.1: Skew x direction x extent.

It is worth reviewing some examples. B's answer in (13) is positively (and negatively) skewed. The default order must be selected to license the conclusion that no computer scientists attended.

(13) **A**: Who attended the last panel?
B: Exactly three linguists. *positive and negative skew*
$\rightsquigarrow$ *no computer scientists*

Similarly, B's answer in (14) is positively (and negatively) skewed. The default order must likewise be selected to license the conclusion that no computer scientists attended.

(14) **A**: Who attended the last panel?
B: Some but not all of the linguists. *positive and negative skew*
$\rightsquigarrow$ *no computer scientists*

B's answer in (15) is unskewed: its positive and negative restrictions consist of the syntacticians and the phonologists, respectively, and are disjoint. Exhaustive interpretation licenses the conclusions that not every linguist attended and that some phonologists did, but nothing is concluded about other linguists or computer scientists. This result calls for the hybrid order.

(15) **A**: Who attended the last panel?
B: Most syntacticians, but not a lot of phonologists. *unskewed*
$\rightsquigarrow$ *not all syntacticians, some phonologists*
$\not\rightsquigarrow$ *conclude nothing about computer scientists or other linguists*

B's answer in (16) is similarly unskewed. This answer licenses a conditional perfection inference, but does not allow conclusions about individuals other than Agnes and Beatrice. This result likewise calls for the hybrid order.

(16) **A**: Who attended the last panel?
B: If Agnes attended, then Beatrice attended. *unskewed*
$\rightsquigarrow$ *if Agnes didn't attend, then Beatrice didn't attend*
$\not\rightsquigarrow$ *conclude nothing about anyone else*

The conclusion to be drawn from this discussion is that the distinction between default and marked answers is not an arbitrary stipulation, but can be systematically identified with the model-theoretic property of positive skew.

Because skew is independent of exhaustivity, one might wonder whether this model-theoretic property has any linguistic relevance elsewhere. One possible connection involves a class of questions that exhibit interesting interactions with quantifiers. Xiang, 2021a, building on Spector, 2008, has argued for a class of *third-order* questions whose abstracts denote sets of generalized quantifiers. The evidence comes from scopal interactions between answer terms and quantifiers contained in a question, as in (17).

(17) **A**: Which classes do you have to take?
B: Syntax or phonology. $\vee > \Box, \Box > \vee$

There are two ways of reading B's answer in (17). The "first-order" reading says that B has to take syntax, or B has to take phonology. This is the predicted reading if the abstract for A's question denotes a first-order property, the property of being a class B has to take, and B's answer denotes a generalized quantifier that takes this property as its argument. Given this function/argument structure, the disjunction in B's answer will necessarily take scope over the modal in A's question.

The "higher-order" reading says that B can fulfill the requirement by taking either syntax or phonology. To generate this reading, the modal in A's question must take scope over the disjunction. Because this reading cannot be generated on the assumption that A's question denotes a first-order property, Spector and Xiang propose that the question can also optionally denote a third-order property of generalized quantifiers ("which quantifiers are such that they must apply to the property of being a class you take?"), which takes B's quantificational answer as its argument. The predicted truth-conditions are that B must take syntax or phonology, but neither disjunct is itself necessary.

Xiang observes that there are polarity-related restrictions on the quantifiers that can interact scopally with questions in this way. For instance, the negation in B's answer in (18) must take scope over the modal; (18) can only mean that B doesn't have to take semantics, not that it must be the case that B doesn't take semantics (i.e., B isn't allowed to take semantics).

(18) **A**: Which classes do you have to take?
B: Not semantics. $\neg > \Box, *\Box > \neg$

The key point in this discussion that indicates a role for skew is Xiang's observation that some nonmonotonic answers can participate in this sort of scopal interaction and others cannot. B's doubly skewed answer in (19) allows both readings: the first-order reading (there are exactly three specific seminars that B must take) and the higher-order reading (B must take three seminars and may not take more) are both available.

(19) **A**: Which linguistics classes do you have to take?
B: Exactly three graduate seminars. $3 > \Box, \Box > 3$

Readers who have difficulty accessing the higher-order reading may find it helpful to imagine a context in which B is enrolled in a yearlong program with strict upper bounds on courseload, so that taking any additional seminars would mean not taking other required classes. In such a context, the prohibition against additional seminars is plausible and pragmatically motivated, so the higher-order reading is accordingly more natural.

However, B's unskewed answer in (20) only allows the first-order reading; there is no reading of (20) that prohibits B from taking phonology.

(20) **A**: Which linguistics classes do you have to take?
B: Syntax, but not phonology. $\neg > \Box, *\Box > \neg$

Xiang surveys a wide range of quantificational answers and concludes that only a restricted class of answers, which empirically coincide with the positively skewed quantifiers, allow higher-order readings. Xiang claims that this class of quantifiers is unified by a different (and nonequivalent) model-theoretic property, which she calls "homogeneous positivity", and she proposes a type-shifting rule that maps a first-order question abstract to a third-order abstract over quantifiers with this property.

Future work on this class of questions should clarify the relation between Xiang's homogeneous positivity and my notion of positive skew. The role of the polarity restrictions should also be substantively motivated. The restriction to homogeneously positive quantifiers encoded in Xiang's type-shifting rule is somewhat arbitrary, and it is tempting to speculate that this compositional operation could be grounded in the notion of a "default answer" that I have argued aligns with positive skew. Moreover, if A's questions in examples like (17)-(20) indeed allow abstracts that denote third-order properties of individuals, the claims I have made in this thesis about type and symmetry breaking should be tested against these examples.

I do not pursue this line here, but I am encouraged that independent researchers have found the same empirical class of answers to play a significant role in the analysis of related linguistic phenomena.

3.1.3 Summary

Whether an answer uses the default or hybrid order is not predictable given previously established formalizations of quantifier polarity such as monotonicity. I proposed to distinguish default and marked answers using the novel model-theoretic property of skew. Skewed quantifiers are quantifiers whose restriction is identical to their positive or negative restriction; they can be thought of as quantifiers whose $\mathbf{R}^+$ set or $\mathbf{R}^-$ set (or both) "dominates" the restriction $\mathbf{R}$. All monotonic quantifiers are skewed, but only some non-monotonic quantifiers are.

The class of default answers is unified by the property of positive skew. Marked answers are those that lack positive skew; both unskewed answers and answers that are only negatively skewed fall under this umbrella. Exhaustivity uses the default order with default answers and the hybrid order with marked answers.

With this generalization in hand, the proposal developed throughout this thesis has reached descriptive adequacy: direction and extent are predictable from skew. The goal of the following section is to reach the larger goal of explanatory adequacy by deriving this correlation from independent pragmatic principles.

3.2 Selection constraints

Why should skew matter? This section shows that the role of skew in order selection follows from the interaction of two factors. One is the pressure to reduce markedness, leading to a preference to use the default order whenever possible. The other is a principle of "information preservation" that requires an answer's literal and exhaustive interpretations to share a restriction. Together, these two principles derive the correlation between positive skew and the default order.

The latter constraint, which I call PRESERVE RESTRICTION, places bounds on how exhaustive interpretation changes an answer's quantificational restriction. Taking restriction as the formal analogue to the informal notion of "aboutness", PRESERVE RESTRICTION guarantees that exhaustive answers are still about the same objects as their literal counterparts. The general idea behind this principle is that semantic content should be recoverable from pragmatic meaning.

In order to make formal sense of this idea it is necessary not only to access generalized quantifier meanings for answers, but also to view exhaustivity as an operation that maps one quantifier to another. To the extent that the proposal is successful, then, it provides another argument for a structured account of the question/answer relation.

3.2.1 Preserve restriction

In chapter 2, I suggested that the default order is less marked than the hybrid order. The default order is a simple inclusion-induced lattice over property extensions. The hybrid order is more complex, since it uses two separate orders $<^+$ and $<^-$, tracks them both simultaneously, and combines them in a fairly sophisticated way that integrates different sources of potentially conflicting information. Whereas the default order depends solely on the question-property, the hybrid order also depends on the particular answer that was given, since whether an object appears in an answer's positive or negative restriction (or both, or neither), thereby affecting $<^+$ or $<^-$, varies according to the answer. The default order is also capable of licensing stronger conclusions, since its extent is not limited to an answer's quantificational restriction.

These factors suggest that language users should have a preference to use the default order whenever possible. Exhaustive interpretation with unrestricted extent is preferred over restricted extent. I propose the pragmatic principle REDUCE MARKEDNESS (21), which drives this preference.

(21) **Reduce markedness.**
Use the default order.

REDUCE MARKEDNESS will be given a more general statement in the following section, where I suggest that the dual of the default order is sometimes exceptionally available. This principle will then need to be stated as enforcing a preference to use the least marked order possible. In a setting with only two orders, such a preference is equivalent to a preference for the default one.

REDUCE MARKEDNESS must be a violable constraint, since there are answers that do not use the default order. There must be some other factor that takes priority in such cases. Together with REDUCE MARKEDNESS, this factor should interact with quantificational skew in such a way as to derive the correlation between skew and order selection: only positively skewed quantifiers use the default order. What would go wrong if the default order were used with an answer that is not positively skewed?

In a sense, we already know the answer to this question: the observation that the default order makes wrong empirical predictions for marked answers was what motivated the development of the hybrid order in the first place. If the default order were used with any of the decreasing answers in (22), the predicted exhaustive reading would be that nobody attended the panel, an inference that is not attested.

(22) **A**: Who attended the last panel?
 a. Not Beatrice.
 b. Not all of the linguists.
 $\not\rightsquigarrow$ *nobody attended*

Because the default order selects only the minimal background extensions consistent with the answer, and all decreasing quantifiers contain the empty set, all decreasing answers would implicate the same thing if the default order were used. They would all implicate that the question-property has an empty extension.

If the default order were used with any of the mixed answers in (23), it would minimize the question-property's extension as much as is consistent with the "positive" component of the answer: (23-a) would implicate that only Agnes attended, and (23-b) would implicate that no computer scientists attended.

(23) **A**: Who attended the last panel?
 a. Agnes but not Beatrice.
 $\not\rightsquigarrow$ *nobody but Agnes*

b. Most linguists but not a lot of computer scientists.
$\not\rightsquigarrow$ *not every linguist, no computer scientists*

The problem is that the default order ignores the "negative" component of these answers, failing to distinguish *Agnes but not Beatrice* from *Agnes* and *Most linguists but not a lot of computer scientists* from *Most linguists*. The same problem arises with the decreasing answers in (22), except because these answers only include a "negative" component they are all mapped to the same output.

These observations are familiar from chapter 1, where they were used to show that Groenendijk and Stokhof, 1984-style exhaustivity makes wrong empirical predictions about the interpretation of marked answers. The hybrid order makes correct predictions about these cases. Since the default order is still available and generally preferred over the hybrid order, however, there must be something infelicitous about the reasoning schema that would yield the hypothetical conclusions in (22) and (23). To rule out the use of the default order with these answers, we need to explain not only why these inferences are unattested, but why they would be pragmatically deviant if they were attested.

The explanation I would like to develop lies in general considerations of "information preservation". It follows from communicative pressure to avoid redundancy. Because minimizing the background extension essentially wipes out the particular semantic contribution of a decreasing answer, or the "negative" part of an unskewed answer, the default order does not make a good communicative heuristic when applied to such answers.

A better heuristic would be one that can distinguish between the members of this large semantic class, without mapping them all to the same output. The hybrid order meets this condition: it maps answers that are not positively skewed to a diverse range of outputs. This is because the hybrid order respects both the "positive" and "negative" semantic components of answers, by distinguishing between positive and negative restrictions.

The default order does not lead to informational redundancy when applied to positively skewed answers. Increasing answers and positively skewed nonmonotonic answers typically have distinct minimal elements—the minimal elements of *Agnes* and *Beatrice* are $\{a\}$ and $\{b\}$, respectively—which is why the default order does not map them all to the same output. For this reason, I suggest that the default order is used whenever it preserves the literal semantics of an answer, and the hybrid order is used otherwise.

How should this principle be implemented? The preceding discussion suggests it should ensure that answers from a particular semantic class are interpreted using an order that can distinguish between the members of that class. This must be a consequence of the analysis, but cannot itself instantiate the analysis. We do not want a principle that says: use an order that maps all the members of some explicitly defined semantic class to different outputs.

There are precedents for such a principle in the literature. In the context of numeral semantics, for instance, Buccola and Spector, 2016 propose a filter on readings of sentences containing numerals where the particular choice of numeral makes no semantic contribution: a logical form S containing numeral n is unavailable if all logical forms generated from S by replacing n with another numeral m have the same truth-conditions. The situation in (22) and (23) is infelicitous for similar reasons: all negative answers of the form *not X*, with X a name, are mapped to the same output by the default order.

It is not obvious whether the requisite notion of semantic class can be defined in a principled way, however. Unskewed answers would seem to form a class, but the default order still distinguishes between them; the problem is rather that it ignores their "negative" component. Moreover, by quantifying over logical forms this strategy sneaks in formal alternatives through the backdoor.

Because the goal of this thesis is to describe exhaustivity in strictly model-theoretic terms, I cannot access this kind of analysis. These methodological commitments place rather strict boundaries on the available solution space. I will accordingly need to define a notion of information preservation in terms of how exhaustivity affects the semantics of an answer, and nothing more.

The principle we want can be defined on the basis of an answer's literal meaning and the exhaustive inferences that would result from using either the default or the hybrid order. A couple points of housekeeping are necessary before developing the analysis. First, because this principle must interact with the notion of quantificational skew, the semantic objects it compares must be generalized quantifiers. In particular, there must be such a thing as an exhaustified quantifier. Rather than taking a structured proposition as input, then, as in the "uncurried" lambda term (24-a), exhaustivity should be defined as an operation on generalized quantifiers, as in the "curried" lambda term (24-b).

(24) a. $\texttt{exh} = \lambda\langle T, P\rangle.T(P_w) \wedge \neg\exists v : T(P_v) \wedge v < w$
b. $\texttt{exh} = \lambda T\lambda P.T(P_w) \wedge \neg\exists v : T(P_v) \wedge v < w$

These terms are semantically equivalent (Curry, 1980; Schönfinkel, 1924), but (24-b) makes it possible to represent the effects of exhaustivity on quantifier meaning without directly mapping the input to a proposition, which is desirable if our goal is to examine how exhaustivity affects the formal properties of generalized quantifiers. The information preservation principle we want can only be defined if it is possible to saturate T without also saturating P. Exhaustivity should then take T and P as separate arguments, in that order. One might take this result as an argument in favor of locating exhaustivity in compositional semantics, as Groenendijk and Stokhof, 1984 did, but I do not view this argument as decisive.

Second, because the second argument of `exh` is over property intensions, we need intensional versions of the restriction operators in order to access the positive and negative restrictions of exhaustified quantifiers. Intensionalizing the restriction operators has not been necessary so far. The new restriction operators are defined in (25).

(25) **Restriction operators (intensional).**
a. $^{\uparrow}\mathbf{R} = \lambda T_{\langle s\langle\tau t\rangle\rangle t}.\bigcap\{S_w | \forall P : T(P) \leftrightarrow T(\lambda v\lambda x.P_v(x) \wedge S_v(x))\}$
b. $^{\uparrow}\mathbf{R}^{+} = \lambda T_{\langle s\langle\tau t\rangle\rangle t}\lambda x_\tau.\exists P, S : \neg T(S) \wedge T(P) \wedge P_w = S_w \cup \{x\}$
c. $^{\uparrow}\mathbf{R}^{-} = \lambda T_{\langle s\langle\tau t\rangle\rangle t}\lambda x_\tau.\exists P, S : \neg T(S) \wedge T(P) \wedge P_w = S_w - \{x\}$

$^{\uparrow}\mathbf{R}$ applies to a quantifier over property intensions, T, and intersects the extensions of all properties that T is closed under intersection with, just like the extensional version of $\mathbf{R}$. $^{\uparrow}\mathbf{R}^{+}$ applies to T and delivers a set of objects whose presence in a property's extension can make a difference as to whether T contains the intension of that property. $^{\uparrow}\mathbf{R}^{-}$ applies to T and delivers a set of objects whose absence from a property's extension can make a difference as to whether T contains the intension of that property. The intensional restriction operators work just like their extensional counterparts; the only difference is that they now have world arguments in the appropriate places in order to successfully extract information about the restrictions of quantifiers over property intensions.

It is now possible to compare the "literal" meanings of generalized quantifiers to their exhaustified versions. Table 3.2 shows a sample of how default ($\texttt{exh}_D(T)$) vs. hybrid ($\texttt{exh}_H(T)$) exhaustivity affects quantifier meaning.

Answer	T	$\mathtt{exh}_D(T)$	$\mathtt{exh}_H(T)$
Skew: $+$			
Agnes	$\lambda P.P(a)$	$\lambda P.P_w = \{a\}$	$\lambda P.P_w(a)$
Some of the cats	$\lambda P.\textbf{cats} \cap P$	$\lambda P.P_w \subset \textbf{cats}$	$\lambda P.\textbf{cats} \cap P_w$ $\wedge \textbf{cats} \not\subseteq P_w$
Every cat	$\lambda P.\textbf{cats} \subseteq P$	$\lambda P.P_w = \textbf{cats}$	$\lambda P.\textbf{cats} \subseteq P$
Agnes and Beatrice	$\lambda P.P(a) \wedge P(b)$	$\lambda P.P_w = \{a, b\}$	$\lambda P.P(a) \wedge P(b)$
Many cats	$\lambda P.n < \lvert P \cap \textbf{cats}\rvert$	$\lambda P.n < \lvert P_w \cap \textbf{cats}\rvert$	$\lambda P.n < \lvert P_w \cap \textbf{cats}\rvert$ $< \lvert\textbf{cats}\rvert$
Agnes or Beatrice	$\lambda P.P(a) \vee P(b)$	$\lambda P.P_w \subset \{a, b\}$	$\lambda P.\{a, b\} \not\subseteq P_w$
Skew: $\pm$			
Exactly 3 cats	$\lambda P.\lvert P \cap \textbf{cats}\rvert = 3$	$\lambda P.\lvert P_w \cap \textbf{cats}\rvert = \lvert P_w\rvert = 3$	$\lambda P.\lvert P_w \cap \textbf{cats}\rvert = 3$
Some but not all of the cats	$\lambda P.\textbf{cats} \cap P_w$ $\wedge \textbf{cats} \not\subseteq P_w$	$\lambda P.P_w \subset \textbf{cats}$	$\lambda P.\textbf{cats} \cap P_w$ $\wedge \textbf{cats} \not\subseteq P_w$
Skew: $-$			
Not Beatrice	$\lambda P.\neg P(b)$	$\lambda P.P_w = \varnothing$	$\lambda P.\neg P_w(b)$
No cats	$\lambda P.\textbf{cats} \not\cap P$	$\lambda P.P_w = \varnothing$	$\lambda P.\textbf{cats} \not\cap P$
Not all of the cats	$\lambda P.\textbf{cats} \not\subseteq P$	$\lambda P.P_w = \varnothing$	$\lambda P.\textbf{cats} \not\subseteq P_w$ $\wedge \textbf{cats} \cap P_w$
Not both Agnes and Beatrice	$\lambda P.\{a, b\} \not\subseteq P$	$\lambda P.P_w = \varnothing$	$\lambda P.\{a, b\} \not\subseteq P_w$ $\wedge P_w \cap \{a, b\}$
Not many cats	$\lambda P.\lvert P \cap \textbf{cats}\rvert < n$	$\lambda P.P_w = \varnothing$	$\lambda P.0 \neq \lvert P_w \cap \textbf{cats}\rvert < n$
Neither Agnes nor Beatrice	$\lambda P.\neg P(a) \wedge \neg P(b)$	$\lambda P.P_w = \varnothing$	$\lambda P.\neg P_w(a) \wedge \neg P_w(b)$
Skew: n/a			
Agnes but not Beatrice	$\lambda P.P(a) \wedge \neg P(b)$	$\lambda P.P_w = \{a\}$	$\lambda P.P_w(a) \wedge \neg P_w(b)$
Many cats but not many hamsters	$\lambda P.n < \lvert P \cap \textbf{cats}\rvert$ $\wedge \lvert P \cap \textbf{hamsters}\rvert < m$	$\lambda P.n < \lvert P_w \cap \textbf{cats}\rvert$ $\wedge P_w \subset \textbf{cats}$	$\lambda P.n < \lvert P_w \cap \textbf{cats}\rvert$ $< \lvert\textbf{cats}\rvert \wedge 0 <$ $\lvert P_w \cap \textbf{hamsters}\rvert < m$
If Agnes, then Beatrice	$\lambda P.P(a) \to P(b)$	$\lambda P.P_w = \varnothing$	$\lambda P.P_w(a) \leftrightarrow P_w(b)$

Table 3.2: Literal vs. exhaustified quantifiers.

If the default order is applied to an increasing quantifier, the output of exhaustivity is a set of properties whose extensions in the actual world contain only the minimal elements of the original quantifier. Default exhaustivity maps *Agnes* to a set of properties whose extensions only contain Agnes. It maps *some of the cats* to a set of properties whose extensions contain some but not all of the cats, and no individuals who are not cats.

If the hybrid order is applied to an increasing quantifier, the output of exhaustivity is either identical to the input, or a new quantifier that only contains properties whose extensions have minimal intersections with the original quantifier's positive restriction. Hybrid exhaustivity has no effect on *Agnes*, but maps *some of the cats* to a quantifier whose elements only contain some but not all of the cats, and possibly other individuals who are not cats.

If the default order is applied to any decreasing quantifier, the output of exhaustivity is a set of properties whose extensions in the actual world are empty. This is because the empty set is the minimal element of all decreasing quantifiers. If the hybrid order is used, the output of exhaustivity is either identical to the input, or a new quantifier that only contains elements whose intersection with the original quantifier's negative restriction is maximal. Hybrid exhaustivity has no effect on *not Beatrice*, but maps *not all of the cats* to a quantifier whose elements contain some but not all of the cats, and possibly other individuals who are not cats.

If the default order applies to doubly skewed quantifiers, the output of exhaustivity is a set of properties whose extensions in the actual world contain only the minimal elements of the original quantifier. The result of applying default exhaustivity to *exactly 3 cats* is a set of properties whose extensions contain 3 cats and no other individuals. If the hybrid order is used, the output is typically identical to the input.

If the default order applies to unskewed quantifiers, the output of exhaustivity is a set of properties whose extensions in the actual world contain only the minimal elements of the original quantifier. In case an unskewed quantifier is equivalent to the intersection of an increasing and a decreasing quantifier, the output of default exhaustivity will be the same as if it had applied to only the increasing quantifier; otherwise the output of default exhaustivity will be a set of properties with empty extensions. Default exhaustivity maps *Agnes but not Beatrice* to a set of properties whose extensions only contain Agnes. It maps *if Agnes then Beatrice* to the empty set.

If the hybrid order applies to unskewed quantifiers, the output of exhaustivity is a set of properties whose extensions have minimal intersections with the original quantifier's positive restriction and maximal intersections with the original quantifier's negative restriction. This operation has no effect on quantifiers that fix the status of every object in the restriction relative to the input property, like *Agnes but not Beatrice*, but does affect other quantifiers. Hybrid exhaustivity maps *many cats but not many hamsters* to a quantifier whose elements contain many but not all cats, and some but not many hamsters. It maps *if Agnes then Beatrice* to a quantifier whose elements either contain both Agnes and Beatrice, or neither Agnes nor Beatrice.

All these results should be familiar. The point is that the formal objects we are working with, exhaustified quantifiers, make it possible to derive the role of skew in order selection. To see how exhaustivity can preserve or eliminate semantic information, it is useful to compare the restrictions of literal vs. exhaustified quantifiers. Table 3.3 shows how exhaustivity with the default order affects an answer's restriction.

Answer	**R**(T)	$^{\uparrow}\mathbf{R}(\mathtt{exh}_D(T))$	$^{\uparrow}\mathbf{R}^{+}(\mathtt{exh}_D(T))$	$^{\uparrow}\mathbf{R}^{-}(\mathtt{exh}_D(T))$
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: $\pm$				
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

Table 3.3: Restrictions for default exhaustified quantifiers.

The pattern that emerges is that the literal quantifier's restriction is formally recoverable from the exhaustified quantifier using ${}^{\uparrow}\mathbf{R}^{+}$ whenever a quantifier is positively skewed, but not otherwise. I claim that this difference drives order selection.

Because the default order distinguishes all possible background property extensions, quantifiers exhaustified by the default order are restricted by the entire domain of discourse. In this way, default exhaustivity always expands a quantifier's total restriction. However, quantifiers vary in whether their positive and negative restrictions are preserved. In particular, because default exhaustivity invariably excludes objects from the extension of the question-property, it typically preserves positive restrictions while expanding negative restrictions. For this reason, the total restrictions of positively skewed quantifiers will remain accessible from their default exhaustified counterparts as long as the default order does not alter a quantifier's positive restriction. This is because the restrictions of positively skewed quantifiers are identical to their positive restrictions, so we can still access the original total restriction whenever the positive restriction is preserved.

By contrast, there is no guarantee that default exhaustivity will preserve the restrictions of quantifiers that are not positively skewed. In the case of decreasing quantifiers, which are negatively skewed and have empty positive restrictions, expanding the negative restriction to the domain of discourse completely wipes out the original restriction. In the case of unskewed quantifiers, whose positive and negative restrictions are distinct and nonempty, default exhaustivity either preserves the positive restriction while expanding the negative restriction, or in the case of conditionals erases both; in either case, some or all of the original restriction becomes unrecoverable.

Table 3.4 shows how exhaustivity with the hybrid order affects an answer's restriction.

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

Table 3.4: Restrictions for hybrid exhaustified quantifiers.

Unlike default exhaustivity, hybrid exhaustivity never changes the total restriction: the second and third columns of Table 3.4 are identical. This is because the hybrid order only regulates the question-property's intersection with an answer's restriction, so objects outside the restriction cannot make a difference. When the hybrid order has an effect, it is to move objects around between the positive and negative restriction. The negative restriction for *some of the cats* is empty, but its hybrid exhaustified counterpart is negatively restricted by the cats; similarly, the positive restriction for *not all of the cats* is empty, but its hybrid exhaustified counterpart is positively restricted by the cats. However, the total restriction of a quantifier exhaustified by the hybrid order cannot change.

I propose the principle Preserve Restriction (26). It compares a GQ's potential exhaustive meanings to its literal meaning, and requires the total, positive, or negative restriction of the former to be identical to the total restriction of the latter.

(26) **Preserve restriction.**
Ensure that:

a. $^{\uparrow}\mathbf{R}(\texttt{exh}(T)) = \mathbf{R}(T)$, or
b. $^{\uparrow}\mathbf{R}^{+}(\texttt{exh}(T)) = \mathbf{R}(T)$, or
c. $^{\uparrow}\mathbf{R}^{-}(\texttt{exh}(T)) = \mathbf{R}(T)$.

Preserve Restriction is what ensures that exhaustive interpretation respects the informational content of an answer's literal meaning. In chapter 2, I proposed to formalize the notion of "aboutness" using quantificational restriction: an answer is about the objects in its restriction. Preserve Restriction says that exhaustivity should not change the objects an answer is about.

The idea behind this principle is that the semantic content of an utterance should play a nontrivial role in its pragmatic meaning. If the default order were used with a decreasing quantifier like *not Beatrice*, the restriction of that quantifier would be irrelevant to the final exhaustive interpretation; it would pointlessly single out Beatrice. Similarly, if the default order were used with an unskewed quantifier like *most linguists but not a lot of computer scientists*, then at least some of the restriction (the computer scientists) would be pointlessly singled out. Preserve Restriction is therefore a reasonable principle to use in communication on the assumption that language users make utterances whose informational content is not arbitrary.

Preserve Restriction does not insist that an exhaustive quantifier share its total restriction with its literal counterpart. If this were the case, all answers would have restricted extent; the default order would never be used, and it would be impossible for exhaustivity to license any conclusions about objects outside an answer's restriction. Rather, Preserve Restriction is satisfied as long as an answer's restriction is preserved in the positive or negative restriction of its exhaustified counterpart. What matters is that the answer's restriction play some role in the partition between objects that have vs. lack the question-property.

This idea recalls von Stechow and Zimmermann's claim that exhaustivity "yields the proposition that everything outside the answer term's quantificational domain behaves just the other way" (von Stechow and Zimmermann, 1984, pg. 11): i.e., their claim that exhaustive interpretation of positive answers excludes all objects outside the answer's restriction from the question-property's extension, and exhaustive interpretation of negative answers excludes all objects outside the answer's restriction from the complement of the question-property's extension. As we have seen, this picture is not accurate: it neglects the asymmetry between positive and negative answers, and it ignores exhaustive inferences regarding objects inside an answer's restriction. Still, the proposal vindicates their more general claim that quantificational restrictions play a role in determining the extension of the question-property.

The conceptual status of PRESERVE RESTRICTION may depend on whether exhaustivity is analyzed as semantic or pragmatic. PRESERVE RESTRICTION cannot be violated: there are no attested exhaustive inferences, marked or unmarked, that render the literal answer's restriction unrecoverable. If we adopt a semantic view, it could then make sense to build this principle into the definition of `exh`, as a sort of presupposition restricting the resolution of the $<$ parameter. However, the idea that semantic content must be recoverable from pragmatic interpretation could have broader applications, and I would like to leave open the possibility that the `exh`-specific implementation of this principle derives from something more general.

One result of PRESERVE RESTRICTION is that no class of quantifiers will all be mapped to the same output. As long as quantifiers within a class have different restrictions, the positive or negative restrictions of their exhaustified counterparts will likewise differ. In this way, the analysis implements the idea that exhaustivity should substantively distinguish between members of a semantic class, and that there should be some "point" to using a particular answer. Explicit appeal to the interchangeability of formal alternatives is not necessary to achieve this result, however. It falls out as a consequence of more basic pressure to preserve information.

3.2.2 Interactions with markedness

The order selection pattern and the correlation with skew follows from how the principles of REDUCE MARKEDNESS and PRESERVE RESTRICTION interact. Specifically, I propose that speakers coordinate on the least marked order that satisfies PRESERVE RESTRICTION.

The correlation with skew is derived from the fact that the default order only preserves the restrictions of positively skewed answers. Typically, the default order expands an answer's restriction without altering the positive restriction. If the original total restriction is identical to the original positive restriction, then preserving the positive restriction is sufficient to preserve the total restriction. Positively skewed answers meet this condition.

For instance, exhaustive interpretation of B's positively skewed, increasing answer in (27) uses the default order because it is possible to do so while preserving the restriction.

(27) **A**: Who attended the last panel?
B: Agnes or Beatrice.
$\rightsquigarrow$ *not both, nobody else*
$^{\uparrow}\mathbf{R}^{+}(\mathtt{exh}_D(T)) = \{a, b\} = \mathbf{R}(T)$

The total restriction of the literal disjunction is the set containing Agnes and Beatrice. The total restriction of the disjunction exhaustified by the default order consists of the entire domain of discourse, but its positive restriction is still the set containing Agnes and Beatrice. The default order satisfies PRESERVE RESTRICTION for this reason. The hybrid order also would have satisfied PRESERVE RESTRICTION, but because the default order is the least marked order, REDUCE MARKEDNESS dictates that it should be used whenever possible. This is why exhaustivity has negative direction and unrestricted extent in (27).

Exhaustive interpretation of B's doubly skewed answer in (28) similarly uses the default order because it can be used while preserving the answer's restriction.

(28) **A**: Who attended the last panel?
B: Exactly three linguists.
$\rightsquigarrow$ *nobody else*
$^{\uparrow}\mathbf{R}^{+}(\mathtt{exh}_D(T)) = \textbf{linguists} = \mathbf{R}(T)$

The total restriction of the literal quantifier *exactly three linguists* is the set of linguists. The total restriction of the quantifier exhaustified by the default order is the entire domain of discourse, but its positive restriction is still the set of linguists. The default order satisfies PRESERVE RESTRICTION for this reason, so REDUCE MARKEDNESS selects it as the least marked order that can be used felicitously. This is why exhaustivity has negative direction and unrestricted extent in (28).

However, the default order does not preserve the restrictions of answers that are not positively skewed. This is because default exhaustivity typically expands an answer's negative restriction. As long as there are objects in the original total restriction that only appear in the negative restriction, formal access to these objects will be lost if the default order is used. Whenever the default order eradicates part or all of the original restriction, PRESERVE RESTRICTION rules it out and the hybrid order is therefore selected.

For instance, exhaustive interpretation of B's negatively skewed, decreasing answer in (29) uses the hybrid order because the default order would fail to preserve the restriction.

(29) **A**: Who attended the last panel?
B: Not all of the computer scientists.
$\rightsquigarrow$ *some of them did*
$\not\rightsquigarrow$ *conclude nothing about the linguists*
$^{\uparrow}\mathbf{R}(\mathtt{exh}_H(T)) = \textbf{linguists} = \mathbf{R}(T)$

The total restriction of the literal quantifier *not all of the computer scientists* is the set of computer scientists. If this quantifier were exhaustified by the default order, the exhaustified quantifier's new total restriction would be the entire domain of discourse. Its new negative restriction would also be the domain of discourse, and its new positive restriction would be empty. None of these sets are identical to the set of computer scientists, so PRESERVE RESTRICTION rules out the default order. By contrast, applying the hybrid order would create an exhaustified quantifier whose total restriction, positive restriction, and negative restriction all consist of the set of computer scientists. The hybrid order thus preserves the answer's restriction and is selected for use. This is why exhaustivity has positive direction and restricted extent in (29).

Exhaustive interpretation of B's unskewed answer in (30) similarly uses the hybrid order because the default order would fail to preserve the restriction.

(30) **A**: Who attended the last panel?
B: Most syntacticians, but not a lot of phonologists.
$\rightsquigarrow$ *not all syntacticians, some phonologists*
$\not\rightsquigarrow$ *conclude nothing about anyone else*
$^{\uparrow}\mathbf{R}(\mathtt{exh}_H(T)) = \textbf{syntacticians} \cup \textbf{phonologists} = \mathbf{R}(T)$

The total restriction of B's literal answer is the set of syntacticians and phonologists. If this quantifier were exhaustified by the default order, the resulting quantifier's new total restriction and negative restriction would both be the entire domain of discourse, while the positive restriction would be the set of syntacticians. None of these sets are identical to the set of syntacticians and phonologists, so Preserve Restriction rules out the default order. Applying the hybrid order would create an exhaustified quantifier whose total, positive, and negative restrictions all consist of the set of syntacticians and phonologists, so the hybrid order preserves the restriction and is selected for use. This is why exhaustivity has mixed direction and restricted extent in (30).

The general pattern that emerges from the interaction of Reduce Markedness and Preserve Restriction is that the least marked order, the default order, can only satisfy Preserve Restriction when applied to positively skewed answers. The default order is therefore selected for use with these answers only. Exhaustive interpretation of decreasing and unskewed answers can only satisfy Preserve Restriction when the more marked hybrid order is used. The hybrid order is therefore selected for use with answers that are not positively skewed.

One might wonder whether default exhaustivity always preserves the restrictions of positively skewed quantifiers, or whether this only holds in specific cases. In fact, this connection does not hold in the general case. The exceptions involve quantifiers like (31) (I omit set brackets around the elements of T).

(31) $T = \{a, ad, b, bd, abc, abcd\}$

T is a quantifier whose positive restriction is larger than the union of its minimal elements. Suppose the domain of discourse is $\{a, b, c, d\}$. Then T is positively skewed: its total and positive (and negative) restrictions are $\{a, b, c\}$. Each element of $\{a, b, c\}$ can be added to a set outside T to create a set in T: a and b can both be added to the empty set, and all three objects can be added to the set containing the other two objects. However, the exhaustified quantifier that would result from use of the default order would only be positively restricted by $\{a, b\}$. The reason for this disparity is that c appears in T's positive restriction, yet there is no minimal element of T that contains c. This example shows that there is no necessary link between positive restriction and the minimal elements of a quantifier.

It is difficult to think of natural linguistic examples that would express the meaning of T in (31) while remaining extensional. The complex biconditional answer in (32) appears to express this meaning, but certainly involves modal quantification and is thus beyond the reach of the "basic" circumscription approach I have been developing.

(32) At least Agnes or Beatrice attended, but not both unless Charmian also attended, and if Charmian was there then both Agnes and Beatrice were there.

If extensional quantificational answers like (32) do exist, however, the proposal in this chapter predicts that exhaustivity should use the hybrid order when applied to such answers. Even though these answers are positively skewed, the default order would fail to satisfy Preserve Restriction.

There is one kind of answer whose restriction is preserved by both orders. These are answers restricted by the entire domain of discourse. Because the default and hybrid orders only regulate the question-property's intersection with the domain of discourse D, exhaustivity cannot fail to preserve the restriction of any answer restricted by D rather than a strict subset of D regardless of which order is chosen. Examples like (33) arguably instantiate this situation.

(33) **A**: Which linguists attended the panel?
B: Not all of them. $\rightsquigarrow$ *some of them did*

The domain of discourse for A's question is the set of linguists, which also restricts B's answer. Exhaustivity has positive direction in (33), so it should use the hybrid order. Yet the default order, which would map B's answer to a quantifier true of any property whose extension does not contain linguists, also satisfies Preserve Restriction.

As defined in (26), the principle of Preserve Restriction is extremely sensitive to domain selection, predicting that the direction and extent of negative and mixed answers should vary depending on whether they are restricted by D or a strict subset of D. However, exhaustivity is not hypersensitive in this way: the hybrid order is still chosen for these answers even when the default order would preserve their restrictions.

This finding should not come as a surprise, since there is always uncertainty about the exact domain of discourse in any communicative context. Because the default order preserves the restrictions of positively skewed answers in any context, it is a "safe" option even if conversational participants are imperfectly coordinated on the domain. By contrast, the default order preserves the restrictions of answers that are not positively skewed only in particular and somewhat exceptional contexts, and is thus a communicatively "risky" option to use, since it might violate Preserve Restriction in case the speaker and addressee are not perfectly coordinated. There will always be some wider domain relative to which the default order does not preserve the restrictions of negative and mixed answers.

One solution to this problem would be to modify Preserve Restriction so as to include universal quantification over all possible domains, or require it to hold relative to the widest domain. However, this requirement would not deliver the desired result for second-order questions, whose answers cannot even in principle be interpreted exhaustively relative to the widest possible domain because of complementation.

Going forward I will assume that conversational participants stick only to orders that are pragmatically "safe". I leave open how best to formalize this idea. Even if default exhaustivity would happen to preserve an answer's restriction in a particular context, the possibility that speakers are imperfectly coordinated on quantifier domains motivates caution in order selection.

3.2.3 Summary

The direction and extent properties of answers are not arbitrary, but can be derived from independent pragmatic principles. I argued for two order selection constraints, Reduce Markedness and Preserve Restriction, that jointly guide order selection. Preserve Restriction requires the restriction of an answer's literal meaning to be accessible from its exhaustive interpretation using one of the restriction operators. It can be thought of as ensuring that exhaustivity does not change what an answer is "about".

Reduce Markedness insists that speakers use the least marked order they can. Because the default order is less marked than the hybrid order, this principle enforces a preference for the default order, but only insofar as it can be used without violating Preserve Restriction. This is indeed the case with positively skewed answers, but not with decreasing or unskewed answers. The result is that positively skewed answers use the default order, while answers that are not positively skewed use the hybrid order.

The overall analysis has a semantic and a pragmatic component. The semantic part comprises the conventional meanings of linguistic answers and the available inventory of orders speakers can use to enrich meaning. The pragmatic part comprises the selection constraints that speakers use to decide on an order. Because the order selection facts do not have to be stipulated, but follow from more general constraints, the proposal has reached the larger goal of explanatory adequacy.

The orders play the role alternatives play in theories that use them. To my knowledge, the analysis I have developed in this thesis is the first formal proposal to derive the facts about scales and symmetry breaking from pragmatic principles while remaining fully explicit, predictive, and falsifiable. In contrast, theories that use lexical scales (Gazdar, 1979; Horn, 1972) are basically descriptive, and even theories that aim for greater predictive power, like the structural theory (Katzir, 2007), typically do not attempt to explain why the descriptive generalizations they hypothesize hold. 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. It is a strength of the circumscription account that it lends itself to such an explanation.

3.3 The dual order

This section discusses the rare, but attested phenomenon of negative answers with unrestricted extent. Usually, it is impossible for exhaustive interpretation of negative answers to have unrestricted extent; this observation is what motivated the analysis I have developed in this thesis. However, there are exceptions to this generalization and it is important to show that the framework can account for them.

Decreasing answers are usually interpreted using the hybrid order, which enforces positive direction and restricted extent. To account for decreasing answers with unrestricted extent, I suggest that a third order over possible worlds, the DUAL ORDER, is applied. This order is the dual of the default order and forms a superset-induced lattice over all possible background extensions. If a decreasing answer is interpreted using the dual order, exhaustivity will have positive direction and unrestricted extent.

The dual order is available but extremely marked, so it is very rarely used. I will argue that the introduction of the dual order is not ad hoc, but that it slots naturally into the independently motivated markedness hierarchy developed in this chapter. In this way, the analysis not only explains the general pattern but provides a principled account of the exceptions and explains why they are exceptional.

3.3.1 Negative answers with unrestricted extent

An important aspect of the analysis I have developed so far involves the inventory of orders over possible worlds that exhaustivity can use, the default order and the hybrid order. These are fairly natural choices for communicative heuristics. The default order is just a subset-induced lattice, arguably the simplest possible algebraic structure made available by the logical foundations of language. The hybrid order is more complex, but represents an effective method of reasoning given multiple sources of positive and negative information. This chapter has shown how order selection can be motivated pragmatically, given the conventional meaning of an answer and the principles of REDUCE MARKEDNESS and PRESERVE RESTRICTION. However, that these are the available orders and not others does not follow from anything deeper in the analysis; they must be taken as basic.

One might wonder why these are the only two orders. In any logic with a negation operation, ordering relations can be reversed. Ideally, the facts about direction, extent, and symmetry breaking would derive entirely from classical logic and the markedness of negation. The question then arises as to why the default order cannot be reversed. If it were possible for exhaustivity to use the dual of the default order, we would expect some answers to have positive direction and unrestricted extent at least some of the time. Although one might suppose that the markedness of negation is what explains why this option is unavailable, this is not a satisfying explanation. Natural languages treat negation as marked, but they still include a negation operation.

In fact, there is evidence that the default order can, in certain contexts, be reversed. von Stechow and Zimmermann, 1984 thought this was possible, and proposed a theory of exhaustivity according to which negative answers always have unrestricted extent. The speakers I consulted mostly gave different judgments, indicating that negative answers usually have restricted extent, but some speakers claim that unrestricted extent is sometimes a possibility. van Rooij and Schulz, 2004 reached similar conclusions based on the speakers they consulted as well.

The contexts in which exhaustive interpretation of negative answers has unrestricted extent are typically contexts in which there is some extra pragmatic factor motivating it. In (34), speakers report that it is possible to understand B's response as implicating that some graduate students were experimented on, as predicted by the hybrid order, but also that unspeakable things happened to the undergraduates.

(34) **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.
$\rightsquigarrow$ *some graduate students*
$\rightsquigarrow$ *the undergraduates were subjected to it*

The "extra" factor in (34) is apparently B's reluctance to give a positive answer out loud. Additionally, because of the horrific nature of the experiment, B should be highly motivated to explicitly state all true and assertable "negative" information. In such a context, speakers are more willing to reach "positive" conclusions via exhaustivity.

This example should be contrasted with (35), where there is no special pragmatic motivation one way or another given A's less loaded question. In this "neutral", information-seeking context, speakers report that nothing can be concluded about the undergraduates from B's answer.

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

One might expect the judgments in (35) to pattern more like (34) in a context where attending the department barbecue means being subjected to physical or psychological torture, or some other awful ordeal.

Suppose A walks into a bike shop and has the exchange in (36) with mechanic B, who gestures to a nearby pair of bikes. Some speakers report that it is possible to understand B's answer as implicating that all the other bikes in the shop are for sale.

(36) **A**: Which of these bikes are for sale?
B: Not those two. $\rightsquigarrow$ *the others are*

The "extra" factor in (36) seems to be A's desire to buy a bike, together with the background assumption that all bikes are for sale unless noted otherwise. In this context, because B is motivated to clearly indicate to A exactly which bikes are off limits, speakers are willing to draw exhaustive inferences with unrestricted extent.

Imagine the (37) exchange in the same scenario. In this information-seeking context, there is no assumption in this context that all bikes are tubeless unless noted otherwise, and the speakers I consulted were not willing to draw exhaustive inferences from B's answer in (37): A's idle curiosity about tubeless tires is apparently not sufficient motivation for unrestricted extent in this context.

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

In (38), however, A's motivation to buy a bike with tubeless tires makes unrestricted extent more natural. At least some speakers were willing to conclude from B's answer in (38) that all the other bikes in the shop have tubeless tires.

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

Although the notion of "extra pragmatic motivation" should be given a more precise characterization in future work, these observations suggest that unrestricted extent is a highly marked, but available option for negative answers.

van Rooij and Schulz, 2004, who only recognized the default order, proposed that when negative answers are interpreted with unrestricted extent, speakers are applying exhaustivity to the complement of the question-property instead of the original question-property. However, they have no explanation for why this should happen, or what triggers the shift from one question-property to another. The account I have developed in this thesis provides the answer: as we have seen, the default order cannot be applied to decreasing answers without violating Preserve Restriction.

In the present framework, there are multiple orders exhaustivity can use, and order selection is systematically predictable from the semantics of an answer. Instead of applying the default order to a different question, I propose that a third order, the dual order, is used. The dual order is just the inverse of the default order and is defined in (39).

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

Like the default order, the dual order ranks worlds by the size of background property extensions in those worlds: a world v is more minimal than w iff the relevant subset of P's extension in v includes the relevant subset of P's extension in w. Like the default order, the dual order has the structure of a Boolean lattice, but larger property extensions are now more minimal rather than smaller ones.

If our domain of entities includes three individuals, Agnes, Beatrice, and Charmian, and there are no postulates restricting the possible extensions of some property P, then the dual order over P-extensions is diagrammed in Figure 3.1, with $\supset$ proceeding from bottom to top.

If the dual order is used in the (34)-(38) examples, it will generate exhaustive implications with positive direction and unrestricted extent. It will generally enforce the implication that everyone has the question-property except those individuals explicitly ruled out by the answer. Applied to B's answer in (36), the dual order allows us to conclude that all the bikes in the shop have the property of being for sale, except the bikes B is pointing at.

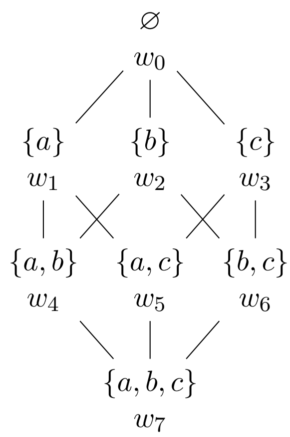

Figure 3.1: The dual order.

Augmented with the dual order, the system developed in this thesis resembles classical logic much more faithfully than it did before. Negation may be marked, but it still exists. Accordingly, there is no reason to suppose that speakers cannot reason using both an order and its dual. Just as generalized quantifiers can be increasing, decreasing, or nonmonotonic, the default, hybrid, and dual orders appear to represent all possible order polarities: an order can be positive, negative, or mixed. Ruling out any of these possibilities would require special stipulation, so it is not surprising that they are all attested.

3.3.2 The markedness hierarchy

The addition of the dual order concludes the presentation of my proposal. It slots into the system as follows: the dual is the most marked order, more marked than the hybrid order, which is in turn more marked than the default (40). The principle of Reduce Markedness should also be revised (41).

(40) **The markedness hierarchy.**
`default` $\succ$ `hybrid` $\succ$ `dual`

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

The central claim of Horn, 1989 is that negation is semantically classical but pragmatically marked. In a formal logic, the role of negation is to reverse the entailment order over truth values. I conjecture that the markedness of linguistic negation reflects this logical property: order reversal is marked in natural language. The hierarchy in (40) can then be grounded in the general markedness of order reversal.

The three orders, default, hybrid, and dual, are natural communicative heuristics to use given an underlyingly classical semantics. Default and dual simply correspond to inclusion and its inverse, which are available relations in almost any logic. The hybrid order is more complex, but represents a plausible response to the fact that it is often necessary to reason about both positive and negative information at once. Just like answers themselves, the heuristics available to enrich answers can have positive, negative, and mixed polarity. This is the classical semantic part of the analysis.

The logical principle of duality says that order is always reversible. If order selection respected the principle of duality, positive, negative, and mixed answers would be pragmatically equal. One might then expect the following order selection pattern: increasing answers would use the default order, decreasing answers would use the dual order, and nonmonotonic answers would use the hybrid order. However, the actually attested order selection pattern does not take this neatly schematic shape. Instead, order selection is sensitive to a hierarchy that prefers the default order to the hybrid order, which is preferred to the dual order. This aspect of the analysis reflects the pragmatic markedness of negation.

In chapter 2, I suggested that the default order is less marked than the hybrid order because it induces a simpler algebraic structure. In light of this claim, it is striking that the dual order is more marked than the hybrid order. Several of the factors that make the hybrid order more marked than the default order would appear to favor the dual order too, which also induces a simpler lattice structure; additionally, the dual order can license stronger conclusions in virtue of enforcing unrestricted extent. However, the normal pattern for negative answers involves restricted extent, not unrestricted extent, indicating that the hybrid order is preferred.

Apparently, the complexity-driven markedness of the hybrid order still does not surpass that of the dual order, which is driven by the absolute markedness of negation. One way to motivate the markedness hierarchy in (40), then, is to view it as tracking the "amount" of negation inherent in an order. The default order is the least marked option because it enforces only minimization, not maximization; it is not negative at all. The hybrid order, which incorporates both minimization and maximization, is partially but not entirely negative, and the dual order, which enforces only maximization, is entirely negative and is therefore the most marked option. In this way, the markedness hierarchy over the inventory of orders is a special case of a much broader generalization, namely the fact that negation is pragmatically marked in natural language.

Although speakers tend to prefer the least marked order consistent with PRESERVE RESTRICTION, a variety of additional pragmatic factors can motivate the switch to a more marked order. This is inevitably the case whenever the dual order is used. Because the hybrid order always preserves an answer's restriction, there must be significant, context-specific pressure to use a more marked order despite the felicity of a less marked one. In particular, the dual order is used in situations where the speaker or the addressee's goals depend in some way on whether all the true, relevant "negative" information is explicitly asserted. I leave it open whether this source of pragmatic sensitivity could be given a precise formalization.

Markedness is not the only factor limiting the use of the dual order. Exhaustivity must also satisfy PRESERVE RESTRICTION. In exactly the same way that decreasing and unskewed answers cannot use the default order, increasing and unskewed answers cannot use the dual order without violating PRESERVE RESTRICTION. The dual order maps all increasing quantifiers to a new set of properties whose extensions include the domain of discourse, an operation that inevitably wipes out information about the input quantifier's restriction.

Typically, the dual order preserves an unskewed quantifier's negative restriction, but fails to preserve the positive restriction. It maps *Agnes but not Beatrice* to a set of properties whose extensions include all individuals but Beatrice, losing the input quantifier's positive restriction (Agnes), and it maps *most linguists but not a lot of computer scientists* to a set of properties whose extensions include not a lot of computer scientists, but all individuals who are not computer scientists, thereby losing the original quantifier's positive restriction (the linguists). Unskewed answers must use the hybrid order: neither the default nor the dual order can be used consistently with PRESERVE RESTRICTION.

By contrast, doubly skewed answers can typically use all three orders. In the same way that the default order expands the negative restrictions of doubly skewed answers while leaving their positive restrictions unchanged, the dual order expands their positive restrictions while leaving their negative restrictions unchanged. The dual order maps *exactly three* *linguists* to a set of properties whose extensions include exactly three linguists and every nonlinguist. This new exhaustified quantifier, which is positively restricted by the domain of discourse, is still negatively restricted by the set of linguists. However, although all three orders are available for doubly skewed answers, the default is the least marked, so Reduce Markedness favors its use.

One might wonder whether there is ever motivation for positively skewed answers to use a more marked order than the default. If Reduce Markedness is violable, this should in principle be possible. If the hybrid order were used with increasing answers, exhaustivity would have negative direction and restricted extent. For instance, in an exchange like *A: Who attended the last panel? B: Most of the linguists*, B's answer would still implicate that not every linguist attended, but nothing would be concluded about computer scientists.

The speakers I consulted gave judgments indicating that this outcome is sometimes possible. There is a strong preference for unrestricted extent, but some (linguistically naïve) speakers commented that the implication that not every linguist attended feels more robust than the implication that no computer scientists attended; the former is almost always attested, but for some speakers the latter feels "optional". If this observation (which should be validated experimentally) holds up, the markedness hierarchy I have proposed can make sense of it: the latter implication is only licensed by the default order, whereas the former persists even if speakers switch to the hybrid order.

Although I follow a long tradition in rejecting the distinction between exhaustive implications and scalar implicatures, attributing both phenomena to the same source, these descriptive categories can be reconstructed within the analysis I have proposed. Scalar implicatures are just exhaustive implications about individuals inside an answer's restriction (*most* $\rightsquigarrow$ *not all*, *not all* $\rightsquigarrow$ *some*, and so on). To the extent that traditional scalar implicatures are more robust than exhaustive implications about individuals outside an answer's restriction, this difference can be attributed to the fact that the default and hybrid orders both license the former, whereas only the default order licenses the latter.

We can ask a similar question about doubly skewed answers: given that all three orders can typically be applied to such answers consistently with Preserve Restriction, is there ever motivation to depart from the default? The hybrid order maps most doubly skewed quantifiers to themselves, so this option should not be ruled out. Most speakers I consulted did not accept inferences that could result from application of the dual order to a doubly skewed answer, however. In an exchange like *A: Which of these bikes are for sale? B: Exactly three mountain bikes*, it is difficult or impossible to hear B's answer as implicating that every nonmountain bike is for sale.

To generate this inference, two less marked orders (the default and the hybrid) would have to be rejected in favor of the dual. Because this option requires moving two levels down on the hierarchy instead of one, it is not surprising that it is virtually impossible.

The explanation for the markedness hierarchy as reflecting the markedness of negation shares some features in common with the structural theory (Katzir, 2007), which also appeals to the markedness of negation to address the symmetry problem. For the structural theory, only equally or less complex sentences count as formal alternatives to some sentence; more complex sentences are ruled out. The notion of markedness that matters is a syntactic one: longer sentences are more marked than shorter ones. Negatively restricted formal alternatives are ruled out in virtue of being marked in this syntactic sense.

Crosslinguistically, it is extremely common for negation to be expressed overtly (Givón, 1978; Greenberg, 1966; Horn, 1989). The syntactic markedness of negation is presumably not accidental, but reflects something about its meaning. Instead of linking the resolution of the symmetry problem to syntactic markedness, which I take to be more symptom than cause, I have proposed that the direction and extent facts follow from a semantic markedness hierarchy. I suggested that the markedness of negation stems from its function as an order-reversing operator, and developed a markedness schema that disfavors partially or completely reversed orders relative to the default relation of inclusion. In this way, the solution to the symmetry problem does lie in the markedness of negation, but the connection is different from how the structural theory envisions it.

In chapter 1, the observation that negative answers are preferentially interpreted with positive direction and restricted extent provided a key argument against the structural theory. Depending on which constraints on the derivation of formal alternatives are adopted, negation either gives rise to NEGATION-INDUCED SYMMETRY among formal alternatives—two formal alternatives that contradict each other—or it does not. To predict restricted extent, the structural theory should not be allowed to bypass negation-induced symmetry. To predict positive direction, however, the structural theory must be allowed to bypass it. I argued that the structural theory therefore cannot account for negative answers, which have positive direction and restricted extent simultaneously.

The introduction of the dual order to the system adds some nuance to this conclusion, however. One might wonder whether the structural theory could account for the observation that negative answers do sometimes have unrestricted extent, by adopting constraints on formal alternatives that are violable. Perhaps negation-induced symmetry is the norm, but in exceptional circumstances speakers can adopt a marked derivation schema that bypasses it—the atomicity constraint proposed by Trinh and Haida, 2015, for instance.

Regardless of how this strategy is implemented, it would not address the basic problem, which is that positive direction and restricted extent are independent. If negation-induced symmetry were the default, but avoidable in special cases, the structural theory could explain why negative answers usually have restricted extent, and sometimes exceptionally allow unrestricted extent. However, it then predicts that positive direction patterns the same way, contrary to the observation that positive direction is the norm.

The system I have proposed makes the proper cut by distinguishing the hybrid order and the dual order. Both enforce positive direction, but only the dual order allows unrestricted extent. In addition to capturing the basic pattern for decreasing answers, the separate availability of the hybrid order is necessary to treat mixed answers like *Most syntacticians but not a lot of phonologists attended*.

Because this answer denotes an unskewed quantifier, PRESERVE RESTRICTION requires it to use the hybrid order: the default order would wipe out the negative restriction, and the dual order would wipe out the positive restriction. Empirically, there are no circumstances in which this answer can convey that no computer scientists attended, or that every computer scientist attended, or anything else about computer scientists.

To allow exceptional unrestricted extent for decreasing answers, but not unskewed answers, the structural theory would have to be augmented with a special mechanism for strategically bypassing negation-induced symmetry, and this mechanism would have to be sensitive to the semantics of the increasing conjunct (*most syntacticians*). In other words, the structural theory would need some way to simulate the model-theoretic property of skew in the object-language syntax. Even if this were possible, it would be a pointless exercise. The theory I have proposed derives the sensitivity to skew without a syntactic middleman.

3.3.3 Summary

Negative answers usually have restricted extent, but can occasionally be interpreted with unrestricted extent in particular circumstances (van Rooij and Schulz, 2004; von Stechow and Zimmermann, 1984). I proposed to account for this observation by adding a third order, the DUAL ORDER, to the inventory of orders exhaustivity can use. The dual order inverts the default order and maximizes the question-property's extension. However, the dual is the most marked order, so it is only rarely used.

The inventory of orders is now logically "complete": exhaustivity can use orders with positive, negative, or mixed polarity. The key to the analysis is that order selection does not simply pair answers with orders that match their polarity, but interacts with a markedness hierarchy that ranks orders by how negative they are. In this way, the resolution of the symmetry problem is directly tied to the markedness of negation.

3.4 Conclusion

This chapter developed a predictive theory of order selection. Answers were distinguished by skew: positively skewed answers use the default order, while decreasing and unskewed answers use the hybrid order. The sensitivity of exhaustive interpretation to skew follows from how two pragmatic principles interact: REDUCE MARKEDNESS, which drives a preference for a less marked order, and PRESERVE RESTRICTION, which requires an answer's exhaustive interpretation to share a restriction with its literal interpretation. I also showed how the rare, but attested cases of negative answers with unrestricted extent can be accounted for using the DUAL ORDER.

The markedness hierarchy over the default, hybrid, and dual orders plays an important role in order selection. Speakers using this hierarchy have a preference for orders that include as little negation—i.e., order reversal—as possible. The hybrid order is partially positive and partially negative, and is for this reason more marked than the default. The dual order is completely negative, and is for this reason more marked than the hybrid.

The main factor that motivates speakers to use a more marked order is the principle of PRESERVE RESTRICTION. This information preservation principle requires the final output of exhaustivity to still be "about" the same objects as the input meaning, and reflects the general conversational presumption that the literal semantic content of the utterances speakers make is not arbitrary. With this principle in hand, it is possible to derive the direction and extent of exhaustivity.

Conclusion

This thesis has proposed an analysis of how an answer's meaning affects its exhaustive interpretation. It is worth reviewing the components of the system in one place. They are as follows.

(42) **Exhaustivity**.
$\mathtt{exh} = \lambda T_{\tau t}\lambda P_{s\tau}.T(P_w) \wedge \neg\exists v : T(P_v) \wedge v < w$

(43) **The default order.**
$v <_{\mathtt{default}} w$ iff $P_v \cap D \subset P_w \cap D$

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

(45) **The dual order**.
$v <_{\mathtt{dual}} w$ iff $P_v \cap D \supset P_w \cap D$

(46) **Restriction operators (extensional).**
a. $\mathbf{R} = \lambda T_{\langle\tau t\rangle t}.\bigcap\{S|\forall P : T(P) \leftrightarrow T(P \cap S)\}$
b. $\mathbf{R}^{+} = \lambda T_{\langle\tau t\rangle t}\lambda x_{\tau}.\exists S : \neg T(S) \wedge T(S \cup \{x\})$
c. $\mathbf{R}^{-} = \lambda T_{\langle\tau t\rangle t}\lambda x_{\tau}.\exists S : \neg T(S) \wedge T(S - \{x\})$

(47) **Restriction operators (intensional).**
a. ${}^{\uparrow}\mathbf{R} = \lambda T_{\langle s\langle\tau t\rangle\rangle t}.\bigcap\{S_w|\forall P : T(P) \leftrightarrow T(\lambda v\lambda x.P_v(x) \wedge S_v(x))\}$
b. ${}^{\uparrow}\mathbf{R}^{+} = \lambda T_{\langle s\langle\tau t\rangle\rangle t}\lambda x_{\tau}.\exists P, S : \neg T(S) \wedge T(P) \wedge P_w = S_w \cup \{x\}$
c. ${}^{\uparrow}\mathbf{R}^{-} = \lambda T_{\langle s\langle\tau t\rangle\rangle t}\lambda x_{\tau}.\exists P, S : \neg T(S) \wedge T(P) \wedge P_w = S_w - \{x\}$

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

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

(50) **The markedness hierarchy.**
$\mathtt{default} \succ \mathtt{hybrid} \succ \mathtt{dual}$

Exhaustivity is implemented as the assumption that a particular answer T, analyzed as a generalized quantifier meaning at the level of information structure, provides the "best" answer to a question about the property P (Groenendijk and Stokhof, 1984; Schulz and van Rooij, 2006; Spector, 2007; van Rooij and Schulz, 2004). An answer is "best" if it contains the actual world and if the actual world is one of the "minimal" worlds according to an order $<$ over possible worlds. This order is a free parameter that can be resolved to the DEFAULT ORDER, HYBRID ORDER, or DUAL ORDER, depending on the answer.

The default order ranks worlds according to the size of P's extension: a world v is more default-minimal than w iff P's extension in v is a subset of P's extension in w. Application of this order will minimize P as much as possible given the truth of the answer. The dual order is just the opposite: v is more dual-minimal than w iff P's extension in v is a superset of P's extension in w. Application of this order will maximize P as much as possible given the truth of the answer.

The hybrid order simultaneously minimizes P's intersection with an answer's positive restriction and maximizes P's intersection with an answer's negative restriction. Because the integration of these two reasoning schemata has the potential to conflict, the hybrid order is defined in a way that can resolve conflicting information. The hybrid order only regulates P's intersection with the positive and negative restrictions of an answer, and cannot license conclusions about objects outside an answer's restriction. The basic explanation for restricted extent with the hybrid order is that it is impossible to assign objects outside an answer's restriction a polarity relative to the meaning of the answer.

The proposal depends on the formal notion of quantificational restriction, which is defined as the intersection of all the sets that a quantifier lives on (Barwise and Cooper, 1981; von Stechow and Zimmermann, 1984). A quantifier's restriction can be further sorted into a positive restriction and a negative restriction whose union makes up the overall restriction. As usual in the logic of generalized quantifiers, I assume that all GQs are themselves extensional (Keenan and Stavi, 1986) but that there are GQs over the extensions and intensions of properties. In order to treat both, I have defined extensional and intensional versions of the restriction operators.

The resolution of the $<$ parameter for exhaustivity is guided by two factors: a markedness hierarchy over the inventory of orders, and the principle of PRESERVE RESTRICTION, which requires an answer's exhaustive interpretation to be restricted by the same objects as its literal interpretation. This information preservation requirement can be seen as ensuring that an answer's "subject matter" plays a role in its pragmatic interpretation; it ensures that the objects an answer is "about" are not arbitrary. The $<$ parameter is then preferentially resolved to the least marked order compatible with PRESERVE RESTRICTION, according to the following markedness hierarchy: the default order is less marked than the hybrid order, which is in turn less marked than the dual order.

I admit that the proposal is somewhat more technical than the norm for work in this area. It is worth keeping in mind what the proposal does not include: there are no lexical scales and no formal alternatives, and there is no appeal to the representational properties of linguistic objects other than question and answer. For this reason, the proposal surpasses theories that use formal alternatives in both descriptive and explanatory adequacy.

The narrow goal of the thesis was to reach a deeper understanding of direction and extent. The directional properties of answers follow from PRESERVE RESTRICTION. The symmetry problem is the problem of explaining why, in an exchange like *A: Who attended? B: Agnes*, B's answer implicates that nobody else attended instead of implicating that everyone else attended. The solution I have proposed is that only the former option respects the informational content of the answer.

The extent properties of answers follow from the inventory of available orders exhaustivity can use. In particular, the restricted extent of negative answers follows from the fact that they use the hybrid order, which enforces restricted extent in virtue of regulating positive and negative restriction simultaneously. The hybrid order does not quantify over objects outside an answer's restriction because there is no way to assign those objects a polarity relative to the answer's meaning. Negative answers sometimes have unrestricted extent when the dual order is used, but strongly prefer restricted extent. This preference follows from the markedness of negation.

The empirical observations targeted by the proposal are also correctly predicted by Spector, 2007, whose work provides an important precedent. However, Spector's analysis hinges on stipulations about formal alternatives in a logical language and does not clearly explain why the facts hold. The proposal developed here is semantically transparent and grounded in independent pragmatic principles.

The broader goal of the thesis was to achieve these results without using formal alternatives. Mainstream theories in this area, like neo-Gricean pragmatics (Gazdar, 1979) or the structural theory (Katzir, 2007), account for exhaustivity by applying sentential negation to alternative sentences. In light of the fact that exhaustive inferences are nondetachable, the identification of alternatives with sentences cannot be maintained. Regardless of whether exhaustivity is semantic or pragmatic, it must be predictable on the basis of content alone. This goal is achievable.

I have deliberately stayed neutral on the question of whether exhaustivity is the result of pragmatic Quantity reasoning or a linguistic convention derived in compositional semantics. Both views are compatible with the proposal.

To adequately characterize the truth-conditions of a sentence like *Every cat purred*, it is not usually considered necessary to refer metalinguistically to other sentences of the object language. In principle, we could analyze this sentence syncategorematically, saying that it is true if every English sentence of the form *Q purred* is true, where *Q* is the name of a cat. However, most scholars would probably find such an analysis bizarre given that it is possible, and preferable, to give quantified sentences a direct interpretation.

The situation with exhaustivity is exactly parallel. This thesis has shown that exhaustivity can be analyzed strictly in terms of quantification over model-theoretic objects. Semantic/grammatical theories should welcome this finding, since it means that exhaustivity can be analyzed like any other sort of natural language quantification. I have proposed an analysis of exhaustivity as an operation on quantifiers. Proponents of silent object-language exhaustivity operators who find my arguments convincing should feel free to view my proposal as targeting the semantics of such operators.

The assumption that exhaustivity differs from ordinary quantification has allowed the literature to ground analyses of polarity sensitivity, free choice, presupposition, homogeneity, and many other things in the mechanics of exhaustivity, together with representational stipulations about formal alternatives. If I am right about how exhaustivity works, this strategy of analysis is on shaky ground and should be reassessed. In some cases, the key ingredients of a particular analysis may be preserved by appealing to independent semantic mechanisms that maximize informative strength. However, it is also possible that the theory of exhaustivity has been inappropriately overextended to cover phenomena that should be analyzed separately.

From a pragmatic perspective, circumscription should be viewed as a characterization of a more involved reasoning process that uses the orders as heuristics. An informal sketch of how this reasoning could go is as follows.

(51) **A**: Who attended the panel?
B: Not many linguists. $\rightsquigarrow$ *some linguists attended*

A reasons:

- B said "not many linguists". Because B is cooperative, B must have made the strongest utterance she could have *relative to a particular order*.
- Cooperative speakers prefer orders that are less marked. The least marked order is the default order.
- B cannot have intended the default order, because the default order would violate Preserve Restriction and a cooperative speaker wouldn't do that.
- The next least marked order is the hybrid order, which satisfies Preserve Restriction. The supposition that B intended the hybrid order is consistent with the assumption that B is cooperative. So, B must have intended the hybrid order.
- Therefore, B believes that some but not many linguists attended the panel.

As this sketch of A's reasoning process shows, the analysis has interesting implications for a pragmatic view of exhaustivity, since it involves a shift in how Quantity is conceived.

Grice's maxim of Quantity requires speakers to make the most informative utterances they can. If all of logical space matters for assessing this requirement, then Quantity cannot be the source of exhaustive inferences, because of the symmetry problem. That is, if q is a stronger alternative than p, it will still never be possible to conclude $\neg q$ from an utterance of p, because we also have to reason about $p \wedge \neg q$; at most, we can conclude that the speaker does not know whether q holds.

The fact that exhaustive inferences cannot be derived by a Quantity maxim implemented relative to the entire set-theoretic universe—what Fox, 2007 calls the "basic" maxim of Quantity—has often been taken to motivate a supplementary theory of formal alternatives. Katzir, 2007 makes this argument explicitly, and Fox, 2007 makes an even stronger argument. Specifically: since formal alternatives are (by assumption) necessary to derive exhaustive inferences, and since it would be unnatural to restrict a pragmatic principle like Quantity to only those propositions that the grammar of formal alternatives makes available, exhaustivity cannot be pragmatic and must be part of grammar.

However, formal alternatives are not necessary to derive exhaustive inferences, so Fox's argument is not valid. Although the analysis I developed in this thesis is compatible with a grammatical view of exhaustivity, I would like to speculate about another possibility: perhaps the "basic" maxim of Quantity does not exist, and instead speakers reason about the informativity of utterances using the least marked order over possible worlds that preserves an utterance's restriction.

After all, the "basic" maxim of Quantity is not a useful communicative heuristic. It is incapable of licensing conclusions that do not explicitly concern a speaker's knowledge state. Reasoning about informativity relative to a particular direction is much more pragmatically useful, because this reasoning is capable of licensing stronger conclusions. A speech community interested in successful communication could find it valuable to adopt a Quantity maxim parameterized for direction, that is, one characterizable by the description I have given for `exh`.

If we are willing to grant that Quantity reasoning is not about all of logical space, but can be focused in a particular direction, then the symmetry problem is no obstacle to a pragmatic account of exhaustivity.

There are many directions in which the framework developed in this thesis should be expanded, including at least the following:

1. The claims I have made about symmetry breaking should be tested against questions with abstracts over a wider range of semantic objects. The prediction is that symmetry should only be breakable in one direction whenever the question-property is first-order, with greater flexibility observed in higher-order domains.

2. For simplicity, I adopted an extensional characterization of circumscription, but the more nuanced epistemic analysis of van Rooij and Schulz, 2004 is necessary to account for answers that contain modals and other intensional operators. The account of order selection I have proposed should be integrated into this more complex circumscription recipe and tested against the exhaustive implications of negative and mixed intensional answers, including more complex conditional answers than the ones I addressed here.

3. I argued that many apparent exceptions to nondetachability should be accounted for in terms of granularity: although direction and extent are predictable on the sole basis of content, the granularity level at which background property extensions are distinguished is sensitive to an answer's linguistic form. If this conjecture is on the right track, the present proposal should be supplemented with an explicit pragmatic theory of granularity regulation.

4. Although there is evidence that the dual order is sometimes used, its exact circumstances of application should be clarified and the notion of "extra pragmatic motivation" that conditions its use given a precise formal characterization.

Appendix A

Facts about generalized quantifiers

This appendix explores the formal properties of the restriction operators, the concept of skew, and the relation between skew and monotonicity. I assume a finite model.

A1 **Definition: live-on sets ($\mathbf{L}$).**
$\mathbf{L} := \lambda T_{\langle\tau t\rangle t}\lambda S_{\tau t}.\forall P : T(P) \leftrightarrow T(P \cap S)$

A2 **Proposition**: every quantifier lives on at least one set.

Proof. Let D be the set of all objects. Because every $P = P \cap D$, we have $\forall P : T(P) \leftrightarrow T(P \cap D)$ for any generalized quantifier T. This proves that there is no T for which $\mathbf{L}(T)$ is empty. $\square$

A3 **Definition: restrictions ($\mathbf{R}$).**
$\mathbf{R} := \lambda T_{\langle\tau t\rangle t}.\bigcap \mathbf{L}(T)$

A4 **Theorem: intersective closure.**
$\mathbf{R}(T) \in \mathbf{L}(T)$ for all T.

Proof. Let T be a generalized quantifier. Because of proposition A2 we know that $|\mathbf{L}(T)| \geq 1$. Then we can place the elements of $\mathbf{L}(T)$ in correspondence with the natural numbers: let $\mathbf{L}(T) = \{S_1, ..., S_n\}$ for $n = |\mathbf{L}(T)|$. The proof is by induction. Because $S_1 \in \mathbf{L}(T)$ it follows that $\forall P : T(P) \leftrightarrow T(P \cap S_1)$. If $n = 1$, we are done.

Otherwise: for the inductive step, suppose that $\forall P : T(P) \leftrightarrow T(P \cap \bigcap\{S_1, ..., S_m\})$ for some $m < n$. Then because $S_{m+1} \in \mathbf{L}(T)$ we have $\forall P : T(P \cap \bigcap\{S_1, ..., S_m\}) \leftrightarrow T(P \cap \bigcap\{S_1, ..., S_m\} \cap S_{m+1})$, therefore also $\forall P : T(P) \leftrightarrow T(P \cap \bigcap\{S_1, ..., S_m\} \cap S_{m+1})$. By associativity of intersection this simplifies to $\forall P : T(P) \leftrightarrow T(P \cap \bigcap\{S_1, ..., S_{m+1}\})$. Because the order over $\mathbf{L}(T)$ was arbitrary, this proves that $\mathbf{L}(T)$ is always closed under intersection. Pick $m + 1 = n = |\mathbf{L}(T)|$ and it additionally follows that $\forall P : T(P) \leftrightarrow T(P \cap \bigcap \mathbf{L}(T)) = \forall P : T(P) \leftrightarrow T(P \cap \mathbf{R}(T))$. Thus $\mathbf{R}(T) \in \mathbf{L}(T)$ for $\mathbf{L}(T)$ of any countable cardinality. $\square$

A5 **Definition: positive restrictions ($\mathbf{R}^+$).**
$\mathbf{R}^+ := \lambda T_{\langle\tau t\rangle t}\lambda x_\tau.\exists S : T(S) \wedge \neg T(S - \{x\})$

A6 **Definition: negative restrictions ($\mathbf{R}^-$).**
$\mathbf{R}^- := \lambda T_{\langle\tau t\rangle t}\lambda x_\tau.\exists S : T(S) \wedge \neg T(S \cup \{x\})$

A7 **Theorem: dualism.**
$\mathbf{R}^+(T) \cup \mathbf{R}^-(T) = \mathbf{R}(T)$ for all T.

Proof. The first step is to show that $\mathbf{R}^+(T) \cup \mathbf{R}^-(T) \subseteq \mathbf{R}(T)$ for any T. Let T be a generalized quantifier and x an object such that $\mathbf{R}^+(T)(x)$. Then there is an S such that $T(S)$ and $\neg T(S - \{x\})$. By intersective closure, $T(S \cap \mathbf{R}(T))$ and $\neg T(S - \{x\} \cap \mathbf{R}(T))$. Suppose for a contradiction that $x \notin \mathbf{R}(T)$. Then $S \cap \mathbf{R}(T) = S - \{x\} \cap \mathbf{R}(T)$, from which it follows that $T(S \cap \mathbf{R}(T)) \leftrightarrow T(S - \{x\} \cap \mathbf{R}(T))$, a contradiction. Thus $x \in \mathbf{R}(T)$. This proves that $\mathbf{R}^+(T) \subseteq \mathbf{R}(T)$.

Next it can be proven that $\mathbf{R}^-(T) \subseteq \mathbf{R}(T)$ in exactly the same way. Let x be an object such that $\mathbf{R}^-(T)(x)$. Then there is an S such that $T(S)$ and $\neg T(S \cup \{x\})$. By intersective closure, $T(S \cap \mathbf{R}(T))$ and $\neg T(S \cup \{x\} \cap \mathbf{R}(T))$. Suppose for a contradiction that $x \notin \mathbf{R}(T)$. Then $S \cap \mathbf{R}(T) = S \cup \{x\} \cap \mathbf{R}(T)$, from which it follows that $T(S \cap \mathbf{R}(T)) \leftrightarrow T(S \cup \{x\} \cap \mathbf{R}(T))$, a contradiction. Thus $x \in \mathbf{R}(T)$. This proves that $\mathbf{R}^-(T) \subseteq \mathbf{R}(T)$. Because we have $\mathbf{R}^+(T) \subseteq \mathbf{R}(T)$ and $\mathbf{R}^-(T) \subseteq \mathbf{R}(T)$, we must have $\mathbf{R}^+(T) \cup \mathbf{R}^-(T) \subseteq \mathbf{R}(T)$.

To show that $\mathbf{R}(T) \subseteq \mathbf{R}^+(T) \cup \mathbf{R}^-(T)$, let x be an object such that $\mathbf{R}(T)(x)$. Suppose for a contradiction that $\neg\mathbf{R}^+(T)(x)$ and $\neg\mathbf{R}^-(T)(x)$. Then $\neg\exists S : T(S) \land \neg T(S - \{x\})$ and $\neg\exists S : T(S) \land \neg T(S \cup \{x\})$, so we have $\forall S : T(S) \to T(S - \{x\})$ and $\forall S : T(S) \to T(S \cup \{x\})$. This simplifies to $\forall S : T(S) \to T(S - \{x\}) \land T(S \cup \{x\})$. Because either $S = S - \{x\}$ or $S = S \cup \{x\}$, we can strengthen this to a biconditional, $\forall S : T(S) \leftrightarrow T(S - \{x\}) \land T(S \cup \{x\})$.

There are now two cases to consider. First, for sets that contain x we have $\forall S : S(x) \to T(S) \leftrightarrow T(S - \{x\}) \land T(S \cup \{x\})$. Because $S = S \cup \{x\}$ whenever $S(x)$ for any S this reduces to $\forall S : S(x) \to T(S) \leftrightarrow T(S - \{x\})$. Second, for sets that do not contain x, $S = S - \{x\}$ whenever $\neg S(x)$ for any S, so we have $\forall S : \neg S(x) \to T(S) \leftrightarrow T(S - \{x\})$.

Because we have both $\forall S : S(x) \to T(S) \leftrightarrow T(S - \{x\})$ and $\forall S : \neg S(x) \to T(S) \leftrightarrow T(S - \{x\})$, it follows that $\forall S : T(S) \leftrightarrow T(S - \{x\})$. But then there is a member of $\mathbf{L}(T)$ that does not include x, namely $\{y | y \neq x\}$, so $x \notin \mathbf{R}(T)$, a contradiction. Thus if $\mathbf{R}(T)(x)$, either $\mathbf{R}^+(T)(x)$ or $\mathbf{R}^-(T)(x)$ must hold. This proves that $\mathbf{R}(T) \subseteq \mathbf{R}^+(T) \cup \mathbf{R}^-(T)$.

Because we have both $\mathbf{R}^+(T) \cup \mathbf{R}^-(T) \subseteq \mathbf{R}(T)$ and $\mathbf{R}(T) \subseteq \mathbf{R}^+(T) \cup \mathbf{R}^-(T)$, it follows that for any T, $\mathbf{R}^+(T) \cup \mathbf{R}^-(T) = \mathbf{R}(T)$ and we are done. □

A8 **Definition: monotonicity.**
A generalized quantifier T is *increasing* iff $\forall A, B : T(A) \land A \subseteq B \to T(B)$.
A generalized quantifier T is *decreasing* iff $\forall A, B : T(A) \land B \subseteq A \to T(B)$.

A9 **Theorem: upward restriction.**
T is increasing iff $\mathbf{R}^-(T)$ is empty.

Proof. Let T be an increasing quantifier. Because $S \subseteq S \cup \{x\}$ for all S, x, we have $T(S) \to T(S \cup \{x\})$ for all S, x, thus $\mathbf{R}^-(T)$ is empty.

To prove the other direction of the biconditional, let T be a quantifier such that $\mathbf{R}^-(T)$ is empty. Then $\neg\exists x, S : T(S) \land \neg T(S \cup \{x\})$, hence also $\forall x, S : T(S) \to T(S \cup \{x\})$. Let A be a set such that $T(A)$ and let $D - A = \{x_1, ..., x_n\}$ for $n = |D - A|$, D the set of all objects. The proof is by induction. Because $T(A)$ holds and $\mathbf{R}^-(T)$ is empty, it follows that $T(A \cup \{x_1\})$. For the inductive step, suppose that $T(A \cup \{x_1, ..., x_m\})$ for some $m < n$. Then because $\mathbf{R}^-(T)$ is empty we also have $T(A \cup \{x_1, ..., x_m\} \cup \{x_{m+1}\})$ which simplifies to $T(A \cup \{x_1, ..., x_{m+1}\})$ by associativity of union. Because the order over $D - A$ was arbitrary, this proves that T applies to all supersets of A. Because A was arbitrary, T is increasing. □

A10 **Theorem: downward restriction.**
T is decreasing iff $\mathbf{R}^+(T)$ is empty.

Proof. The proof is parallel to the proof of upward restriction. Let T be a decreasing quantifier. Because $S - \{x\} \subseteq S$ for all S, x, we have $T(S) \to T(S - \{x\})$ for all S, x, thus $\mathbf{R}^+(T)$ is empty.

To prove the other direction of the biconditional, let T be a quantifier such that $\mathbf{R}^+(T)$ is empty. Then $\neg\exists x, S : T(S) \wedge \neg T(S - \{x\})$, hence also $\forall x, S : T(S) \to T(S - \{x\})$. Let A be a set such that $T(A)$ and let $A = \{x_1, ..., x_n\}$ for $n = |A|$. The proof is by induction. Because $T(A)$ holds and $\mathbf{R}^+(T)$ is empty, it follows that $T(A - \{x_1\})$. For the inductive step, suppose that $T(A - \{x_1, ..., x_m\})$ for some $m < n$. Then because $\mathbf{R}^+(T)$ is empty we have $T(A - \{x_1, ..., x_m\} - \{x_{m+1}\})$ which simplifies to $T(A - \{x_1, ..., x_{m+1}\})$ by de Morgan's law. Because the order over A was arbitrary, this proves that T applies to all subsets of A. Because A was arbitrary, T is decreasing. □

A11 **Definition: skew.**
A generalized quantifier T is *positively skewed* iff $\mathbf{R}(T) = \mathbf{R}^+(T)$.
A generalized quantifier T is *negatively skewed* iff $\mathbf{R}(T) = \mathbf{R}^-(T)$.

A12 **Proposition:** if T is increasing, T is positively skewed.

Proof. Follows from dualism and upward restriction. □

A13 **Proposition:** if T is decreasing, T is negatively skewed.

Proof. Follows from dualism and downward restriction. □

A14 **Proposition:** if T is positively and negatively skewed, either T is nonmonotonic or T is both increasing and decreasing.

Proof. Let T be a positively skewed and negatively skewed quantifier. Then $\mathbf{R}^+(T) = \mathbf{R}(T) = \mathbf{R}^-(T)$. There are two cases. In case $\mathbf{R}(T)$ is empty, $\mathbf{R}^+(T) = \mathbf{R}(T) = \mathbf{R}^-(T) = \varnothing$. By upward restriction, T is increasing. By downward restriction, T is decreasing.

In case $\mathbf{R}(T)$ is not empty, $\mathbf{R}^+(T) = \mathbf{R}(T) = \mathbf{R}^-(T) \neq \varnothing$. By upward restriction, T is not increasing. By downward restriction, T is not decreasing. $\mathbf{R}(T)$ is either empty or not empty, so these are all the cases and we are done. □

A15 **Proposition:** if T is unskewed, T is nonmonotonic.

Proof. Let T be an unskewed quantifier. Then $\mathbf{R}^+(T) \neq \mathbf{R}(T) \neq \mathbf{R}^-(T)$. Suppose for a contradiction that T is increasing: then by upward restriction, $\mathbf{R}^-(T)$ is empty. By dualism, $\mathbf{R}(T) = \mathbf{R}^+(T)$, a contradiction. Thus T is not increasing.

Suppose for a contradiction that T is decreasing: then by downward restriction, $\mathbf{R}^+(T)$ is empty. By dualism, $\mathbf{R}(T) = \mathbf{R}^-(T)$, a contradiction. Thus T is not decreasing. □

References

Acquaviva, Paolo (2019). “Two studies on the internal syntax of complex names”. In: *Italian Journal of Linguistics* 31.2, pp. 3–36.

Atlas, Jay and Stephen Levinson (1981). “It-clefts, informativeness and logical form”. In: *Radical Pragmatics*. Academic Press.

Bach, Emmon (1986). “The algebra of events”. In: *Linguistics and Philosophy* 9.1, pp. 5–16.

Balogh, Kata (2009). “Theme with variations: a context-based analysis of focus”. PhD thesis. University of Amsterdam.

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

Belnap, Nuel (1975). “How a computer should think”. In: *Proceedings of the Oxford International Symposium on Contemporary Aspects of Philosophy*.

Bergen, Leon, Roger Levy, and Noah Goodman (2016). “Pragmatic reasoning through semantic inference”. In: *Semantics and Pragmatics* 9.20.

Bernard, Timothée and Lucas Champollion (2024). “Negative events and compositional semantics”. In: *Journal of Semantics* 40.4, pp. 585–620.

Birkhoff, Garrett (1940). *Lattice theory*. American Mathematical Society.

Bledin, Justin (2025). “Composing menus”. Ms., Johns Hopkins University.

Block, Eliza (2008). “Is the symmetry problem really a problem?” Ms., NYU.

Breheny, Richard (2008). “A new look at the semantics and pragmatics of numerically quantified noun phrases”. In: *Journal of Semantics* 25, pp. 93–139.

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, Isabelle Dautriche, and Emmanuel Chemla (2018). “Competition and symmetry in an artificial word learning task”. In: *Frontiers in Psychology* 9.2176.

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

Buccola, Brian and Benjamin Spector (2016). “Modified numerals and maximality”. In: *Linguistics and Philosophy* 39, pp. 151–199.

Carlson, Greg (1977). “A unified analysis of the English bare plural”. In: *Linguistics and Philosophy* 1, pp. 413–456.

Carston, Robyn (1998). “Informativeness, relevance and scalar implicature”. In: *Relevance theory: applications and implications*. John Benjamins.

Charlow, Simon (2019). “Scalar implicature and exceptional scope”. Ms., Rutgers University.

Chierchia, Gennaro (2004). “Scalar implicatures, polarity phenomena, and the syntax/ pragmatics interface”. In: *Structures and beyond 3*. Oxford University Press.

Chierchia, Gennaro (2006). "Broaden your views: implicatures of domain widening and the logicality of language". In: *Linguistic Inquiry* 37.4, pp. 535–590.
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.
Cormen, Thomas, Charles Leiserson, Ronald Rivest, and Clifford Stein (1990). *Introduction to algorithms*. MIT Press.
Cresswell, Maxwell and Arnim von Stechow (1982). "De re belief generalized". In: *Linguistics and Philosophy* 5, pp. 503–535.
Cummins, Chris, Uli Sauerland, and Stephanie Solt (2012). "Granularity and scalar implicature in numerical expressions". In: *Linguistics and Philosophy* 35, pp. 135–169.
Curry, Haskell (1980). "Some philosophical aspects of combinatory logic". In: *Proceedings of the Kleene Symposium*.
Davey, B.A. and H.A. Priestley (1990). *Introduction to lattices and order*. Cambridge University Press.
Davidson, Donald (1967). "The logical form of action sentences". In: *The logic of decision and action*. University of Pittsburgh Press.
Davies, Mark (2008). *The Corpus of Contemporary American English (COCA)*.
Dionne, Danielle and Elizabeth Coppock (2022). "Differences in implicature across languages stem from differences in salience of alternatives". In: *Glossa Psycholinguistics* 1.8, pp. 1–27.
Fara, Delia Graff (2000). "Shifting sands: an interest-relative theory of vagueness". In: *Philosophical Topics* 20, pp. 45–81.
Feinmann, Diego (2023). "Ignorance is a problem". In: *Snippets* 45.
— (2025). "On breaking symmetry by complexity". In: *Snippets* 47.
Fine, Kit (2017). "Truthmaker semantics". In: *A companion to the philosophy of language, second edition*. Wiley.
Fitting, Melvin (1991). "Bilattices and the semantics of logic programming". In: *Journal of Logic Programming* 11, pp. 91–116.
Fox, Danny (2007). "Free choice and the theory of scalar implicatures". In: *Presupposition and implicature in compositional semantics*. Palgrave Macmillan.
Fox, Danny and Martin Hackl (2006). "The universal density of measurement". In: *Linguistics and Philosophy* 29, pp. 537–586.
Fox, Danny and Roni Katzir (2011). "On the characterization of alternatives". In: *Natural Language Semantics* 19, pp. 89–107.
Fox, Danny and Benjamin Spector (2018). "Economy and embedded exhaustification". In: *Natural Language Semantics* 26.1, pp. 1–50.
Frank, Michael C. and Noah Goodman (2012). "Predicting pragmatic reasoning in language games". In: *Science* 336.6084, pp. 998–998.
Franke, Michael (2011). "Quantity implicatures, exhaustive interpretation, and rational conversation". In: *Semantics and Pragmatics* 4.1, pp. 1–82.
Franke, Michael and Gerhard Jäger (2016). "Probabilistic pragmatics, or why Bayes' rule is probably important for pragmatics". In: *Zeitschrift für Sprachwissenschaft* 35.1, pp. 3–44.
Frixione, Marcello (2001). "Tractable competence". In: *Minds and Machines* 11, pp. 379–397.
Gallin, Daniel (1975). *Intensional and higher-order modal logic*. North-Holland.
Gazdar, Gerald (1979). *Pragmatics*. Academic Press.
Geis, Michael and Arnold Zwicky (1971). "On invited inferences". In: *Linguistic Inquiry* 2, pp. 561–566.

Geurts, Bart (2006). "Take *five*". In: *Non-definiteness and plurality*. Ed. by Svetlana Vogeleer and Liliane Tasmowski. Benjamins, pp. 311–329.

— (2010). *Quantity implicatures*. Cambridge University Press.

Ginsberg, Matthew (1988). "Multivalued logics: a uniform approach to reasoning in artificial intelligence". In: *Computational Intelligence* 4, pp. 265–316.

— (1990). "Bilattices and modal operators". In: *Journal of Logic and Computation* 1.1, pp. 41–69.

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

Givón, Talmy (1978). "Negation in language: pragmatics, function, ontology". In: *Syntax and semantics 9: pragmatics*. Academic Press.

Greenberg, Joseph (1966). *Language universals, with special reference to feature hierarchies*. Mouton.

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

Groenendijk, Jeroen and Martin Stokhof (1982). "Semantic analysis of wh-complements". In: *Linguistics and Philosophy* 5, pp. 175–233.

— (1984). "Studies on the semantics of questions and the pragmatics of answers". PhD thesis. University of Amsterdam.

Halpern, Joseph and Yoram Moses (1984). "Towards a theory of knowledge and ignorance". In: *Proceedings 1984 non-monotonic reasoning workshop*.

Hamblin, C.L. (1973). "Questions in Montague English". In: *Foundations of Language* 10, pp. 41–53.

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

Hausser, Roland (1983). "The syntax and semantics of English mood". In: *Questions and answers*. Springer.

Hausser, Roland and Dietmar Zaefferer (1979). "Questions and answers in a context-dependent Montague grammar". In: *Formal semantics and pragmatics for natural languages*. Reidel.

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.

— (1984). "Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature". In: *Meaning, form and use in context*. Georgetown University Press.

— (1989). *A natural history of negation*. University of Chicago Press.

— (2000). "From if to iff: conditional perfection as pragmatic strengthening". In: *Journal of Pragmatics* 32, pp. 289–326.

Jacobs, Joachim (1980). "Lexical decomposition in Montague grammar". In: *Theoretical Linguistics* 7, pp. 121–136.

— (1983). *Fokus und Skalen, zur Syntax und Semantik von Gradpartikeln im Deutschen*. Niemeyer.

Jacobson, Pauline (2016). "The short answer: Implications for direct compositionality (and vice versa)". In: *Language* 92.2, pp. 331–375.

Johnson-Laird, Philip (2008). "Mental models and deductive reasoning". In: *Reasoning: studies in human inference and its foundations*. Cambridge University Press.

Karttunen, Lauri (1977). "Syntax and semantics of questions". In: *Linguistics and Philosophy* 1, pp. 3–44.

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

Keenan, Edward and Jonathan Stavi (1986). "A semantic characterization of natural language determiners". In: *Linguistics and Philosophy* 9, pp. 253–326.

Kennedy, Christopher (2007). "Vagueness and grammar: the semantics of relative and absolute gradable adjectives". In: *Linguistics and Philosophy* 30, pp. 1–45.

— (2015). "A "de-Fregean" semantics (and neo-Gricean pragmatics) for modified and unmodified numerals". In: *Semantics and Pragmatics* 8.10, pp. 1–44.

Kisielewicz, Andrzej (1988). "A solution of Dedekind's problem on the number of isotone boolean functions". In: *Journal für die Reine und Angewandte Mathematik* 386, pp. 139–144.

Kratzer, Angelika (1989). "An investigation of the lumps of thought". In: *Linguistics and Philosophy* 12, pp. 607–653.

— (1995). "Stage-level and individual-level predicates". In: *The generic book*. University of Chicago Press.

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

— (1999). "At least some determiners aren't determiners". In: *The semantics/pragmatics interface from different points of view*. Amsterdam: Elsevier, pp. 257–291.

— (2009). "Approximate interpretations of number words: a case for strategic communication". In: *Theory and evidence in semantics*. CSLI Publications.

Kroch, Anthony (1972). "Lexical and inferred meanings for some time adverbs". In: *Quarterly progress reports of the research laboratory of electronics 104*. MIT.

Levinson, Stephen (2000). *Presumptive meanings: the theory of generalized conversational implicature*. MIT Press.

Lewis, David (1973). *Counterfactuals*. Blackwell.

Li, Haoze (2021). "Mandarin wh-conditionals: a dynamic question approach". In: *Natural Language Semantics* 29, pp. 401–451.

Link, Godehard (1983). "The logical analysis of plurals and mass terms: a lattice-theoretical approach". In: *Meaning, use and interpretation of language*. de Gruyter.

Łukasiewicz, Jan (1951). *Aristotle's syllogistic, from the standpoint of modern formal logic*. Oxford University Press.

Marr, David (1982). *Vision: a computational investigation into the human representation and processing of visual information*. W. H. Freeman and Company.

Mascarenhas, Salvador (2021). "A note on the cardinalities of sets of scalar alternatives". In: *Journal of Semantics* 38, pp. 473–482.

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

Matushansky, Ora (2008). "On the linguistic complexity of proper names". In: *Linguistics and Philosophy* 31.5, pp. 573–627.

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

— (1986). "Applications of circumscription to formalizing common sense knowledge". In: *Artificial Intelligence* 28, pp. 89–116.

Merchant, Jason (2004). "Fragments and ellipsis". In: *Linguistics and Philosophy* 27.6, pp. 661–738.

Muskens, Reinhard (1989). "Going partial in Montague grammar". In: *Semantics and contextual expression*. Foris.

Nelken, Rani and Nissim Francez (2002). "Bilattices and the semantics of natural language questions". In: *Linguistics and Philosophy* 25, pp. 37–64.

Parikh, Prashant (2001). *The use of language*. CSLI Publications.

Parsons, Terence (1990). *Events in the semantics of English*. MIT Press.

Penka, Doris (2011). *Negative indefinites*. Oxford University Press.

Quackenbush, Robert (1986). "Unsolved problems: Dedekind's problem". In: *Order* 2, pp. 415–417.

Reiter, Raymond (1978). "On closed world data bases". In: *Logic and data bases*. Plenum Press.

— (1980). "A logic for default reasoning". In: *Artificial Intelligence* 13, pp. 81–132.

Rett, Jessica (2020). "Manner implicatures and how to spot them". In: *International Review of Pragmatics* 12.1, pp. 44–79.

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.

Rooth, Mats (1985). "Association with focus". PhD thesis. University of Massachusetts Amherst.

— (1992). "A theory of focus interpretation". In: *Natural Language Semantics* 1, pp. 75–116.

— (1996). "Focus". In: *The Handbook of Contemporary Semantic Theory*. Blackwell, pp. 271–297.

Russell, Benjamin (2012). "Probabilistic reasoning and the computation of scalar implicatures". PhD thesis. Brown University.

Sauerland, Uli (2004). "Scalar implicatures in complex sentences". In: *Linguistics and Philosophy* 27, pp. 367–391.

Scharten, Rosemarijn (1997). "Exhaustive interpretation: a discourse-semantic account". PhD thesis. Katholieke Universiteit Nijmegen.

Schönfinkel, Moses (1924). "Über die Bausteine der mathematischen Logik". In: *Mathematische Annalen* 92.3-4, pp. 305–316.

Schöter, Andreas (1996). "Evidential bilattice logic and lexical inference". In: *Journal of Logic, Language, and Information* 5, pp. 65–105.

Schulz, Katrin and Robert van Rooij (2006). "Pragmatic meaning and non-monotonic 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*.

Sevi, Aldo (2005). "Exhaustivity: a semantic account of quantity implicatures". PhD thesis. Tel Aviv University.

Smith, Ryan Walter (2020). "Similative plurality and the nature of alternatives". In: *Semantics and Pragmatics* 13.15, pp. 1–44.

Soames, Scott (1982). "How presuppositions are inherited: a solution to the projection problem". In: *Linguistic Inquiry* 13, pp. 483–545.

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

— (2008). "An unnoticed reading for *wh*-questions: elided answers and weak islands". In: *Linguistic Inquiry* 39.4, pp. 677–686.

Spector, Benjamin (2016). "Comparing exhaustivity operators". In: *Semantics and Pragmatics* 9.11, pp. 1–33.
Swanson, Eric (2010). "Structurally defined alternatives and lexicalizations of XOR". In: *Linguistics and Philosophy* 33.1, pp. 31–36.
— (2017). "Omissive implicature". In: *Philosophical Topics* 45.2, pp. 117–137.
Szabolcsi, Anna (1981a). "The semantics of topic-focus articulation". In: *Formal methods in the study of language pt. 2*. Mathematisch Centrum.
— (1981b). "Compositionality in focus". In: *Folia Linguistica* 15.1-2, pp. 141–161.
— (1983). "Focusing properties, or the trap of first order". In: *Theoretical Linguistics* 10.1-3, pp. 125–146.
Tarski, Alfred (1956). "The concept of truth in formalized languages". In: *Logic, semantics, metamathematics*. Clarendon Press, pp. 152–278.
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 Benthem, Johan (1989). "Semantic parallels in natural language and computation". In: *Logic Colloquium '87*.
van der Auwera, Johan (1997). "Pragmatics in the last quarter century: the case of conditional perfection". In: *Journal of Pragmatics* 27, pp. 261–274.
van der Hoek, Wiebe, Jan Jaspars, and Elias Thijsse (1999). "Persistence and minimality in epistemic logic". In: *Annals of Mathematics and Artificial Intelligence* 27, pp. 25–47.
— (2000). "A general approach to multi-agent minimal knowledge". In: *Proceedings of JELIA 2000*.
van Rooij, Robert and Katrin Schulz (2004). "Exhaustive interpretation of complex sentences". In: *Journal of Logic, Language and Information* 13, pp. 491–519.
— (2007). "*Only*: meaning and implicatures". In: *Questions in dynamic semantics*. Elsevier.
van Tiel, Bob, Emiel van Miltenburg, Natalia Zevakhina, and Bart Geurts (2016). "Scalar diversity". In: *Journal of Semantics* 33.1, pp. 137–175.
von Fintel, Kai (1994). "Restrictions on quantifier domains". PhD thesis. University of Massachusetts Amherst.
— (2001). "Conditional strengthening". Ms., MIT.
von Stechow, Arnim (1991). "Focusing and backgrounding operators". In: *Discourse Particles, Pragmatics and Beyond*. John Benjamins.
von Stechow, Arnim and Thomas Ede Zimmermann (1984). "Term answers and contextual change". In: *Linguistics* 22, pp. 3–40.
Weir, Andrew (2020). "Negative fragment answers". In: *The Oxford Handbook of Negation*. Oxford University Press.
West, Douglas (2001). *Introduction to graph theory*. Prentice Hall.
Westera, Matthijs (2017a). "QUDs, brevity, and the asymmetry of alternatives". In: *Proceedings of the 21st Amsterdam Colloquium*.
— (2017b). "Exhaustivity and intonation: a unified theory". PhD thesis. University of Amsterdam.
Westerståhl, Dag (1984). "Determiners and context sets". In: *Generalized quantifiers in natural language*. Dordrecht: Foris.
— (1988). "Quantifiers in formal and natural languages". In: *Handbook of Philosophical Logic: Volume IV: Topics in the Philosophy of Language*. Dordrecht: Springer Netherlands.
Winter, Yoad (2001). *Flexibility principles in boolean semantics: the interpretation of coordination, plurality, and scope in natural language*. Cambridge: MIT Press.

Xiang, Yimei (2021a). "Higher-order readings of *wh*-questions". In: *Natural Language Semantics* 29, pp. 1–45.
— (2021b). "A hybrid categorial approach to question composition". In: *Linguistics and Philosophy* 44, pp. 587–647.
Zimmermann, Thomas Ede (1985). "Remarks on Groenendijk and Stokhof's theory of indirect questions". In: *Linguistics and Philosophy* 8, pp. 431–448.