STRENGTHENING PREDICATES

Mathieu Paillé

Department of Linguistics
McGill University, Montréal

August 2022

A thesis submitted to McGill University
in partial fulfillment of the requirements of the degree of
Doctor of Philosophy

Abstract

Sentences in natural language are routinely interpreted as stronger than would be expected from the lexical meanings of the overt lexical items alone. This has led to the postulation of exhaustification (strengthening) mechanisms in pragmatics and semantics. Such exhaustivity effects have largely been discussed for logical vocabulary, focused expressions, and predicates forming entailment scales with other predicates. Relying on recent work on additive particles, I argue that exhaustivity is at play in a significantly broader array of meanings than previously appreciated: all predicates are exhaustified, in all sentences. That is, the intuited meanings of predicates in sentences are stronger than their lexical–conceptual meanings. I focus on 'taxonomic' predicates, which do not form entailment scales with other predicates. I make this case first and foremost based on apparently banal contradictions like *This comedy is a tragedy* or *The white flag is green*. While these contradictions are intuitively due to the meanings of the predicates, the interaction of these predicates with additive particles (*This comedy is also a tragedy*) and conjunction (*This play is both a comedy and a tragedy*) is argued to show that the predicates are underlyingly consistent. As such, the contradiction observed in the basic case must result from exhaustification.

In addition to demonstrating the existence of exhaustification in the meaning of taxonomic predicates, I also show that this exhaustification behaves in a hitherto undescribed way. The exhaustification of a given predicate is not only obligatory, but it is also obligatorily local to the predicate. Modelling exhaustification through an Exh(aust) operator, roughly equivalent to a covert *only*, predicates are claimed to 'control' Exh: they both require its presence and roughly dictate its syntactic locus. These constraints on Exh give its semantic output the flavour of lexical meaning. I argue that the locality requirement on Exh is best understood as it needing to be in the predicate's maximal projection, and I model this by postulating an Agree relation between derivational morphemes (n^0, a^0, etc.) and Exh.

For Exh to exhaustify predicates in a non-trivial way, predicates must come with alternatives; similarly to expressions like *some* or *or*, they bear alternatives even without being focused. I make two claims about alternatives. First, concerning the alternatives borne by predicates, I suggest as a first approximation that these are the sisters of the predicate in a given conceptual taxonomy. I then propose a notion of 'predicational jurisdiction'—the kind of information provided by a predicate—to suggest that predicates are alternatives iff they share a jurisdiction. For example, *green* and *table* are not interpreted as mutually exclusive (i.e., are not alternatives for controlled exhaustivity) because they contribute different kinds of information; but *table* and *chair*, *comedy* and *tragedy*, and *green* and *white* are alternatives because they share a jurisdiction. This both explains why taxonomic sisters are alternatives, and, as I will show, manages to capture a broader range of data. The second claim about alternatives pertains to how Exh and additive particles interact. One of the key datapoints motivating the view that taxonomic predicates undergo exhaustification is

their interaction with additive particles. Building on work suggesting that additives serve to avoid unwanted exhaustivity effects, I suggest that additives are directly involved in pruning alternatives from the domain of Exh. They do not prevent exhaustification by removing Exh, but can weaken Exh by making it exclude fewer alternatives.

The claim that there is a systematic and principled mismatch between the lexical–conceptual meaning of taxonomic vocabulary items and the meaning intuited from these expressions in actual sentences challenges what appears to be a tacit consensus in linguistics, psychology, cognitive science, and philosophy. Work on concepts takes for granted that the nature of concepts can be researched from the meanings of predicates in natural-language sentences. This thesis shows that this is not straightforwardly the case, because grammar systematically interferes with the basic meanings of predicates.

Résumé

Les phrases dans le langage naturel sont régulièrement interprétées comme étant plus fortes que l'on ne s'y attendrait en considérant uniquement le sens lexical des éléments lexicaux prononcés. Cela a mené à la postulation de mécanismes d'exhaustification (renforcement) en pragmatique et en sémantique. On a surtout discuté de tels effets d'exhaustivité pour le vocabulaire logique, les expressions focalisées, et les prédicats formant des échelles d'implication avec d'autres prédicats. Faisant fond sur de la recherche sur les particules additives, j'argumente que l'exhaustivité est à l'œuvre dans une gamme d'effets sémantiques bien plus large qu'apprécié jusqu'à présent : tous les prédicats sont exhaustifiés, dans toutes les phrases. En autres mots, le sens intuité des prédicats dans les phrases est plus fort que leur sens lexico-conceptuel. Je me concentre sur les prédicats « taxinomiques », qui ne forment pas d'échelle d'implication avec d'autres prédicats. Cet argument provient avant tout de contradictions apparemment banales comme *Cette comédie est une tragédie* ou *Le drapeau blanc est vert*. Tandis que ces contradictions sont intuitivement causées par le sens des prédicats, j'argumente que l'interaction de ces prédicats avec les particules additives (*Cette comédie est aussi une tragédie*) et la conjonction (*Cette pièce est et une comédie et une tragédie*) démontre que les sens sous-jacents de ces prédicats sont consistants. Ainsi, les contradictions observées dans les cas de base doivent être le résultat d'exhaustification.

En plus de démontrer l'existence d'exhaustivité dans le sens des prédicats taxinomiques, je démontre aussi que cette exhaustivité se comporte de façon indécrite jusqu'aujourd'hui. En effet, l'exhaustification d'un prédicat quelconque est non seulement obligatoire, mais elle est aussi nécessairement calculée localement au prédicat. J'argumente que l'exhaustivité, modélisée à l'aide d'un opérateur Exh(austivité) à peu près équivalent à un *seulement* non-prononcé, est « contrôlée » par les prédicats : ceux-ci requièrent la présence d'Exh et dictent sa position syntaxique approximative. Ces contraintes sur Exh donnent à son résultat une saveur lexicale. La contrainte de localité d'Exh consiste en son apparition obligatoire dans la projection maximale du prédicat, ce que je modélise au travers d'une relation d'Accord entre les morphèmes dérivationnels (n^0, a^0, etc.) et Exh.

Pour qu'Exh exhaustifie les prédicats de façon non triviale, les prédicats doivent avoir des alternatives ; comme des expressions telles que *certain* ou *ou*, ils portent des alternatives même sans être focalisés. Je fais deux suggestions sur les alternatives. Premièrement, en ce qui concerne les alternatives portées par les prédicats, je suggère en première approximation que celles-ci sont les sœurs du prédicat dans une taxinomie conceptuelle quelconque. Or, je propose ensuite une notion de « juridiction prédicationnelle » – la sorte d'information contribuée par un prédicat – afin de suggérer que les prédicats sont des alternatives ssi leur juridiction est la même. Par exemple, *vert* et *table* ne sont pas interprétés comme mutuellement exclusifs (c'est-à-dire que ces prédicats ne sont pas des alternatives pour l'exhaustivité contrôlée) puisqu'ils contribuent de différentes sortes

d'information ; mais *table* et *chaise*, *comédie* et *tragédie*, et *vert* et *blanc* sont des alternatives puisque ces prédicats partagent une juridiction. Cela explique pourquoi les sœurs taxinomiques sont des alternatives, ainsi que d'avoir l'avantage d'illuminer une gamme plus large de données, comme je démontrerai. Ma seconde suggestion sur les alternatives concerne la façon dont Exh et les particules additives interagissent. Après tout, l'interaction entre les prédicats taxinomiques et les particules additives est une des données centrales motivant l'hypothèse que ces prédicats sont exhaustifiés. Prenant comme point de départ des travaux suggérant que les additifs ont la capacité de contourner des effets d'exhaustification indésirables, je suggère que les additifs s'occupent directement de restreindre les alternatives d'Exh. Ils ne préviennent pas l'exhaustification en enlevant Exh de la syntaxe, mais ils peuvent affaiblir Exh en le faisant exclure moins d'alternatives.

La suggestion qu'il existe un écart systématique et réglementé entre le sens lexico-conceptuel du vocabulaire taxinomique et du sens intuité de ces expressions dans les phrases où ils se trouvent met à l'épreuve un consensus apparent tacite dans la linguistique, la psychologie, la science cognitive et la philosophie. Le travail sur les concepts prend pour acquis que la nature des concepts peut être éclairée depuis le sens des prédicats tel qu'observé dans des phrases simples. Or, la présente thèse démontre que ceci n'est pas le cas sans complications : en réalité, la grammaire s'ingère systématiquement dans le sens élémentaire des prédicats.

Acknowledgements

This thesis has benefited from insights and support from a number of people and institutions.

I have been continuously stunned by the dedication, insight, help, encouragement, and mentorship I have received from my co-supervisors, Bernhard Schwarz and Luis Alonso-Ovalle. I cannot thank them enough. It has been humbling to witness their knowledge of the semantics literature and their ability to carefully disentangle various assumptions in my thinking. They have been a role-model for the kind of semanticist I should aspire to be, and have completely surpassed my expectations in how much they have helped me with all components of my professional life as a graduate student. Their encouragement and support have been central in producing research such as the present thesis, and in facing the challenges of life in academia.

Luis and Bernhard both unreservedly agreed to take up this project with me during my second year at McGill when I was still working on my first evaluation paper in syntax. I showed up at their offices with a piece of paper with a few intriguing datapoints about additive particles, but little concrete idea about where the project would go. Many professors might have reacted to this by expecting me to come up with clear reasons why the data were interesting and the sorts of theoretical contributions I was expecting to make; if Luis and Bernhard had done this, this project might never have taken off. Instead, they encouraged me from the start. This project has been the best source of drive I have had academically, and they both have my appreciation for believing in me enough to embark on it based on initially little more than some intriguing datapoints.

I keep trying to write separate paragraphs for Luis and Bernhard, which they both deserve. But everything I can think of saying about the one, I can say for the other—for the time they have put in this project, the impact they have had on it intellectually, and the support they have given me. In short, Bernhard and Luis have been a wonderful part of my life these past years, and this thesis would not have gone far without them.

In addition to Bernhard and Luis, I thank the other members of my committee, Junko Shimoyama and Aron Hirsch. This project has roots in discussion about copular clauses in a syntax seminar given by Junko in winter 2018, and I'm grateful for the discussions we have had about it since. Aron came on this project later, but has had an outsized impact on it; chapter 3 in particular is due largely to discussions with him. He has been incredibly generous with his time and with sharing his ability to think clearly through complex empirical and theoretical puzzles, and his ability to make connections to areas of the semantics literature not related to my own work in immediately obvious ways.

I also wish to thank other colleagues of various types at McGill, including Michael Wagner for challenging various components of my analysis and suggesting some important empirical observations (especially the 'double focus' on colour terms discussed in chapter 2), and Justin Royer and Jacob Hoover for their interest in talking with me about my project. I'm also thankful to

Lisa Travis, Jessica Coon, and Nico Baier for supervising my work in syntax (Jessica also has my thanks for discussing some of the syntax in chapter 6 of the present thesis), as well as graduate students involved in syntax, in particular Henrison Hsieh and again Justin Royer—they all greatly affected the quality of my time at McGill.

Parts of this work have been presented at various places. I'm thankful for the feedback of reviewers for and audiences at WCCFL 38 at UBC, SALT 30 at Cornell, Sinn und Bedeutung 26 at the University of Cologne (and the reviewers for the upcoming Sinn und Bedeutung 27 at Charles University in Prague), and NELS 52 at Rutgers. I'm also thankful to the audiences at the Syntax/Semantics reading group at McGill, the *50 ans de linguistique à l'UQÀM* colloquium, Humboldt University Berlin, the University of Göttingen, and the Workshop on Part–Whole Structures in Natural Language at Masaryk University in Brno (with particular thanks to Viola Schmitt, Nina Haslinger, and Marcin Wągiel for the invitations to the last three, respectively).

I thank Nadine Bade, Viola Schmitt, Uli Sauerland, and Alexander Göbel for meeting with me about various parts of this thesis, Marcin Wągiel and Marcin Morzycki for helpful discussion by email, and Nina Haslinger for very many interesting and fruitful discussions touching on both the content of this thesis and a number of related areas of semantics.

I thank everyone who was involved in some way in securing a post-PhD academic continuation for me: Bernhard Schwarz, Luis Alonso-Ovalle, Lisa Travis, Aron Hirsch, Jessica Coon, Benjamin Spector, Philippe Schlenker, Viola Schmitt, Elizabeth Ritter, Martina Wiltschko, and Dimitrios Skordos.

This research was supported by McGill University through the Faculty of Arts Doctoral Fellowship in Indigenous Cultural Studies, and by the Social Sciences and Humanities Research Council of Canada through a Master's Canada Graduate Scholarship followed by a Vanier CGS. I thank these institutions as well as everyone who contributed to my funding applications; this includes many people who have already been named, as well as George Fulford and Peyton Veitch.

I can't thank Douglas Lebo enough for their presence during the six months I spent writing this thesis. Ça aurait été une expérience complètement différente sans toi.

Contents

Chapter 1

Introduction

1.1 Introduction

Sentences in natural language are routinely interpreted as having stronger meaning than is expected from the lexical meanings of the overt lexical items alone. For example, in simple sentences, disjunction is interpreted as exclusive, and existential quantifiers are interpreted as incompatible with universal meanings:

(1) a. It's raining or windy.
$\rightsquigarrow$ It's not raining and windy.
b. Aisha ate some of the apples.
$\rightsquigarrow$ Aisha did not eat all of the apples.

Neither of the inferences in (1) can be attributed to the lexical meaning of the expressions *or* and *some*. The inferences disappear in downward-entailing (DE) environments, such as the antecedents of conditionals:

(2) a. If it's raining or windy, the cat will play indoors.
$\not\rightsquigarrow$ The cat will not play indoors if it's raining and windy.
b. If Aisha ate some of the apples, I will buy coffee.
$\not\rightsquigarrow$ I will not buy coffee if Aisha ate all of the apples.

This has led to the postulation of strengthening mechanisms in natural language.

In this thesis, I investigate the distribution of such strengthening effects. To understand the 'distribution' of strengthening, I will focus on two main questions:

(3) **Two questions about the distribution of strengthening:**
a. Which expressions are subject to strengthening?
b. In what syntactic loci is such strengthening computed?

Of course, my intent is not to answer either of those questions exhaustively, but simply to contribute toward understanding them. I will make one main claim about each of these questions, both of which come out to meaning that sentences are subject to significantly more strengthening effects than previously believed.

About the first question, I show that strengthening is at play in the meaning of content predicates. While this has already been claimed for predicates that form entailment scales with other predicates (Horn 1972) or for predicates that are contrastively focused, I argue that strengthening is in fact a very general feature of predicates. I build this claim from apparently banal contradictions like the sentences in (4).

(4) a. #This comedy is a tragedy.
b. #The white flag is green.

These contradictions might most simply be thought of as stemming from the lexical or lexical–conceptual meanings of these predicates. Comedies are necessarily not tragedies (the sets denoted by *comedy* and *tragedy* have an empty intersection), and *white* and *green* are lexically contradictory due to their universal quantificational force (*green* means that *all* parts of its argument are green, not just that *some* parts are green; and likewise for *white*). However, this simple approach does not hold up to scrutiny. In particular, the contradictions disappear with various Boolean conjunctive elements, including additive particles and conjunctions:

(5) a. This comedy is **also** a tragedy.
(or: A tragicomedy is a comedy that is #(**also**) a tragedy)
b. The white flag is **also** green.

(6) a. A tragicomedy is a play that is both a comedy **and** a tragedy.
b. The flag is both white **and** green.

I will show that possible alternative analyses trying to maintain that the predicates in (4) are truly underlyingly inconsistent do not hold up. As such, the contradictions in (4) must be the result of the predicates having stronger meanings than their underlying lexical or conceptual meanings. Going back to our two questions, this constitutes a contribution to what would ultimately be the exhaustive answer to question (3a), significantly broadening the array of expressions which are claimed to undergo strengthening.

As for the question in (3b), the strengthening of predicates in fact provides an important research space precisely because it motivates a different kind of answer from other types of strengthening effects. What is special about data like (4) is that the postulated strengthening effect occurs even if it leads to the creation of a sentence-internal contradiction out of consistent lexical material. It is not unheard of to posit that strengthening can lead to contradictions *across sentences*; for example, Bade (2014, 2016) provides an analysis of data like (7) claiming that the additive *too* is required because each sentence in B's answer would otherwise be strengthened to mean that *only* Aisha/Ben sang (more on this shortly below).

(7) A: Who sang at the party?
B: Aisha sang. Ben sang #(too).

But strengthening is not usually posited to create contradictions within sentences (but see Chierchia 2013). This is an important distinction. It is one thing to claim that *sentences* are obligatorily strengthened without consideration of other sentences in the discourse; it is quite another to claim that *certain constituents of sentences* are obligatorily strengthened without consideration of other constituents of the same sentence.

Let's unpack this a bit. To claim that strengthening can target particular constituents of a sentence to the exclusion of other constituents, we need to think of it as taking place at some point in the semantic derivation, rather than post-semantically. If strengthening was a pragmatic phenomenon due to the interpretation of sentences qua speech acts, it would take the meaning of entire utterances into consideration; syntactic constituency does not have theoretical status in speech acts. As such, to model that certain constituents can be strengthened without consideration of others, we need to adopt the semantic or 'grammatical' theory of strengthening of Chierchia et al. (2012) (and others), hence providing a new kind of argument in favour of this theory.[1] This semantic theory posits that at least some strengthening effects are the result of an Exh(aust) operator. As a syntactically present operator, it can be embedded so as to take scope over only certain parts of a sentence. Hence, it can strengthen a particular syntactic constituent without taking into consideration the meaning of other constituents in the sentence. We return to this in more depth in section 1.2.

However, even on the semantic theory of exhaustivity, at least as spelled out by Chierchia et al. (2012), nothing ever forces Exh to only take a particular constituent as its prejacent. The syntax of Exh is 'free'; an expression that triggers alternatives can be exhaustified locally or not. Yet, to capture (4) as an exhaustivity effect, it must be stipulated that Exh is necessarily very local to the predicates causing the contradiction. An Exh taking the entire clause as its prejacent would fail to create a contradiction, because (as we will see) Exh is defined so as not to exclude any alternatives that are entailed by its prejacent. A hypothetical LF like (8a) must be ruled out; only an LF like (8b) can be available, if we are to derive the contradictory meaning through strengthening.

(8) a. Exh_{ALT} [this comedy is a tragedy].
b. This [Exh_{ALT} comedy] is a [Exh_{ALT} tragedy].

My contribution to answering the question (3b), then, is to show that there is an entire class of exhaustivity effects where the syntax of Exh is systematically constrained. Certain alternative-triggering expressions require not only the presence of an Exh operator, but also constrain the syntactic position of this operator.

The rest of this introductory chapter is organized as follows. In section 1.2, I overview pragmatic and semantic theories of strengthening. Then, in section 1.3, I turn to work on additive particles, which shows that when additive particles are obligatory in a discourse, it is to avoid an unwanted exhaustivity effect. This will set the stage for the rest of this thesis, where obligatory additive particles will be used to detect exhaustivity effects in language. Section 1.4 briefly discusses some more general assumptions and conventions, and section 1.5 provides a chapter-by-chapter overview of the thesis.

1.2 Background on strengthening and alternatives

In this section, I first give a brief overview of how linguists have analyzed some effects widely accepted to involve strengthening in language, including a pragmatic theory (section 1.2.1), a semantic theory (section 1.2.2) based on an Exh(aust) operator, and a theory that merges lexical

[1] I will be using the word 'strengthening' theory-neutrally in this thesis, while reserving the term 'exhaustivity' specifically for semantic strengthening.

ambiguity with pragmatics (section 1.2.3). I will focus my discussion on the strengthening of *scalar* items ('scalar implicatures' SIs in neo-Gricean parlance)—expressions that generate alternatives forming an entailment scale. These include {*or*, *and*} and {*some*, *all*}.[2,3] I then finesse the meaning of semantic strengthening in section 1.2.4, discussing its proposed properties of 'innocent' exclusion and inclusion, and of excluding alternatives that are neither stronger nor weaker than its prejacent. Finally, I comment specifically on the kind of expressions that can serve as alternatives to the prejacent of Exh in section 1.2.5. Section 1.2.6 concludes by providing a working definition for Exh.

1.2.1 The pragmatic theory

The pragmatic theory of scalar implicatures takes as a starting point Grice's (1975, 1989) insight that the plain semantic meaning of sentences is enriched in conversation through abductive reasoning ('inference to the best explanation'), and expands from there (e.g., Horn 1972, Gazdar 1979, Levinson 1983, Blutner 2002, 2004, Spector 2003, Sauerland 2004, van Rooij & Schulz 2004, Russell 2006, Geurts 2010). Grice (1975:45) points out that conversations are "characteristically, to some degree at least, cooperative efforts," and posits the following principle:

(9) **The Cooperative Principle:** (Grice 1975:45)
Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.

He breaks down this principle into several maxims that cooperative speakers follow (Grice 1975:45–46):

(10)
a. **The maxim of Quantity:**
   (i) Make your contribution as informative as is required (for the current purposes of the exchange).
   (ii) Do not make your contribution more informative than is required.
b. **The maxim of Quality:** Try to make your contribution one that is true.
   (i) Do not say what you believe to be false.
   (ii) Do not say that for which you lack adequate evidence.
c. **The maxim of Relation:** Be relevant.
d. **The maxim of Manner:** Be perspicuous.
   (i) Avoid obscurity of expression.
   (ii) Avoid ambiguity.
   (iii) Be brief (avoid unnecessary prolixity).
   (iv) Be orderly.
   (v) etc.[4]

On the pragmatic approach to strengthening, literal semantic meaning is strengthened due to ab-

[2]Expressions forming entailment scales are shown as sets, simply to avoid redundancy: there is no need to stipulate scalar ordering when it falls out from the meaning of the expressions.

[3]The latter scale could be expanded as, e.g., {*some*, *many*, *most*, *all*}, but I will focus exclusively on *some* and *all* for simplicity.

[4]He writes: "And one might need others" (Grice 1975:46).

ductive reasoning based on the assumption that speakers adhere to these maxims (Geurts 2010). Listeners ask themselves why speakers spoke as they did, and strengthen the plain meaning of their interlocutor's speech with the resulting inferences. These are a sentence's *implicatures*. Such a theory of strengthening is *pragmatic* rather than *semantic* because implicatures are calculated based on the actual use of sentences in a conversation; as such, a sentence's implicatures are computed from the meaning of the *entire speech acts*, and therefore entire sentences as a unit, rather than particular syntactic constituents of sentences.

In addition to these conversational maxims, an important consideration is which sentences listeners consider as *alternatives* to the speaker's assertion. In order to reason about why the speaker said what they said, and not something else, listeners must consider what that 'something else' could be. I will return to alternatives in section 1.2.5, but for now, I simply note that we need a notion of *scales* providing a set of alternatives to certain expressions. It does not matter for our purposes whether scales are lexically encoded as primitives or simply something that comes out of their members' entailment relations. Assuming the scale $\{\textit{some}, \textit{all}\}$ (see footnote 3), the sentence in (11a), repeated from (1b), has as an alternative the sentence in (11b).

(11) a. Aisha ate some of the apples.
b. Aisha ate all of the apples.

The plain, non-strengthened semantic meaning of (11a) is as in (12a), assuming a simple existential meaning for *some* (12b).

(12) a. $\exists x[x \sqsubseteq \iota y[y \in \max_{\sqsubseteq}(*\mathsf{apple})] \wedge \mathsf{ate}(a,x)]$.
b. $[\![\text{some}]\!] = \lambda P.\lambda Q.\exists x[P(x) \wedge Q(x)]$.

But clearly, this does not capture the intuited meaning of (11a), which is that Aisha ate *only* some of the apples: she did not eat all of them. On the pragmatic approach to strengthening, this holds because the plain meaning of (11a), (12a), is enriched through the negation of the alternative in (11b). The listener infers that (11b) is false because it is more informative than (11a), and as such the speaker would have said it if they could (by the maxims of Quantity and Quality). The speaker must therefore not believe (11b) to be true.

This is consistent with the speaker either believing that (11b) is false, or simply being uncertain about the status of (11b) (Soames 1982:521, Horn 2001[1989]:233–234, Sauerland 2004:382ff, Paillé & Schwarz 2018). As such, there must be an 'epistemic step' (Sauerland 2004) strengthening 'it is not the case that the speaker believes (11b) to be true' to 'the speaker believes (11b) to be false.' The epistemic step yields the negation of (11b), strengthening (11a) as desired.

1.2.2 The semantic theory

The pragmatic theory predicts the computation of strengthening to always take place based on entire sentences, because the phenomenon is due to post-semantic abductive reasoning about speakers' intentions. The last two decades have seen a flurry of research problematizing this view (see Chierchia 2004 for early work on this) and suggesting that strengthening should be computed by a covert operator, called Exh(aust) (Chierchia 2006; Fox & Hackl 2006; Fox 2007; Chierchia et al. 2012; Sauerland 2012; Bar-Lev & Fox 2017; Fox & Spector 2018). Exh is present syntactically and can scope over subconstituents of sentences; as such, it affects the semantic rather than prag-

matic meaning of sentences. This is presumably not meant to *replace* the notion of pragmatic strengthening, but its explanatory power overlaps with many of the effects discussed in the pragmatic strengthening literature.

One kind of evidence that there exists semantic strengthening is the apparent 'strengthening' of a particular constituent taking place despite this leading to a global *weakening* of the sentence. Indeed, scalar implicatures typically reverse in DE environments (e.g., Horn 1972, Fauconnier 1975, Levinson 2000, Chierchia 2004) such as the antecedents of conditionals. In (13), for example, *or* is interpreted inclusively.

(13) If you take salad or dessert, you'll be real full. (Chierchia et al. 2012:2306)

However, as Chierchia et al. (2012) observe, it is in fact possible to intuit exclusive disjunctions in antecedents:

(14) If you take salad or dessert, you pay $20; but if you take both there is a surcharge. (Chierchia et al. 2012:2306)

The exclusiveness of *or* means that the constituent *you take salad or dessert* is stronger than it would have been on the inclusive meaning, but the entire sentence is weaker. This is unexpected from the pragmatic theory of strengthening.

On the other hand, if strengthening comes from a semantic operator Exh (written in LFs as Exh_{ALT}, given that it takes a set of alternatives as its first argument), we can strengthen only a particular constituent regardless of the consequences for the entire sentence. For now, let's simply define Exh as asserting both the truth of its prejacent (the constituent it scopes over) and the falsity of all stronger alternatives. Let's assume, for the time being, that (14)/(15) has the alternatives in (15a) and that Exh excludes all alternatives stronger than its prejacent. With this in place, we obtain the meaning in (15b) for the antecedent and in (15c) for the entire sentence.

(15) If $[\text{Exh}_{\text{ALT}}$ [you take salad or dessert]], you pay $20; but if you take both there is a surcharge.

a. $\text{ALT} = \{$you take salad or dessert, you take salad and dessert$\}$
b. $[\![\text{Exh}_{\text{ALT}}$ [you take salad or dessert]$]\!] = 1$ iff
you take salad or dessert $\wedge$ you do not take salad and dessert.
c. $[\![$If $[\text{Exh}_{\text{ALT}}$ [you take salad or dessert]], you pay 20]\!] = 1$ iff
you pay $20 if you take salad or dessert but not both.

This correctly captures the intuited meaning, something which could not be done from the pragmatic theory as laid out in section 1.2.1.

Another argument from Chierchia et al. (2012) for semantic strengthening comes from the embedded computation of exhaustivity in upward-entailing (UE) environments. Some of their discussion focuses on Hurford's Constraint (16), the observation that in sentences with disjunction, neither disjunct can entail the other (Hurford 1974; Singh 2008b).

(16) The joining of two sentences by *or* is unacceptable if one sentence entails the other; otherwise the use of *or* is acceptable. (Hurford 1974:410)

The constraint is well-motivated from simple contrasts like (17).

(17) a. Aisha is an American or a Russian. (Hurford 1974:410)
b. #Aisha is an American or a Californian.

An apparent counterexample to Hurford's Constraint comes from examples with scalar expressions in one of the disjuncts:

(18) a. Aisha solved the first or the second problem or both.
b. Aisha read some or all of the books.
(Chierchia et al. 2012:2309; they cite Gazdar 1979 for (b))

These data can be reconciled with Hurford's Constraint if they involve an exclusive *or* in the first disjunct of (18a) and the 'some but not all' meaning for *some* in (18b). But a pragmatic strengthening effect, laid out in (19) for ease of presentation with an Exh at the root of the sentence, would fail to derive this:[5]

(19) Exh_{ALT} [Aisha read some or all of the books].

a. $\text{ALT} = \left\{ \begin{array}{c} \text{Aisha read some or all of the books,} \\ \text{Aisha read all or all of the books } (\equiv \text{Aisha read all of the books}), \\ \text{Aisha read some and all of the books } (\equiv \text{Aisha read all of the books}) \end{array} \right\}$

b. $[\![(19)]\!] = 1$ iff Aisha read some or all of the books $\wedge \neg$[Aisha read all of the books].
$\equiv$ Aisha read some but not all of the books.

This result not only fails to make the data compatible with Hurford's Constraint, but in fact yields a meaning that is entirely alien to the intuited meaning of the sentence.[6] In contrast, Chierchia et al. (2012) suggest to strengthen the first disjunct of the sentences in (18) through an embedded Exh taking only that disjunct as its prejacent (in (20), I strike-through elided material, and ignore for simplicity the strengthening of *or* to being exclusive):

(20) [Exh_{ALT} [Aisha read some ~~of the books~~]] or ~~Aisha read~~ all of the books.

a. $\text{ALT} = \{$Aisha read some of the books, Aisha read all of the books$\}$
b. $[\![\text{Exh}_{\text{ALT}}$ [Aisha read some ~~of the books~~]$]\!] = 1$ iff
Aisha read some of the books $\wedge \neg$[Aisha read all of the books].
c. $[\![(20)]\!] = 1$ iff Aisha read some but not all of the books, or Aisha read all of the books.

As can be seen in (20), this embedded Exh scoping above only the first disjunct means that (18b) can in fact be reconciled with Hurford's Constraint. Assuming the validity of the constraint, this constitutes an argument in favour of embedded, and therefore necessarily semantic, exhaustification.

1.2.3 No lexical–pragmatic alternative to the grammatical theory

Sauerland (2012) asks whether the data that the semantic theory is meant to cover might alternatively be explained by augmenting the pragmatic theory of SIs with a lexical ambiguity among

[5] In (19), the alternatives are inspired by Sauerland (2004). I have left out single-disjunct alternatives (see section 1.2.4) for simplicity; nothing hinges on this.

[6] This means that something must block the LF in (19); presumably it is precisely the fact that Exh's prejacent violates Hurford's Constraint.

weak scalar expressions. On this alternative view, non-maximal scalar terms would be ambiguous between their weak meanings and a stronger meaning corresponding to a lexicalization of an implicature. In other words, *some* would be ambiguous between its existential meaning and a stronger 'some but not all' meaning. Sauerland (2012) suggests that the data brought up by Chierchia et al. (2012) as evidence for the grammatical theory could just as well be explained in this way. For instance, the meaning of (21a) (repeated from (13)) would come from the lexically inclusive version of *or*, while the meaning of (21b) (repeated from (14)) would come from the lexically exclusive version of *or*.

(21) a. If you take salad or dessert, you'll be real full.
b. If you take salad or dessert, you pay $20; but if you take both there is a surcharge.

It would not do to *only* use lexical ambiguity to capture SIs; this would mean that implicatures could only be local or not be at all. Sometimes, however, what is needed is a global implicature together with a weak lexical meaning for a non-maximal scalar expression. Consider (22):

(22) Aisha doesn't like all of Beethoven's symphonies. (Sauerland 2012:41)
$\rightsquigarrow$ Aisha likes some of Beethoven's symphonies.

On the pragmatic view, the 'indirect implicature' in (22) that Aisha *does* like *some* of Beethoven's symphonies arises from the negation of the alternative *Aisha doesn't like some/any of Beethoven's symphonies* (Sauerland 2012:42). The lexical-ambiguity approach, unless reinforced with the possibility of global pragmatic implicatures, is a non-starter for (22): the inference must be calculated above *not*, so that the double negation yields the positive meaning that Aisha does like some of Beethoven's symphonies. Hence, to deal with the existence of both embedded and globally computed SIs, linguists wishing to reject the semantic approach to strengthening would need both lexical ambiguity and the pragmatic computation of implicatures.

Sauerland (2012) points out that this effort at avoiding semantic strengthening cannot deal with any 'intermediate' computation of SIs—computation which is neither global nor so local as to be possibly captured by the availability of a strong lexical meaning. His examples make use of Hurford's Constraint too. He points to the following example as requiring an intermediate implicature in order not to violate Hurford's Constraint (citing personal communication with Benjamin Spector):

(23) Either Aisha must read at least three of the books or she must read at least four of them.

The plain meaning of the second disjunct entails the plain meaning of the first; this would violate Hurford's constraint, if there was no intermediate strengthening on at least the first disjunct. Strengthening must be computed above *must* but below *or*:

(24) a. Either [Exh_{ALT} [Aisha must read at least three of the books]] or she must read at least four of them.
b. $[\![\text{Exh}_{\text{ALT}}$ [Aisha must read at least three of the books]$]\!] =$

$$1 \text{ iff } \begin{cases} \text{Aisha must read at least three of the books} \wedge \\ \neg[\text{Aisha must read at least four of the books}] \wedge \\ \neg[\text{Aisha must read at least five of the books}] \wedge \\ \qquad\qquad \neg \ldots \end{cases}$$

With the first disjunct exhaustified in this way, the second disjunct no longer entails it, making the sentence compatible with Hurford's Constraint. The presence of intermediate implicatures is predicted by the semantic theory, but not by the pragmatic theory augmented with lexical ambiguity. As Sauerland (2012) concludes, this is a strong argument in favour of the semantic theory.

1.2.4 Non-weaker alternatives, and innocent exclusion and inclusion

We have just concluded that strengthening is best captured through a grammatical operator, Exh, which is defined as asserting both its prejacent and the negation of stronger alternatives. In this section, I tweak Exh's meaning in two ways, and point to a third possible tweak that has been suggested in the literature but which I will remain agnostic about. First, there is evidence that Exh excludes non-weaker alternatives, not just stronger ones; in other words, logically independent alternatives are excluded too. Second, I follow Fox's (2007) proposal to make Exh only exclude alternatives which can all be excluded consistently ('innocently excludable' alternatives). Finally, I outline why Bar-Lev & Fox (2017) and Bar-Lev (2018, 2021) depart more radically from Exh's traditional meaning in also taking it to assert that those alternatives which *cannot* be excluded are in fact *true* (as long as no contradiction arises from this 'inclusion'). Exh does not only exclude innocently excludable alternatives but also *includes* innocently includable ones. I will remain agnostic about this last proposal; it will come up a few times in the thesis, but it will play almost no role in my own argumentation.

Non-weaker alternatives

So far, I have described strengthening as involve the negation (exclusion) of *stronger* alternatives. As pointed out by Chierchia et al. (2012), however, this makes wrong predictions for weak scalar expressions in non-monotonic environments:

(25) Exactly one student solved some of the problems. (Chierchia et al. 2012:2325)

(25) means that exactly one student solved some *but not all* the problems. To obtain the 'not all' meaning for *some*, (25) must involve the exclusion of the alternative in (26).

(26) Exactly one student solved all of the problems.

Together, (25) and the exclusion of (26) mean that exactly one student solved at least one of the problems, and it is not the case that exactly one student solved all of the problems. It cannot be that more than one student solved all of the problems, because if this was the case, it would also be the case that more than exactly one student solved at least some of the problems. Therefore, it must be that exactly one student solved at least some of the problems and no student at all solved all of the problems. The student who solved some of the problems, it follows, did not solve all of them.

But (26) is neither stronger nor weaker than (25). In order for Exh to negate (26), it must be defined so as to exclude not only *stronger* alternatives, but all *non-weaker* alternatives. If so, alternatives which are logically independent from the prejacent of Exh are excluded too.

Innocently excludable alternatives

The second tweak to Exh is to define it as an 'intelligent' operator that is designed not to create contradictions out of non-contradictory lexical material. This is well motivated. To see this, let's first consider what the alternatives for disjunctions are once again. Consider disjunctions under universals, as in either of the examples in (27).

(27) a. Every student went to Winnipeg or Montréal.
b. Aisha must go to Winnipeg or Montréal.

If disjunctions only had one alternative obtained by replacing *or* with *and*, the only inferences we would obtain for the examples in (27) are the following:

(28) a. ¬[Every student went to Winnipeg and Montréal]
b. ¬[Aisha must go to Winnipeg and Montréal]

But this is not enough to capture the intuited meanings of the sentences in (27). Indeed, (27a) does not only convey that not all students went to both cities; it also conveys that not every student went to Winnipeg, and not every student went to Montréal. Likewise, (27b) does not only convey that Aisha does not have to go to both cities; it also conveys that Aisha does not have to go to Winnipeg, and she does not have to go to Montréal (but she does have to go to one of them). The exclusions in (28) do not capture this. If Exh only resulted in the exclusions in (28), it could be that all the students went to Winnipeg, but only half to Montréal; and it could be that Aisha must go to Winnipeg, and can decide whether she goes to Montréal.

The shortcoming of only having an alternative with *and* can be overcome by claiming that each disjunct is itself an alternative, as in (29) for (27a). The hypothesis that disjuncts are alternatives to disjunctions was initially put forward by Sauerland (2004), but due to different empirical concerns.

(29) $\textsc{Alt} = \left\{ \begin{array}{c} \text{Every student went to Winnipeg or Montréal,} \\ \text{Every student went to Winnipeg,} \\ \text{Every student went to Montréal,} \\ \text{Every student went to Winnipeg and Montréal} \end{array} \right\}$

Exhaustifying (27a) with the alternatives in (29) excludes that every student went to Winnipeg and that every student went to Montréal, as desired.

On the other hand, if disjunctions have each disjunct as an alternative, a puzzle emerges for simpler sentences like (30), where the disjunction is not embedded under a universal.

(30) Aisha went to Winnipeg or Montréal.

Since each disjunct is stronger than the assertion, one would expect Exh to negate each disjunct. But the negation of both disjuncts contradicts Exh's prejacent, according to which at least one of the disjuncts must be true.

(31) $[\![\text{Exh}_{\text{ALT}}\ [\text{Aisha went to Winnipeg or Montréal}]]\!]$

$$= 1 \text{ iff } \begin{cases} \text{A. went to Winnipeg or Montréal} \wedge \\ \quad \neg[\text{A. went to Winnipeg}] \wedge \\ \quad \neg[\text{A. went to Montréal}] \wedge \\ \neg[\text{A. went to Winnipeg and Montréal}] \end{cases} \Rightarrow \text{contradiction}$$

To avoid each disjunct being negated in sentences like (30) while still giving rise to excludable alternatives in cases like (27), Fox (2007:§6.1), reworking a pragmatic proposal from Sauerland (2004), suggests to define Exh so as to make it avoid contradictions (pace Chierchia 2013). On this view, Exh only excludes alternatives that can be *consistently* negated with one another and with the assertion of the prejacent. This is called 'innocent exclusion'; only 'innocently excludable' alternatives are excluded by Exh.

(32) **Innocent Exlusion procedure**: (Bar-Lev & Fox 2017:99)

a. Take all maximal sets of alternatives that can be negated consistently with the prejacent.

b. Only exclude (i.e., negate) those alternatives that are members in all such sets—the **Innocently Excludable** (= IE) alternatives.

In other words, the subset of alternatives that are innocently excludable must be both consistent and chosen non-arbitrarily. For a disjunction *A or B* (like the simple sentence in (30)), which has the set of alternatives $\{A \textit{ or } B, A, B, A \textit{ and } B\}$, only the alternative *A and B* is innocently excludable. *A* and *B* cannot both be negated consistently with the prejacent, so they are not negated. (30) therefore has the non-contradictory truth conditions in (33).

(33) $[\![(30)]\!] = 1$ iff A. went to Winnipeg or Montréal $\wedge$ A. did not go to Winnipeg and Montréal.

Innocently includable alternatives

The 'innocent exclusion' property of Exh is an important factor in designing it so as to avoid the creation of contradictions, but it is essentially aligned with the general view that strengthening (whether pragmatic or semantic) involves the negation of stronger or even non-weaker alternatives. A more substantial break from the typical meaning of Exh is proposed by Bar-Lev & Fox (2017) and Bar-Lev (2018, 2021), who claim that Exh asserts that non-excluded alternatives are *true*, rather than simply 'leaving them be.' That is, Exh *includes* non-excluded alternatives. To avoid the creation of contradiction, like with exclusion, we use a notion of *innocently* includable alternatives—alternatives whose truth can be asserted consistently.

(34) **Innocent Inclusion procedure**: (Bar-Lev & Fox 2017:102)

a. Take all maximal sets of alternatives that can be asserted consistently with the prejacent and the falsity of all [innocently excludable] alternatives.

b. Only include (i.e., assert) those alternatives that are members in all such sets—the **Innocently Includable** (= II) alternatives.

Bar-Lev & Fox (2017) make the case for this to capture Free Choice effects like (35).

(35) You can play hockey or lacrosse.

With an Exh capable of Innocent Inclusion, Free Choice can be captured simply by having an Exh operator above the disjunction:

(36) a. Exh_{ALT} [You can play hockey or lacrosse].

b. $\text{ALT} = \left\{ \begin{array}{c} \text{You can play hockey or lacrosse,} \\ \text{You can play hockey,} \\ \text{You can play lacrosse,} \\ \text{You can play hockey and lacrosse} \end{array} \right\}$

(i) IE alternatives: {You can play hockey and lacrosse}

(ii) II alternatives: $\left\{ \begin{array}{c} \text{You can play hockey or lacrosse,} \\ \text{You can play hockey,} \\ \text{You can play lacrosse} \end{array} \right\}$

c. $[\![(36a)]\!] = 1$ iff you can play hockey $\wedge$ you can play lacrosse $\wedge$ you can't play both.

Bar-Lev & Fox (2017:102 fn. 4) defend the view that exclusion takes place before inclusion from the simple datapoint in (37) (modified slightly):

(37) Some of the boys came.

(37) has the alternative *all of the boys came*. If inclusion took place before exclusion, (37) would end up meaning that all of the boys came. Thus, exclusion must take place first.

Innocent Inclusion will not play a big part in this thesis; I mention it here for the few sections where it will. While I personally find that the idea is appealing, little in my own argumentation will hinge on it, so for the purposes of this thesis I will for the most part remain agnostic.

Interim summary

In this subsection, we have reviewed three components of Exh's meaning. It excludes all non-weaker alternatives rather than only stronger ones; it does not knowingly create contradictions by excluding alternatives whose negations are not consistent with one another or with the prejacent; and (if Bar-Lev & Fox (2017) and Bar-Lev (2018, 2021) are right) it includes alternatives which are not excluded and which can be asserted consistently with the prejacent and the exclusion of the excludable alternatives.

1.2.5 Alternatives and syntactic complexity

The last piece of background on exhaustivity I discuss has to do with the nature of the alternatives that Exh might exclude. I follow Katzir (2007) and Fox & Katzir (2011), who argue that alternatives are restricted syntactically: they can be equally syntactically complex as the prejacent, or less complex, but they cannot be more syntactically complex.

Katzir (2007) builds this argument due to the 'symmetry problem' (a term he ascribed to class notes by Kai von Fintel and Irene Heim). As described by Katzir (2007:673):

> for any ϕ' that is stronger than ϕ, and that we would like to reason about [i.e., consider as an alternative to ϕ], there is another alternative, $\phi'' = \phi \wedge \neg\phi'$, which is also stronger than ϕ, and which would license an inference in the opposite direction. Combined, ϕ' and ϕ'' license only ignorance inferences, contrary to fact.

What this means in the case of the expression *some* is that it would not be strengthened to mean 'not all' if it had both *all* and *some but not all* as alternatives, as in (38b).

(38) a. Exh_{ALT} [Aisha ate some of the apples]. (≈ Katzir 2007:673)

b. $\text{ALT} = \left\{ \begin{array}{c} \text{Aisha ate some of the apples,} \\ \text{Aisha ate all of the apples,} \\ \text{Aisha ate some but not all of the apples} \end{array} \right\}$

There is no innocently excludable alternative in (38b); *Aisha ate all of the apples* and *Aisha ate some but not all of the apples* are both stronger than Exh's prejacent, but they produce inconsistent results if they are both negated. Katzir's solution to the symmetry problem is to claim that in fact, *Aisha ate some but not all of the apples* is not an alternative to Exh's prejacent in (38a). He suggests that it cannot be an alternative because it is syntactically more complex than Exh's prejacent. Katzir claims that alternatives can be created by replacing lexical items with other lexical items, or deleting parts of the syntax, but not adding new structure to the syntax.

There is one important exception to this: Katzir (2007:§5) posits that expressions that are syntactically more complex than Exh's prejacent can be alternatives if they are contextually provided. Specifically, a phrase ϕ' can be an alternative to another more simplex phrase ϕ if ϕ' is a syntactic constituent of the same sentence that ϕ is in. Katzir introduces this to deal with examples like (39):

(39) It was warm yesterday, and it is a little bit more than warm today. (Matsumoto 1995:44)
$\rightsquigarrow$ It was not a little bit more than warm yesterday.

Katzir (2007:687) points out that "more or less any . . . inference can be triggered if the relevant material is already part of the structure," as in (40), which can be read as carrying the inference that it was not sunny with gusts of wind yesterday.

(40) It was warm yesterday and it is warm and sunny with gusts of wind today.
(Katzir 2007:687)

I will follow Katzir (2007) in this thesis in assuming that, unless contextually provided, an expression's alternatives must be no more syntactically complex than the expression. On the other hand, while Katzir assumes that contextually provided alternatives must be provided by the sentence rather than prior discourse (without defending this claim), I will not follow this assumption: complex alternatives can generally be contextually provided, whether by the same sentence or a previous sentence.

1.2.6 Interim conclusion: a working definition for Exh

Taking stock, we have seen in this section that strengthening effects are the result of a semantic Exh operator, present in the syntax. I assume that it excludes non-weaker innocently excludable alternatives, and it might also include innocently includable alternatives. Its alternatives are syntactic objects which can be at most as complex as its prejacent, unless a syntactically more complex phrase is contextually provided.

Putting this together, we can define Exh as in (41), closely following Bar-Lev & Fox (2017:104). The set of alternatives is written as a subscript on Exh. I include Innocent Inclusion in (41) for

thoroughness (as written above, I remain agnostic for this thesis).[7]

(41) a. $[\![\text{Exh}_{\text{ALT}}\ \text{S}]\!]^w = 1$ iff
$[\![\text{S}]\!]^w = 1 \wedge \forall \text{S}' \in \text{IE}(\text{S}, \text{ALT})[[\![\text{S}']\!]^w = 0] \wedge \forall \text{S}'' \in \text{II}(\text{S}, \text{ALT})[[\![\text{S}'']\!]^w = 1]$.

b. $\text{IE}(\text{S}, \text{ALT}) = \bigcap\{\text{ALT}' \subseteq \text{ALT} : \text{ALT}'$ is a maximal subset of ALT, such that
$\{w : [\![\text{S}]\!]^w = 1 \wedge \forall \text{S}' \in \text{ALT}'[[\![\text{S}']\!]^w = 0]\} \neq \emptyset$.

c. $\text{II}(\text{S}, \text{ALT}) = \bigcap\{\text{ALT}'' \subseteq \text{ALT} : \text{ALT}''$ is a maximal subset of ALT, such that
$\{w : [\![\text{S}]\!]^w = 1 \wedge \forall \text{S}'' \in \text{ALT}''[[\![\text{S}'']\!]^w = 1] \wedge \forall \text{S}' \in \text{IE}(\text{S}, \text{ALT})[[\![\text{S}']\!]^w = 0]\} \neq \emptyset$.

While this is the working definition I will use for most of this thesis, I will redefine Exh in chapter 6 according to the trivalent semantics of Bassi et al. (2021). And again, the notion of Innocent Inclusion will play no role in the original arguments I will make in this thesis; I only include it now because it will surface in other linguists' arguments later. The only part of my own argumentation that will hinge on the notion of inclusion will come when I will tentatively adopt Bar-Lev's analysis of plural homogeneity effects in chapter 6.

1.3 Additive particles and their interaction with exhaustivity

We have just built up a theory of strengthening which I will use as the basis for this thesis. We now turn to recent work on additive particles, which has opened a new route to finding strengthening effects in language. Bade (2014, 2016) (cf. Krifka 1998; Sæbø 2004; Aravind & Hackl 2017) discusses additive particles (the focus particles *also* and *too*), and more specifically cases where these are obligatory, as in (42). The observation that additives are sometimes obligatory goes back to Green (1973) and Kaplan (1984).

(42) A: Who sang at the party?
B: Aisha sang. Ben sang #(too).

Bade argues that such 'obligatory additive' effects arise when unwanted exhaustification would take place without the additive. Without the additive, B's answer in (42) would mean that *only* Aisha sang, and *only* Ben sang—a contradiction in discourse. Turning Bade's discussion on its head, we can use obligatory additives to uncover exhaustivity effects: if an additive is necessary in sentence ϕ, it must be there because the version of ϕ without the additive, ϕ', has an Exh operator creating a semantic problem. That is, we can use obligatory additive particles to uncover exhaustivity effects we might not have otherwise noticed. This will be one of the basic tools used in this thesis.

In this section, we simply review arguments about additive particles and what makes them ever be obligatory. I start with the view that additives are necessary to avoid unwanted exhaustivity effects (section 1.3.2), then turn to an alternative approach based on the principle of *Maximize Presupposition* (section 1.3.3). Finally, section 1.3.4 overviews arguments in favour of the first of these approaches over the second. But before heading into these theories, let's make sure we have a bit of ground to stand on by better understanding the lexical meaning of additives.

[7] The fact that Exh asserts its prejacent in (41) actually falls out from the Innocent Inclusion operation (Bar-Lev & Fox 2017:104 fn. 6), but I have written it out separately anyway for clarity.

1.3.1 The meaning of additives

Additive particles are presuppositional focus particles. In the discourse in (43), the second sentence (S2) carries a presupposition that is paraphrasable as 'someone other than Ben went to the party' or perhaps more specifically 'Aisha went to the party.'

(43) Aisha went to the party. Ben also went to the party.

Of my paraphrases, the former is only 'existential' whereas the second is stronger due to being anaphoric. Either meaning for additives' presupposition predicts *also* to be licensed in (43). As stated above, additives are focus particles; in (43), *Ben* is focused and it is *also*'s 'associate.' If *also* is anaphoric, *Aisha (went to the party)* would be *also*'s 'antecedent.'

Kripke (2009[1990]) shows that the presupposition of additives is not just existential: (44) requires a salient proposition of the form *x is having dinner tonight* (where $x \neq$ Aisha). It is not enough for conversational participants to know others are dining too. (44) cannot be uttered out of the blue; for it to be felicitous, the conversational participants must have a particular individual in mind who is having dinner that night in addition to Aisha.

(44) Aisha is also having dinner tonight.

The view that additives are more than merely existential has been widely adopted (Soames 1989, Heim 1992, Zeevat 1992, 2002, Beaver 1997, Asher & Lascarides 1998, van der Sandt & Geurts 2001, Geurts & van der Sandt 2004, Chemla 2008, Bade 2016, Göbel 2019, pace Ruys 2015),[8] although not all proposals that are 'more than existential' are the same in the details.

On the one hand, there are truly anaphoric proposals for additives, which essentially treat them like pronouns. One such proposal comes from Heim (1992), who deals with additives' anaphoricity through indexation. (45) follows the spirit of her proposal in having *also* co-indexed with some proposition $g(i)$ from the alternatives of the prejacent.

(45) $[\![\text{also}_i]\!]^g = \lambda\, \text{ALT}_{\langle st,t\rangle}.\lambda p.\lambda w : g(i) \in \text{ALT} \wedge g(i)(w) \wedge g(i) \neq p.\ p(w).$

Thus, assuming an assignment function g where $[1 \rightarrow \lambda w.\ \mathsf{sing}(a)(w)]$, the truth conditions in (46b) hold for the S2 in (46a). Note that it is required that Aisha sang in order for the S2 to be true, even if only Ben is overtly mentioned in the sentence; as such, with this indexation, S2 entails that Aisha sang.

(46) a. [Aisha sang]$_1$. Ben also$_1$ sang.

b. $[\![\text{also}_1\ [\text{Ben}_F\ \text{sang}]]\!]^g = \begin{cases} 1 \text{ if } \mathsf{sing}(a) \wedge \mathsf{sing}(b), \\ 0 \text{ if } \mathsf{sing}(a) \wedge \neg\mathsf{sing}(b), \\ \quad \# \text{ otherwise} \end{cases}$

Some authors have weaker 'anaphoric-like' presuppositions only requiring the existence of a true alternative proposition in the set of alternatives. This is a sort of 'specific existential' presup-

[8] Karttunen & Peters (1979) are cited by Kripke (2009[1990]:371, fn. 9) as providing an existential analysis of additives; this is true in their formalism (p. 35), but in their prose (p. 33), they write that "'John drinks too' entails that there is someone else *under consideration* other than John who drinks" (my emphasis). That is, the fact that (44) is infelicitous out of the blue is aligned with what Karttunen & Peters (1979) write in their prose. Looking ahead, their prose corresponds to the 'specific existential' lexical entry in (47).

position. For example, Göbel (2019:289) gives (47).

(47) $[\![\text{also}]\!] = \lambda\, \text{ALT}_{\langle st,t\rangle}.\lambda p.\lambda w : \exists q[q \in \text{ALT} \wedge p \not\Rightarrow q \wedge q(w) = 1].\ p(w).$

While not anaphoric per se, this is more than a mere existence presupposition because there must be an identifiable true proposition in the sentence's alternatives. Kripke's example (44) is correctly predicted to be infelicitous if there is not another individual who is an alternative to Aisha and whom the common ground entails to be having dinner that night. The fact that someone somewhere is having dinner in addition to Aisha is not enough to satisfy the presupposition in (47).

I will follow Heim's approach in this thesis, and bring up this 'specific existential' alternative when relevant in chapter 3.

1.3.2 Obligatory additives as avoiding unwanted strengthening effects

The first set of theories about obligatory additives relies on the claim that *Aisha* and *Ben* in discourses like (48), repeated from (42), are *contrastive topics* (see Büring 2016).

(48) A: Who sang at the party?
B: Aisha sang. Ben sang #(too).

Krifka (1998:121) calls each of the two sentences in B's answer a 'contrastive answer,' because each only provides information on a proper subset of the contrastive topics (e.g., *Aisha sang* remains mum on Ben). He posits a 'distinctiveness' constraint on contrastive answers; in (49), T is the contrastive topic and C is the focus of the comment (corresponding to the entire comment, viz. *sang*, in these examples).

(49) **Krifka's distinctiveness condition** (Krifka 1998:122)
If $[\ldots \text{T}_F \ldots \text{C}_F \ldots]$ is a contrastive answer to a question Q, then there is no alternative T′ of T such that the speaker is willing to assert $[\ldots \text{T}' \ldots \text{C} \ldots]$.

That is, if *Aisha* and *Ben* are contrastive topics, *Aisha sang* implies (due to (49)) that there is no other true answer of the form *x sang*. Interestingly, Krifka (1998:122) suggests that this constraint has its roots not from the maxim of Quantity but the maxim of Manner: "if the speaker could assert $[\ldots \text{T}' \ldots \text{C} \ldots]$, the speaker would have asserted it right away by way of conjoining T and T′"; indeed, "the answer $[\ldots \text{T} \wedge \text{T}' \ldots \text{C} \ldots]$ is shorter than the answer $[\ldots \text{T} \ldots \text{C} \ldots] \wedge [\ldots \text{T}' \ldots \text{C} \ldots]$."

Given that *Aisha sang* in (48) is a contrastive answer, the distinctiveness condition ensures that the sentence means that Aisha is the only individual among the contrastive topics such that the speaker is willing to assert that she sang. That is, without the additive, the first sentence in B's answer means that Ben did *not* sing. For this theory to work, we need to claim that additives "allow us to get around the distinctiveness constraint" (Krifka 1998:122). How exactly additives allow the circumvention of unwanted strengthening effects is the focus of my chapter 3, so I put this question aside for the time being.

Sæbø (2004) provides some data that are problematic for the details of Krifka's analysis. In particular, some instances of obligatory additives can be observed where no conjunction ($[\ldots \text{T} \wedge \text{T}' \ldots \text{C} \ldots]$ in Krifka's notation) would have been possible or desirable. One such case is if the first sentence is strictly stronger than the second:

(50) To the north [Swift Deer could see] the yellow-brown desert, a lot belt of green cactus-covered ridges and distant blue mountain ranges with sharp peaks. **To the south #(too) he could see mountains.** (Sæbø 2004:206)

The maxim of Manner would not require *too* here, because 'To the north and to the south lay the yellow brown desert, . . . ' is not a viable alternative (being false, assuming that to the south, there are only mountains). Returning to the empirical content of (48), (51) makes the same point as (50):

(51) Aisha sang badly. Ben #(also) sang.

In (51), *Aisha and Ben sang badly* is not a viable alternative if Ben did not sing badly, so Krifka's approach does not predict the necessity of the additive.

The second case Sæbø (2004) brings up is when some distance in discourse separates the sentence with the obligatory additive and its antecedent, as in (52). In such cases, the antecedent is often not intuited as a contrastive topic. As Sæbø (2004:207) writes, in (52), "the first paragraph does not suggest that we can only form things out of Lego blocks."

(52) So now you see what I meant about Lego blocks. They have more or less the same properties as those which Democritus ascribed to atoms. And that is what makes them so much fun to build with. They are first and foremost indivisible. Then they have different shapes and sizes. They are solid and impermeable. They also have 'hooks' and 'barbs' so that they can be connected to form every conceivable figure. These connections can later be broken so that new figures can be constructed from the same blocks. . . .

We can form things out of clay #(too), but clay cannot be used over and over, because it can be broken up into smaller and smaller pieces. (Sæbø 2004:207)

Since clay was not even under consideration at the point in the discourse where Lego blocks are discussed, the maxim of Manner once again cannot explain the required presence of the additive in (52): we cannot posit an alternative of the form *Now you see what I meant about Lego blocks and clay* without significantly restructuring the discourse.[9]

These concerns can be alleviated by Bade's (2014; 2016) reformulation of these theories as claiming that additives are obligatory to avoid unwanted *semantic exhaustification* rather than Gricean implicatures. On Bade's approach, contrastive topics, being focused, are exhaustified, with other contrastive topics as alternatives. Without the additive, B's answer in (48) has the following LFs and truth conditions (assume Aisha, Ben, and Carrie are the contrastive topics in the discourse):

(53) a. $[\![\text{Exh}_{\text{ALT}}\ [\text{Aisha}_F\ \text{sang}]]\!] = 1$ iff Aisha sang $\wedge$ Ben didn't sing $\wedge$ Carrie didn't sing.
b. $[\![\text{Exh}_{\text{ALT}}\ [\text{Ben}_F\ \text{sang}]]\!] = 1$ iff Ben sang $\wedge$ Aisha didn't sing $\wedge$ Carrie didn't sing.
$\Rightarrow$ contradiction in discourse

[9] Sæbø's (2004:213) solution to these difficulties is to define additives so that they effectively conjoin the topic of their sentence with some alternative. In (50), for example, the sentence with *too* means 'To the south and to the north, he could see mountains.' Additives result in 'aggregate contrastive topics' in their own sentences (Sæbø 2004:214), without affecting prior sentences. I will show in chapter 3 that it is insufficient to understand additive particles as *only* affecting their own sentences; they must also be able to affect prior ones.

Again, we will discuss how additives fix the problem in (53) in chapter 3. Bade simply suggests that additives allow Exh to disappear; let's go with this proposal for now as a placeholder hypothesis until chapter 3.

On the exhaustification approach, Sæbø's examples are unproblematic. Let's first consider (54), repeated from (51), where the issue for the maxim of Manner analysis is that the comment in the additive's antecedent is strictly stronger than the comment in its host sentence.

(54) Aisha sang badly. Ben #(also) sang.

The exhaustification analysis predicts the additive to be obligatory due to the second sentence alone, which has the meaning in (53b) and therefore contradicts the first. The same goes for examples like (52), where only the sentence with the additive involves a contrastive topic. When Lego blocks are first mentioned, they are not a contrastive topic and therefore not focused; if the sentence is exhaustified at all, it would not create entailments about alternatives to the Lego blocks. However, the sentence bringing in clay as an alternative to Lego blocks *does* exhaustify clay as a topic contrasting with (at least) Lego blocks. Without the presence of *also*, this would contradict the preceding discourse:[10]

(55) $[\![\text{We can form things out of clay}_F]\!]$
$= 1$ iff $\begin{cases} \text{we can form things out of clay} \wedge \\ \text{we can't form things out of Lego blocks} \wedge \\ \text{we can't form things out of} \ldots \end{cases}$

1.3.3 Obligatory additives as the result of *Maximize Presupposition*

An alternative approach to obligatory additive effects is the claim that additives are obligatory due to being presuppositional; they are obligatory whenever their presupposition is met. This theory makes use of the principle of *Maximize Presupposition* (e.g., Heim 1991, Sauerland 2008, Chemla 2008, Singh 2011, Schlenker 2012, Bade 2021), henceforth MP:

(56) **Maximize Presupposition**: (Chemla 2008:142)
Among a set of alternatives, use the felicitous sentence with the strongest presupposition.

This principle is observable due to a variety of cases where presuppositionally strong material must be used instead of presuppositionally weaker or non-presuppositional alternatives. Examples of this phenomenon include the obligation to use definite articles over indefinites when their uniqueness presupposition is met (Heim 1991) as in (57), as well as the obligation to use *both* when its 'cardinality of two' presupposition is met (58) and to use *know* (which has a factive presupposition) rather than *believe* when something is known by the speaker to be true (59).

(57) a. The sun is in my eyes. (Heim 1991)
b. #A sun is in my eyes.

(58) a. Both of Aisha's eyes are open. (Singh 2011:150)
b. #All of Aisha's eyes are open.

[10] Note that, on Katzir's (2007) approach to alternatives, while *Lego blocks* is more syntactically complex than *clay*, it can still function as an alternative due to being contextually provided.

(59) SCENARIO: The speaker has a sister. (≈ Chemla 2008:141)
a. Aisha knows that I have a sister.
b. #Aisha believes that I have a sister.

MP involves positing scales of alternatives according to presuppositional strength: {*a*, *the*}, {*all*, *both*}, and {*believe*, *know*}.

Several authors (Amsili & Beyssade 2006, Chemla 2008, Sauerland 2008:590 fn. 2, Singh 2008a, 2011) have explained obligatory additives as stemming from MP. Doing so has the advantage of collapsing the obligatory nature of additives with the obligatory nature of other presupposition triggering material (57)–(59). It also means that there is no puzzle about how additives can fix problems associated with exhaustification (which makes our lives as linguists easier, but is not an argument in favour of the MP approach).

A challenge for the MP approach is the question of what additives' non-presuppositional scalemate is. On this approach, after all, in discourse like (60) (repeated from (42) and (48)), the second sentence in B's answer involves competitions between the two sentences in (61).

(60) A: Who sang at the party?
B: Aisha sang. Ben sang #(too).

(61) a. Ben sang.
b. Ben sang too.

Chemla (2008) simply posits the scale {∅, *too*}. But this concerns Singh (2008a): if *Ben sang too* is said to compete with *Ben sang*, this means that the *Maximize Presupposition* account of obligatory additives must come with a rejection of Katzir's (2007) claim that alternatives are no more syntactically complex than the assertion. If Katzir (2007) is right about alternatives, a speaker should be able to say (61a) without (61b) even being considered as an alternative, therefore failing to trigger any problem due to MP. To ensure that *Ben sang* and *Ben sang too* have the same syntactic complexity, Singh claims that the scale *too* is in is not {∅, *too*}, but rather {∼, *too*}, where ∼ is Rooth's (1992) focus interpretation operator. On the standard approach, there is a ∼ in both the sentence *Ben sang* and *Ben sang too*, so to ensure equal complexity, Singh modifies Rooth's theory so that certain expressions like *too* can interpret focus without ∼.[11]

If Singh's solution to the problem of alternatives is accepted, MP offers a simple way to capture obligatory additive effects, which does not require a theory of how an additive could fix an unwanted exhaustification effect.

1.3.4 Comparing the two theories of obligatory additives

We have seen that additives might either be obligatory when they are required to circumvent an unwanted exhaustivity effect, or because of the MP principle requiring presuppositional expressions to be used instead of non-presuppositional ones when the presupposition is met.

[11] Another way to avoid the problem identified by Singh (2008a) would be to take ∅ in Chemla's scale to refer not to the absence of linguistic material, but to a covert meaningless lexical item with the same syntactic profile as *too*. It is not clear to me whether this alternative really holds up given the rich array of possible syntactic loci for additive particles; moreover, it is generally not conceptually appealing to posit phonologically and semantically null lexical items (e.g., Chomsky 2001).

There are several ways to tell the theories apart, and these favour the unwanted-exhaustivity theory. Bade (2016) starts off by considering how additives behave in negative sentences. On the MP account, negation is not expected to change whether an additive is obligatory, because presuppositions project past negation. Indeed, we can see that under negation, presuppositionally strong expressions are generally still required (62a), while exhaustivity effects are well-known (e.g., Bassi et al. 2021) to disappear under negation (62b).

(62) a. With my sunglasses, I don't have {#a, the} sun in my eyes.
b. Aisha didn't talk to Ben or Carrie. (#She talked to both.) (Bassi et al. 2021)
$\not\rightsquigarrow$ it is not true that Aisha talked to just one of them

Based on work by Österle (2015) and her own experiments (Bade 2016:§3.2.3), Bade (2016) shows that the prediction of the MP account of obligatory additives is not borne out. Additives are no longer obligatory under negation (see also Bade & Tiemann 2016 and Bade & Renans 2021):

(63) a. Aisha sang. Ben #(also) sang.
b. Aisha sang. Ben didn't (also) sing.

Additives also become optional in other DE contexts where exhaustivity is known to optionally (in fact, preferrably) disappear:

(64) Aisha came. (Bade 2021:12)
a. If Ben came (too), the party was fun.
b. Did Ben come (too)?

Since the presupposition of *too* projects in (64), just like the data with negation (63b), the MP account predicts *too* to be obligatory in (64) just like in (63b).

Another test proposed by Bade (2016) is to look at whether additives remain obligatory when two clauses in a discourse are conjoined into a single sentence:

(65) Aisha sang and Ben (also) sang.

She shows experimentally that the additive is no longer required in these conjunctions. This is expected from the unwanted-exhaustivity account: Exh could take global scope (66), from where it would not create a contradiction needing to be fixed by an additive.

(66) $[\![\text{Exh}_{\text{ALT}}\ [\text{Aisha}_F \text{ sang and Ben}_F \text{ sang}]]\!]$
$= 1$ iff Aisha sang $\wedge$ Ben sang $\wedge$ Carrie didn't sing.

On the MP account, however, this is surprising: the additive's presupposition is met in the second conjunct, so it should be obligatory.

As such, I follow Bade's conclusion that additives are obligatory as a way of circumventing unwanted exhaustification effects, rather than due to their being presuppositionally stronger than a putative non-presuppositional alternative. Moreover, new arguments in favour of the unwanted-Exh account of obligatory additives will be provided in chapter 2.

1.4 Basic assumptions in this thesis

The previous discussion sums up the relevant background for the thesis. In this section, I discuss a few more general assumptions about semantics and syntax; later, in section 1.5, I will give a chapter-by-chapter overview of the body of the thesis.

In this thesis, I will generally assume Fregean ('formal') semantics as spelled out by Heim & Kratzer (1998), including their lambda notation and semantic types, and the generative syntax associated with the Minimalist Program (Chomsky 1995b), in particular as spelled out by Chomsky (2001). I will mark examples with * if they are syntactically ill-formed and with # if they suffer from any problem in meaning, whether it is general weirdness, falseness in a given scenario, or (as will often be the case) internal problems in consistency (i.e., sentence-internal contradictions). My formalism will depart from convention in a few specific ways, but these are intended as presentational shortcuts rather than substantial disagreement. Perhaps most saliently, it is standard in formal semantics to overtly intensionalize expressions through world parametres or variables; since intension plays an explicit role in very little of my discussion, I generally leave out world variables, unless it is helpful to give an expression a presupposition (as a condition on worlds) or a particular example requires them for empirical reasons.

In the rest of this section, I discuss the nature of predicates in formal semantics (section 1.4.1), as well as some relevant syntactic assumptions (section 1.4.2).

1.4.1 Predicates in formal semantics

This thesis investigates strengthening in language specifically in the domain of the meaning of predicates. I will largely focus on syntactically simplex, one-place ($\langle et \rangle$) predicates like *comedy* or *red*, with a focus on nouns and adjectives. I will occasionally bring up syntactically complex predicates (VPs and PPs) like (for example) *visits Aisha* (*visits* is a two-place predicate, but *visits Aisha* is already partially saturated) or *on the table*—in particular in chapter 4.

Focusing again on one-place simplex predicates, these have sets of individuals as their extensions (see e.g. Montague 1973; Heim & Kratzer 1998; Coppock & Champollion in progress). *Aisha is a dog* means that $a \in \{x : x \text{ is a dog}\}$, which we can abbreviate as $\text{dog}(a)$. At first approximation, then, the denotation of *dog* is the set of dogs or its characteristic function, as in (67):

(67) $$[\![\text{dog}]\!] = \begin{bmatrix} a & \rightarrow & 1 \\ b & \rightarrow & 0 \\ c & \rightarrow & 1 \end{bmatrix} \equiv \{a, c\}$$

However, Heim & Kratzer (1998) comment rather common-sensically that denotations similar to (67) are inaccurate representations of human linguistic knowledge. Writing on the example of the verbal predicate *smokes*, they find that displaying its meaning in a function akin to (67) "would have required more world knowledge than we happen to have. We do not know of every existing individual whether or not (s)he smokes. And that's certainly not what we have to know in order to know the meaning of 'smoke'" (Heim & Kratzer 1998:21). This is quite right; if the goal of semantics/linguistics is to model the linguistic knowledge of an idealized speaker, (67) is clearly inadequate. The extension of the predicate is not its 'meaning'; one can know the meaning of a predicate without knowing its full extension. How, then, can the meaning of predicates be

captured?

In early approaches to formal semantics, it was noted that defining predicates by their real-world extension leads to a problem for predicates with no extension in the real world. The two predicates in (68) have the same extension, so if all semanticists have to deal with predicates' meanings is their extension, these predicates would mean the same thing (see Goodman 1949, 1953, 1968; Carnap 1955; Rescher 1959).

(68) a. $[\![\text{unicorn}]\!] = \{\ \}$
b. $[\![\text{fairy}]\!] = \{\ \}$

This is, of course, as wrong as it looks. For instance, "we cannot justifiably substitute the term 'unicorn' for the extensionally empty term which occurs in the sentences" in (69) (Rescher 1959:626):

(69) a. A panacea can cure gold.
b. A centaur has the body of a goat.
c. A dragon can breathe fire.

It was proposed that predication should be intensionalized to deal with this issue. Rather than only providing information about real-world extensions, predicates map individuals to the set of worlds in which the predicate is true of them. If there are three animals, a, b, and c, and three worlds, w_1, w_2, and w_3, the denotation of *dog* might be as in (70) (e.g., Carnap 1947, 1955, 1963):

(70) $$[\![\text{dog}]\!] = \begin{bmatrix} a & \rightarrow & \{w_1, w_2\} \\ b & \rightarrow & \{w_2, w_3\} \\ c & \rightarrow & \{w_1\} \end{bmatrix}$$

The issue of *unicorn* and *fairy* is solved because there are imaginary worlds where these exist, and the predicates have different extensions in those worlds.

Intensionalizing predicates solves the issue of predicates with identical extensions in the real world, but it does not provide a theory of the meaning of predicates. It would be easy to overstate the difference that possible worlds make; (70) is just a function determining different *extensions* in possible worlds. This can be brought out by rewriting the function in (70) as in the equivalent (71), which defines a property as a function from worlds to sets of individuals (Gärdenfors 2000:62):

(71) $$[\![\text{dog}]\!] = \begin{bmatrix} w_1 & \rightarrow & \{a, c\} \\ w_2 & \rightarrow & \{a, b\} \\ w_3 & \rightarrow & \{b\} \end{bmatrix}$$

This "shows the correspondence between the extensional and the intensional definition of a property because the value of the function representing a property is a set of objects as in the extensional case" (Gärdenfors 2000:62). Stalnaker (1981) emphasizes this equivalence by calling this theory an 'extensional' view of predication—"extensional in the sense that properties are defined by their extensions in different possible worlds" (Stalnaker 1981:346).

As such, if we are interested in the meaning of predicates, intensionalizing predicates is not a substantial improvement over the extensional approach of early work in formal semantics; it does not provide a theory of their meaning. What we need is a sort of rule determining whether an individual should count as a member of the set denoted by the predicate, without speakers of a

language needing to be pre-equipped with knowledge of the membership of that set. As Stalnaker (1981:347) writes, "What the standard semantics lacks is an account of properties that defines them independently of possible worlds and of individuals. . . . a property must be not just a rule for grouping individuals, but *a feature of individuals in virtue of which they may be grouped*: not just a propositional function, but something that determines a propositional function" (my emphasis; I assume that by "a rule for grouping individuals," he means a pre-defined, lexicalized extensional function). Heim & Kratzer (1998) deal with the problems of a purely extensional approach to predicate meanings by claiming that predicates' meaning is a *condition* that must be met for them to be true of an individual:

(72) $[\![\text{dog}]\!] = [f : D \to [W \to \{0,1\}]$. For all $x \in D$ and $w \in W$, $f(x)(w) = 1$ iff x is a dog in $w]$
(adapted from Heim & Kratzer 1998:19, 20)

The idea is that the predicate *dog* is true of an individual x if the condition that x is a dog holds. Clearly, (72) requires no memorization or lexicalization of who is and isn't a dog in which world. The fact that it specifies the extension of *dog* through a *condition* is what constitutes the meaning of *dog* (see Heim & Kratzer 1998:22 on differences made by the 'mode of presentation' of extension chosen by the semanticist).

At the same time, Heim & Kratzer's proposal is not fully transparent because the condition they posit is entirely reliant on metalanguage ('is a dog') whose status is not clarified. This can be fixed easily by stating explicitly that the meta-language 'dog' involves a link between language and the concept DOG (e.g., Pietroski 2018). (72) is tacitly relying on an unspecified theory of concepts or world knowledge; we can simply make this explicit. We can define predicates like *dog* quite simply as in (73a), but on the understanding that the meta-language constant dog refers to the set of exemplars of the concept DOG (73b)—and take for granted that concepts allow for the active categorization of exemplars, so that no memorization is necessary to know whether something is an exemplar of a given concept.

(73) a. $[\![\text{dog}]\!] = \lambda x.\lambda w.\mathsf{dog}(x)(w)$.
b. $\mathsf{dog}(x)(w) \leftrightarrow x \in \{y : y \text{ is an exemplar of the concept DOG}\}$ in w

This small change makes it explicit that, while our goal as semanticists is to understand the *composition* of lexical items and phrases into larger phrases, ultimately the symbols that combine have meanings that make reference to mental modules that are non-linguistic in nature—here, concepts. (73) is essentially a standard definition for predicates in formal semantics, but fleshes out a tacit assumption.

Of course, I have not taken a stance on what exactly concepts are in (73); see Margolis & Laurence (2021) for an overview of the literature on concepts. In chapter 2, I will briefly discuss 'geometric' approaches (van Fraassen 1967; Lambert & van Fraassen 1970; Stalnaker 1981; Gärdenfors 2000); while I will reject a particular aspect of at least one such theory, it seems to me like a promising path forward, and for that reason I will sometimes talk about 'conceptual spaces' to refer to mental spaces populated by particular families of related concepts.

In this thesis, I will alternate rather freely between talking about 'exemplars of concepts' and 'individuals in sets,' and indeed between 'concepts' and 'predicates.'

1.4.2 Some syntactic assumptions

A significant focus of this thesis is a locality constraint that I will posit on Exh operators, with a focus on copular sentences. Syntactic assumptions will occasionally affect the nature of this discussion; in this section, I briefly discuss the syntax of copular sentences and some assumptions about phrase structure.

Copular sentences

While much of the data in this thesis will come from copular sentences, this is a thesis on the meaning of predicates, not copular sentences. I will make the case for this in chapter 2 (section 2.2.2). Nonetheless, let me make some very basic points about copular sentences.

The copular sentences I will be concerned with in this thesis are 'predicational' as in (74); the expression following *be* is a predicate, not an individual (as in the identificational *This is Aisha*, for example).

(74) This piece of cloth is a flag.

I follow the common assumption that *to be* is semantically vacuous at least in predicational copular sentences. In (74), *flag* is a one-place predicate taking the *e*-type individual *this piece of cloth* as its argument.

There are several possibilities about the syntactic locus of the subject, *this piece of cloth*. In (74), while *this piece of cloth* is spelled-out at Spec-TP, I follow the widely accepted predicate-internal subject hypothesis (Zagona 1982, Kitagawa 1986, Contreras 1987, Kuroda 1988, Speas 1986, Sportiche 1988, Koopman & Sportiche 1991, Burton & Grimshaw 1992, Guilfoyle et al. 1992, McNally 1992, McCloskey 1997), according to which subjects are base-generated lower than Spec-TP. Where is the phrase initially generated? Semantically, it does not matter much. It could be the specifier of the DP *a flag* or of a PredP (commonly posited for copular sentences); I will simply put subjects in Spec-*v*P, assuming a stative *v* corresponding to the verb 'to be.' This assumption will occasionally have implications for the relative scope of subjects and adjuncts like additive particles or Exh: there will sometimes be cases where the relative height of reconstructed subjects and such adjuncts will not be the same depending on whether subjects in copular sentences are taken to be base-generated in Spec-*v*P or somewhere lower.

This brings us to the next syntactic issue worth mentioning, viz. the relative height of adjuncts and specifiers.

Phrase structure

I will assume some version of Bare Phrase Structure (Chomsky 1995a), and as such not draw any bar levels: non-complex heads X are dominated immediately by XP, not X′. Nothing hinges on this. One question that will arise at a few points is the relative height of adjuncts and specifiers. Adjuncts in standard X′-theory are located below specifiers: they are both immediately dominated by and sister to X′. But adjuncts have also sometimes been taken to adjoin to maximal projections above specifiers (see e.g. Hornstein et al. 2005:196). The few times that this comes into play, I will simply consider both possibilities.

1.5 Outlook

This thesis is organized as follows.

Chapter 2 constitutes the central empirical contribution of this thesis. Much as overviewed in the introduction of the present chapter, I make two points. First, I argue that exhaustivity is at play in the meaning of 'taxonomic' predicates like *comedy* or *green*, due to data like (75), repeated from (5).

(75) a. This comedy is #(also) a tragedy.
b. The white flag is #(also) green.

Second, I argue that this exhaustivity effect displays novel behaviour. Predicates are not only obligatorily exhaustified, but also insist that the Exh operator associating with them be syntactically local to them. I call exhaustivity effects that are subject to these twin requirements 'controlled exhaustivity' effects; I contrast controlled exhaustivity with the 'free' exhaustivity standardly assumed in the literature (e.g., Chierchia et al. 2012). The evidence for the locality requirement does not only come from clause-internal contradictions like these sentences in (75) without the additive; in fact, I will show that it is generally observable in the meanings of taxonomic predicates even in non-contradictory sentences like (76).

(76) a. Every play on this shelf is a comedy.
b. Every flag here is green.

In chapter 3, I discuss one of the central datapoints motivating the claim that exhaustivity is involved in the observed meanings of predicates, namely additive particles, and their interaction with exhaustivity. Taking for granted Bade's (2014, 2016) claim that additives are obligatory when and only when they serve to avoid unwanted exhaustivity effects, the question I discuss in this chapter is how exactly they do this. There are two parts to this question. Consider again (77), repeated from (7).

(77) A: Who sang at the party?
B: Aisha sang. Ben sang #(too).

Assume that in the absence of an additive, both sentences in B's answer are exhaustified in a problematic way; without the additive, B's answer contains the LFs in (78).

(78) Exh_{ALT} [Aisha$_F$ sang]. Exh_{ALT} [Ben$_F$ sang].

There are two problems that *too* must fix: the fact that the first sentence contradicts the plain meaning of the second, and the fact that the second sentence contradicts the plain meaning of the first. In the chapter, I discuss several possibilities from the literature about how additives could interact with Exh operators to ensure consistency in discourse: Exh operators might disappear entirely, additives could weaken them by scoping underneath them syntactically, or additives could involve pruning certain problematic alternatives from the domains of the Exh operators. I give new evidence in favour of the latter approach. In the presence of an additive, (78) can be unpacked as in (79) (where in the second sentence, the relative scope of *also* and Exh is immaterial).

(79) [$\text{Exh}_{\{\text{Aisha sang, }\sout{\text{Ben sang}}\text{, Carrie sang}\}}$ Aisha$_F$ sang].

also [Exh$_{\{\sout{\text{Aisha sang}}, \text{Ben sang, Carrie sang}\}}$ Ben$_F$ sang].

Chapter 4 discusses the types of expressions that control exhaustivity and what causes predicates to be alternatives for the purposes of controlled Exh. The first half of the chapter moves beyond simplex predicates like *comedy* or *green* to make the case that controlled exhaustivity can also be found with complex phrases. It is observed with PPs (80a) as well as constituents which might correspond to the VP or some larger clausal constituent (80b).

(80) a. The book about cats is #(also) about bicycles.
b. The train crashed that killed Aisha #(also) killed Ben.

Sentences like (80b) in particular offer a promising empirical bridge between the sort of obligatory-additive data discussed by Bade (2016) (e.g., (77)) and the data discussed in this thesis (which mostly involves non-focused predicates, rather than contrastive topics). However, I leave for future research how exactly to reconcile the two empirical domains. As for the paradigm with PPs in (80a), I discuss this at much more length in the chapter; indeed, these data suggest that the phenomenon of 'thematic uniqueness' well-known in the literature on thematic roles should be thought of as part of a broader pattern, namely controlled exhaustivity. (80a) and data like it are descriptively a thematic uniqueness effect, and best understood as the result of controlled Exh.

The second half of chapter 4 refocuses on the controlled-exhaustivity data with simplex predicates, asking what predicates are alternatives to which. Why are *green*/*white* and *comedy*/*tragedy* interpreted as mutually incompatible (meaning that they are alternatives for controlled Exh), but not, for example, *green*/*table* or, say, *mermaid*/*figure-skater*? What is the relation that determines whether two predicates are alternatives for controlled Exh? This relation is assumed in chapter 2 to be the relation of 'co-hyponymy,' but several types of data complicate this, including the controlled exhaustivity effects from PPs or VPs, as well as data with predicates that act as alternatives despite not clearly being cohyponyms. There is a generalization that does seem to hold, namely that predicates (including complex phrases) behave as alternatives if and only if they contribute the same kind of information in a given sentence. I call the kind of information contributed by a given predicate the 'jurisdiction' of the predicate. Following this line of thinking, I end up suggesting that the acceptability of usually-incompatible predicates in certain circumstances can be understood as a core part of the paradigm. For instance, *fork* and *spoon* are contradictory in (81a) due to sharing jurisdictions, but are consistent in (81b) due to having different jurisdictions (the form vs. the function of the object).

(81) a. This fork is #(also) a spoon.
b. Poor Aisha! This fork is her spoon.

Moving on, chapter 5 turns to a particular subset of the data discussed throughout this thesis, namely so-called 'summative' predicates (predicates which are true of an individual by virtue of being true of that individual's parts). Colour adjectives are such predicates; (82) is true if each part of the flag is green.

(82) The flag is green.

Summative predicates control exhaustivity:

(83) The white flag is #(also) green.

Throughout the thesis, colour terms are used as an example among other 'controllers' of exhaustivity; they are simply part of the general paradigm of taxonomic predicates controlling exhaustivity. Analyzing them as such necessarily involves two claims: colour terms are lexically weak, and they come with the exclusion of other colour terms.

In this chapter, I compare my analysis of summative predicates with two theories from the literature on *homogeneity effects* that take sentences like (82) to involve universal quantification over parts (rather than existential meaning paired with the exclusion of other predicates). Homogeneity is best known from examples with pluralities, where the observation is that pluralities in positive sentences (84a) behave as *universals*, whereas pluralities in negative sentences (84b) are interpreted as negated *existentials*.

(84) a. The students sang. (plural homogeneity)
≈ all the students sang
b. The students did not sing.
≈ none of the students sang

The same goes within atoms for summative predicates:

(85) a. The flag is green. (subatomic homogeneity)
≈ all parts of the flag are green
b. The flag is not green.
≈ no part of the flag is green

The homogeneity paradigm has received a number of analyses in the literature, most of which are tailor-made for pluralities, sometimes with the assumption that they should carry over to the subatomic level as well. In the chapter, I focus on the 'excluded-middle presupposition' analysis of Löbner (2000), Gajewski (2005) and others, as well as the Exh-based theory of Bar-Lev (2018, 2021).

These theories consistently run into problems with conjunction data like (86).

(86) The flag is white and green.

I show that the conjunction is necessarily interpreted as Boolean in sentences like (86). Both the theories I focus on problematically predict (86) to mean that *all* parts of the flag are both white *and* green. Whatever their status for plural homogeneity, these theories do not carry over to subatomic homogeneity. Thus, chapter 5 effectively defends my controlled-exhaustivity account of predicates' strong meanings against potential competitors for a subset of the data (colour terms, and more generally summative predicates).

Finally, chapter 6 turns to trying to explain the controlled nature of Exh with predicates. There are two basic explananda: why predicates require the presence of Exh, and why they require Exh to be local (in some sense) to them. To even begin answering these questions, we will need to deal with two questions: whether it is really true that Exh is always present, and what the nature of its locality horizon is. Building in part on the discussion in chapter 3 on the interaction between Exh and additives, I argue that it is at least possible (and in some cases empirically necessary) to claim that Exh is *always* present with taxonomic predicates. The three domains where Exh might appear

absent are sentences with additives, with conjoined predicates, and under sentential negation. But I will show that the data with both additives and conjunction require us to postulate one or more Exh operators, whose problematic effects are simply neutralized by the domain-restriction of its alternatives (in the case of additives) or its syntactic position (as I will argue is the case for alternatives). As for the negation data, here there is no need whatsoever to posit an Exh operator, but it is *possible* to claim that there is one if we adopt the presuppositional Exh operator ('Pexh') of Bassi et al. (2021). I will therefore do so, making it possible to claim that there is no exception whatsoever to the obligatory and local presence of an Exh operator with taxonomic predicates.

From there, I will characterize the locality constraint on Exh as being that it must occur within the XP headed by the predicate it associates with. This will involve analyzing the syntax of conjunctions as involving the inheritance of the conjuncts' category to the entire conjunction. The goal is to capture that a single Exh operator can occur above both conjuncts, as in (87).

(87)

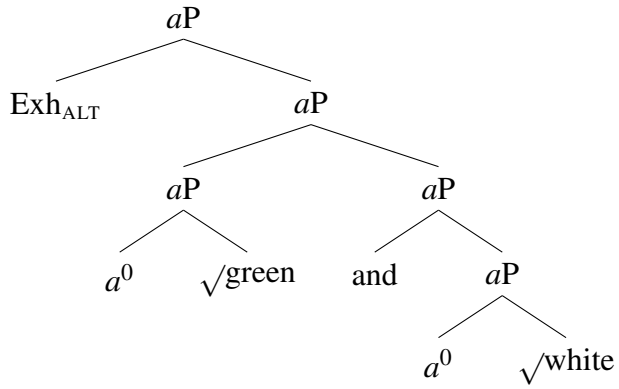

To capture that Exh is obligatorily present in taxonomic predicates' XP, I argue that there is a [*u*Exh] feature on derivational morphemes (a^0, n^0, etc.): they must find an Exh operator to Agree with. The [*u*Exh] feature can percolate upward along with labels, but no higher. Following Chomsky (2001) in assuming that (*i*) the failure to value unvalued features leads to a crash in the syntactic derivation, and (*ii*) there is no upward agreement, positing a [*u*Exh] feature on derivational morphemes captures that Exh is both obligatory and necessarily local to them. The consequence of this proposal is that derivational morphemes effectively clean up conceptual space: when they take a root, they immediately require it to be exhaustified vis-à-vis similar roots, ensuring that whatever overlap is presence in the conceptual module does not carry over linguistically (modulo *and* and *also*).

Chapter 7 concludes by providing a brief summary of the thesis, as well as some further discussion of additive particles' interaction with predication.

Chapter 2

On taxonomic predicates

2.1 Introduction

Since the work of Horn (1972), it is generally accepted that the meaning of certain predicates involves exhaustification. For example, *warm* in (1) is interpreted not only as meaning that its subject meets the standard for being warm, but also that it does not meet the standard for being hot or boiling.

(1) The water is warm.
$\rightsquigarrow$ The water is not hot or boiling.

This not due to the lexical meaning of *warm*, but to exhaustivity. Indeed, the 'not hot' meaning of *warm* disappears in downward-entailing (DE) environments:

(2) If the water is warm, Aisha will put it outside.
$\approx$ if the water is warm, hot, or boiling, Aisha will put it outside

As such, it must be that *warm* has certain other predicates as alternatives.[1] These predicates form entailment scales like (3):

(3) {*warm*, *hot*, *boiling*}

Horn (1972:47) provides sets of predicates like the following as involving strengthening:

(4)
a. {*pretty*, *beautiful*}
b. {*cool*, *cold*}
c. {*intelligent*, *brilliant*}
d. {*good*, *excellent*}
e. {*happy*, *ecstatic*}
f. {*like*, *love*}
g. {*dislike*, *hate*}

I will refer to such predicates as 'scalar' predicates—not because they are gradable (most of them are, but not e.g. *like*/*love*), but because they form entailment scales with other predicates.

In this chapter, I show that exhaustivity is not only involved in the meaning of scalar predicates. Rather, language also routinely exhaustifies predicates that do not take part in an entailment scale—

[1] For Horn (1972), the scales formed by predicates like *warm* and *hot* are stored in the lexicon; more recent work takes them to fall out from entailment and/or relevance (e.g., Katzir 2007, Chierchia et al. 2012).

what I will call 'taxonomic' (or 'non-scalar') predicates. The evidence for this comes from the way these predicates interact with conjunction and additive particles. Consider a contradictory sentence like (5).

(5) #Some comedies are tragedies.

The contradiction in (5) certainly comes across, prima facie, as the result of lexical-conceptual incompatibility between *comedy* and *tragedy*. It is surprising, then, that the contradiction disappears by means of an additive or through conjunction (6). Note that (6a) is a *clause-internal* additivity effect; *also* refers anaphorically to *comedies*, not to prior material in the discourse.

(6) a. Some comedies are also tragedies.
b. Some plays are both comedies and tragedies.

Building on this observation, I will claim that taxonomic predicates have their meaning delimited by other content words in language, specifically due to exhaustivity. Classes of taxonomic predicates form sets of alternatives, much like those in (3) and (4) but without entailment relations (setting aside complications addressed in due time).

In this sense, the empirical contribution of this chapter is to introduce a new empirical domain in which to observe exhaustivity. But this foray into a new domain comes with a surprising theoretical consequence for Exh. With taxonomic predicates, Exh is constrained in a way not previously discussed in the literature. Specifically, taxonomic predicates not only make Exh obligatory, but they also dictate its approximate syntactic position. Indeed, taxonomic predicates require an Exh operator taking them in its immediate scope. This can be observed in a variety of ways, including sentences like (5). If Exh was free to scope anywhere, we would expect the possibility of a global Exh (7). I change *some comedies* to *the comedy* in (7) to avoid complications from the strengthening of *some* to *not all*.

(7) Exh_{ALT} [the comedy is a tragedy].

A global Exh as in (7) does not create a contradiction out of non-contradictory lexical material: Exh only excludes alternatives that are not entailed by its prejacent. Assuming only the predicative *tragedy* triggers alternatives,[2] (7) would come out to the truth conditions in (8c) (recall from chapter 1 that Exh excludes logically independent alternatives).

(8) a. Exh_{ALT} [the comedy is a tragedy].
b. $\text{ALT} = \left\{ \begin{array}{c} \text{the comedy is a tragedy,} \\ \text{the comedy is an epic,} \\ \text{the comedy is a comedy} \end{array} \right\}$
c. $[\![(8a)]\!] = 1$ iff the comedy is a tragedy $\wedge \neg$[the comedy is an epic].

Notably, the alternative *The comedy is a comedy* is not excluded due to being entailed by Exh's prejacent; hence, no contradiction is created. To derive a contradiction, what we need is actually an Exh on each predicate (in fact, at least one predicate); due to its narrow scope, such an Exh can 'unknowingly' create contradictory meanings:

[2] With a global Exh, if the noun *comedy* in the subject also triggered alternatives, this would lead to alternatives with entailments about other plays altogether ('the epic is a tragedy').

(9) $[\![\text{The } [\text{Exh}_{\text{ALT}} \text{ comedy}] \text{ is a } [\text{Exh}_{\text{ALT}} \text{ tragedy}]]\!]$
$= 1 \text{ iff the } \begin{pmatrix} \text{comedy \&} \\ \text{not a tragedy \&} \\ \text{not an epic} \end{pmatrix} \text{ is a } \begin{pmatrix} \text{tragedy \&} \\ \text{not a comedy \&} \\ \text{not an epic} \end{pmatrix}.$

Given that taxonomic predicates not only require the presence of an Exh operator but also dictate its syntactic position, I will refer to them as 'controlling' Exh; but note that this sense of 'control' is not to be confused with the syntactic notion of control.

This chapter is organized as follows. I begin in section 2.2 by discussing sense relations among predicates, in particular the cohyponymy relation and the observation that cohyponyms are inferred as mutually exclusive. Then, in section 2.3, I show that the observed exclusivity is in fact a product of grammar, and specifically exhaustivity, as evidenced from their behaviour with additive particles and conjunction. In section 2.4, I turn to a theoretical consequence of claiming that Exh is the cause of taxonomic predicates' mutual exclusivity. In particular, deriving the right meanings involves stipulating that taxonomic predicates not only require the presence of an Exh operator, but specifically require it to take them in its immediate scope. The controlled nature of Exh with taxonomic predicates leads to a new desideratum for a theory of the syntactic distribution of exhaustivity. In the present chapter, I only point out the existence of such a constraint on Exh and give it some characterization; I will return to it more thoroughly in chapter 6.

2.2 Cohyponymic exclusivity

In the sense-relations literature (see Cann 2011 for a recent overview), hyponymy refers to the 'kind of' relation. For example, *poodle* is a hyponym of *dog* and *red* is a hyponym of *colour*. 'Cohyponyms' are sisters in a taxonomy: *poodle* and *labrador* are cohyponyms, as are *red* and *blue*.

In this section, we begin with an empirical observation: cohyponyms are interpreted as mutually exclusive in basic sentences.[3] On this point, Cann (2011:459) gives the example (10), involving hyponyms of *animal* or *mammal*.[4]

(10) #That sheep is a horse.

Sheep and *horse* are cohyponyms and are interpreted as mutually exclusive. Call this *cohyponymic exclusivity*. Sentences of the form in (10), where two cohyponyms are predicated of the same individual or set of individuals, will be the main empirical paradigm in this thesis, so it is worth introducing some terminology to talk about them: I will variously refer to examples like (10) involving as *double-predication* or *pseudo-repetition*: it is not the *same* predicate that is repeated, but there are two predicates *from a given class* applied to the same individual or set of individuals.

There is nothing particularly interesting about (10), since the effect can be entirely derived from

[3]What I mean by 'basic sentence' will become clear once we discuss in sections 2.3 and 2.4.2 the linguistic environments in which it is no longer the case that cohyponyms are interpreted as inconsistent, most saliently with conjunctions, additives, and sentential negation.

[4]I will not be too concerned in this chapter about how fine-grained taxonomies should be. In this example, it will not matter whether there is a node *mammal* between *animal* and *sheep*. Perhaps there is flexibility, with speakers creating taxonomies at the right level of granularity for the purposes of a given situation.

world kowledge. We know that, biologically, animals cannot be two species at once. No theory of predicates or sense relations would have predicted (10) to be acceptable. In spite of this, in this section, I show that cohyponymic exclusivity is in fact not only a matter of world knowledge. True, there are cases like (10), where world knowledge is involved and suffices as an explanation (although I will soon suggest that even here, world knowledge might not be enough). But when we turn to examples where world knowledge is not a reasonable culprit for cohyponyms' mutual incompatibility, we observe that cohyponymic exclusivity is in fact still observed. I motivate this empirically in section 2.2.1, and elaborate further on the empirical claim in section 2.2.2 by showing that the exclusivity effect is not an accident of the particular examples I focus on, and again in section 2.2.3 by comparing cohyponyms (taxonomic predicates) with predicates forming entailing scales, showing they behave differently. With these empirical observations in place, I will then conclude this section by outlining two theories that attempt to derive cohyponymic exclusivity as a fact of conceptual structure or of the lexicon (section 2.2.4).

2.2.1 Cohyponymic exclusivity is not reducible to world knowledge

It is not always the case that cohyponymic exclusivity can be explained from world knowledge, suggesting that something more arbitrary (about language or about concepts) is at play. Moving on from (10), consider examples like (11). These all involve attempts at predicating two cohyponyms of a particular individual or set of individuals, and they all result in contradictions just like (10). Roughly speaking, the cohyponyms come from the taxonomies of colour, film type, literary genre, utensils, jurisdiction, kinds of morphemes, and vehicles, respectively.

(11) a. #The white flag is green.
b. #Some animated films are live-action.
c. #Some comedies are tragedies.
d. #This fork is a spoon.
e. #Some federal responsibilities are provincial.
f. #Some inflectional morphemes are derivational.
g. #This train is a plane.

In contrast to (10), world knowledge is not sufficient to explain these examples, and is very clearly not a factor at all in most of them. Let's take them one by one.

At least two of these examples, (11a) and (11b), touch on part-whole structure. The contradiction effect in these examples has to do with both predicates applying to all parts of their argument: (11a) is paraphrasable as 'the entirely white flag is entirely green,' and (11b) as 'the entirely animated film is entirely live-action.' It is true that the fact that a given surface cannot be entirely covered by two different colours comes from our conceptualization of the world (and likewise with film types). Yet, there is no language-independent reason for why these predicates should necessarily modify all parts. It is possible for objects to have white parts and green parts (11a), and there are films with animated parts and live-action parts (either one after the other or at the same time on a particular frame).[5] Hence, whatever explains (11a) and (11b) must be something other than

[5]For a list of such films, see the category FILMS WITH LIVE ACTION AND ANIMATION on Wikipedia (`https://en.wikipedia.org/wiki/Category:Films_with_live_action_and_animation`, accessed December 27, 2021).

world knowledge. It could be that the lexical items are lexically specified as being true of all parts of their argument, for example—but this would be a fact of language, not world knowledge.

The same holds for the rest of these examples. Notably, (11c) and (11d) are intuited as contradictions despite the existence of portmanteau predicates (*tragicomedy* and *spork*) which refer to objects that have the properties of both the predicates in these examples. That is, a tragicomedy has properties of comedies as well as tragedies, and a spork has properties of both forks and spoons. Thus, we know that, at some level of abstraction, the set of things describable as comedies is not strictly disjoint from the set of things describable as tragedies (and likewise with forks and spoons). Yet, (11c) teaches us that tragicomedies are not describable as 'comedies that are tragedies.' Again, something other than world knowledge must be at play.

As for (11e), despite the legal separation of responsibilities between different jurisdictions, the example is not clearly at odds with some real-world examples. For example, healthcare in Canada is technically a provincial jurisdiction, but with significant federal involvement. Indeed, "Canada's health care system consists of 13 provincial and territorial health insurance plans" that are "individually administered on a provincial or territorial basis, within guidelines set by the federal government";[6] in addition to federal regulation, the federal government transfers funds earmarked for healthcare expenses to the provinces through the Canada Health Transfer.[7] Given such complications in federalism, there is no basis for (11e) to be analysed as contradictory due to a clash with extralinguistic knowledge. The same goes for (11f), which cannot (but, naively, 'should') describe portmanteau morphemes that contribute both derivation and inflection (I don't know whether such morphemes exist, but they could, and (11f) would still be contradictory). Finally, possibly like (11f), (11g) is an ambiguous case. On the one hand, there are no 'train-plane' vehicles in the real world. On the other hand, the sentence is intuitively unacceptable even if we imagine an adequate science-fiction possible world.

In spite of this, the sentences in (11) are all judged as contradictions. As such, even if (10) can be understood in terms of world knowledge, this does not hold across the board: there are cases where world knowledge is not a possible source of cohyponymic exclusivity. Something else must impose such exclusivity among cohyponyms. Presumably, this would either be the structure of concepts (how we classify exemplars, forcing them into one or another category) or something more linguistic.

Before moving on, one word of caution is in order. The sentences in (11) are all contradictions on their most basic, literal readings; but it is possible to obtain non-contradictory meanings by complicating things in various ways. For example, a sentence like (12), repeated from (11a), is quite acceptable if *white* is interpreted as referring to the past state of the flag, and *green* to the present state of the flag. A scenario could be if the speaker put an exclusively white flag in a washing machine, and it came out exclusively green.

(12) The white flag is green (now).

In the same vein, many of the sentences in (11) can also be rescued by giving one of the predicates a 'function-as' reading. In (13), the fork (which is not a spoon or even a spork) is claimed to be

[6] Wikipedia, s.v. MEDICARE (CANADA) (`https://en.wikipedia.org/wiki/Medicare_(Canada)`, accessed December 27, 2021).

[7] See for example Wikipedia, s.v. CANADA HEALTH TRANSFER (`https://en.wikipedia.org/wiki/Canada_Health_Transfer`, accessed December 27, 2021).

usable as if it was a spoon.

(13) This fork is a good spoon.

Further, many of the sentences in (11) also become acceptable if one of the predicates is interpreted 'meta-linguistically.' What I mean by this is that instead of a predicate *P* meaning '*x* is *P*,' it means 'people *say* that *x* is *P*' or 'people use the word *P* to describe *x* (perhaps incorrectly).'

(14) Some 'inflectional' morphemes are derivational.

As the scare-quotes in (14) are meant to indicate, here, the speaker is claiming that some morphemes that are usually viewed as inflectional (and not derivational) are, in fact, derivational (and not inflectional). Another way to make sentences with pseudo-repetitions acceptable is if they are interpreted for possible worlds with rules different from ours. Indeed, even (15), repeated from (10), can be understood as referring to a cartoon scenario where a horse has been turned into a sheep. The animal has the body of a sheep (and not a horse) but the spirit of a horse (and not a sheep).

(15) That sheep is a horse.

All these possible interpretations of the sentences I have used to motivate cohyponymic exclusivity are not truly counterexamples to the exclusivity effect. In all cases, the cohyponyms are still incompatible once we control for the particular complications brought by the chosen interpretations. In (12), the colours are only 'compatible' due to being predicated of the flag at different points in time; in (13), the fork is not a spoon, but it can be used as a spoon that is not a fork; in (14), the speaker is rejecting that the morphemes are inflectional at all and claiming they are only derivational; in (15), each predicate holds of the animal in a different way (the outer body vs. the inner essence of the animal), and does so in a way that is exclusive of the other (the body is that of a sheep that is not a horse, and the spirit is that of a horse that is not a sheep). I will therefore ignore these sorts of readings for this chapter (we will return to this issue in chapter 4); they are essentially unhelpful complications on the data. We will be engaging in what you could call semantic *literalism* as a methodological principle: when we observe sentences with pseudo-repetitions like (11), we will ignore complexified readings where the predicates hold at different times, in different ways, metalinguistically, or in whatever other non-straightforward way one may think of. This methodology corresponds to the basic scientific principle of controlling variables and isolating the sort of data that one is interested in—in our case, this is the meaning of predicates, not tense semantics or metalinguistic uses of expressions.[8]

Viewing the incompatibility of the predicates in (11) as the result of the cohyponymy relation (i.e., sisterhood in a taxonomy) is intuitively appealing from the evidence seen so far. But to see that this really is a fact of cohyponymy, we also need to observe that similar judgments do *not* hold with non-cohyponymic predicates. To this purpose, the sentences in (16) take predicates from different taxonomies rather than the same taxonomy: the predicates are not cohyponyms. Crucially, the contradiction effect disappears (other kinds of semantic weirdness may remain; I return to this

[8]In fact, the fact that predicates can hold 'in different ways' as in (13) and perhaps (15) is not clearly reducible to other parts of language like tense semantics or the metalinguistic use of expressions. We will return to this, and integrate the observation as a core part of our theory, in chapter 4.

immediately below). (16a–d) mix and match predicates from different domains identified in the prose immediately preceding (11); (16d–f) explicitly attempt pragmatically unlikely combinations, still not reaching the contradiction effect of (11).[9]

(16) a. This train is green. cf. (11a) and (11g)
b. Some animated films are tragedies. cf. (11b) and (11c)
c. ?? Some forks are provincial. cf. (11d) and (11e)
d. ??Some forks are tragedies. cf. (11c) and (11d)
e. This spider is an accountant.
f. ??Green ideas sleep. (cf. Chomsky 1957)

Some comments are in order. First, I assume that human beings have the ability not just to detect whether or not a sentence is deviant, but also to reason about what sort of deviance is involved. Intuitively, the sentences I have marked with ?? in (16) are more or less deviant, at least on their literalist interpretations.[10] But none of these sentences are *contradictions* in the way that the sentences in (11) are. That is, the meaning of some of these sentences is obscure,[11] but not contradictory like (15). I have marked the deviant sentences with ?? rather than # to emphasize the distinction, rather than to make a claim about the strength of the deviance.

The second point about (16) pertains to (16e) specifically. The acceptability of (16e) highlights just how easy it is to jump into 'cartoon mode'—our willingness to accept sentences that completely violate our knowledge of the real world, because we can picture a cartoon world in which the sentence would be true. (16e) requires no effort whatsoever to accept as a possible sentence. This reaffirms the limitations of using world knowledge to explain cohyponymic exclusivity. Indeed, I provided the examples in (11) in order to avoid using examples like (17) (ignore for a moment the non-literalist reading where a horse has been turned into a sheep), repeated from (10), about which I wrote that the predicates' exclusivity is essentially uninteresting because it can be explained by world knowledge alone. Yet, (16e) too violates world knowledge about spiders, but without this resulting in any deviance. As such, without reference to the cohyponymy relation, it is not clear what distinguishes examples like (16e) from (17).

(17) #That sheep is a horse.

In other words, why can I effortlessly imagine a cartoon where a spider is an accountant, but not one where a sheep is a horse?[12]

The conclusion so far is that the following generalization holds:

(18) **Cohyponymic exclusivity**
Cohyponyms are interpreted as mutually incompatible. In some but not all cases, this may

[9] The same might go for (16c), although I suppose a province could own a special fork, making it 'provincial.'

[10] One example where a non-literalist reading is particularly salient is (16f). If *green* is taken to mean 'ecological' and *sleep* to mean 'lie dormant,' (16f) is in fact completely acceptable to mean that ecological ideas are lacking in popularity at a certain time and place.

[11] In the case of (16d), the deviance is akin to a presupposition failure: *tragedy* refers to the logic in a story, but forks are not stories, so the predication 'does not go through,' as it were.

[12] Again, I am putting aside once again a non-literalist interpretation of (17), according to which a horse has been turned into a sheep. This is fully justified because (16e) is fully acceptable on its literalist reading; the spider is literally a spider and it is literally an accountant (I am not even sure what a non-literalist reading would be).

coincide with world knowledge.

2.2.2 Cohyponymic exclusivity is only due to the meaning of cohyponyms

I have just framed cohyponymic exclusivity purely in terms of the meaning of cohyponyms themselves. Is this right, or are there other factors involved in creating the apperance of cohyponymic exclusivity (e.g., the verb *to be*)? In this short section, I defend my view that the sorts of infelicities under discussion arise due to the meanings of the cohyponyms, and not something else about the sentences we are considering. Cohyponymic exclusivity does not arise due to the verb *to be*, due to the sort of determiner used in the subject position, or due to the information structure of pseudo-repetitions (where one cohyponym is in the subject and the other in the predicate of the sentence).

So far, all the sentences we looked at involved the verb *to be*. A sceptic could suggest that cohyponymic exclusivity is the result of something about *to be*. Yet, the exclusivity persists without *to be*: adjectival cohyponyms that are stacked on a single noun are interpreted as mutually exclusive without there being a copula.

(19) a. #This derivational inflectional morpheme stumped many grad students.
b. #The green white flag hung low in the rain.
c. #I watched three animated live-action movies yesterday.
d. #Canada's federal provincial healthcare system collapsed following budget cuts.

Even scenarios like (20), which could explain why one adjective is stacked above the other, do not improve the judgment.

(20) SCENARIO: *We are discussing two jointly funded projects; one is funded by the federal and municipal governments, and the other is funded by the federal and provincial governments.*
A: The federally funded project is important.
B: #Are you talking about the MUNICIPAL federal funds or the PROVINCIAL federal funds?

Since the effect persists without *to be*, blaming the copula for the contradictions we have been discussing would likely be the wrong direction to take.

Second, the reader may have noticed that I have been using a handful of different syntactic structures for the sentences under discussion, as exemplied by the sentences in (21).

(21) a. #**This** comedy is a tragedy.
b. #**The** comedy is a tragedy.
c. #**Some** comedies are tragedies.

We could add relative clauses to this set:

(22) #A tragicomedy is a comedy **that** is a tragedy.

Using examples with a variety of determiners and structures is meant to avoid, once again, heading in the wrong direction by blaming something irrelevant, such as the choice of determiner and

whether it introduces asserted or presupposed material, or whether the sentence has a relative clause or not. In fact, cohyponymic exclusivity is observed regardless of these sorts of changes. Thus, I maintain that the relevant part of the discussion is the meaning of the cohyponyms, not the sort of determiner they occur with or whether or not they are in a relative clause.

More evidence that we should not blame cohyponymic exclusivity on particular syntactic environments, or something about the information structure or syntax of pseudo-contradictions, comes from predicates which do *not* behave as cohyponyms. This is what we turn to in this next subsection.

2.2.3 Taxonomic vs. scalar predicates

In this section, we put a caveat on the claim that cohyponyms control exhaustivity: this is only true if we understand cohyponymy as necessarily involving taxonomies, and not scales.

In the sense-relations literature, the notion of cohyponymy does not only include sets of predicates that form what I will call *taxonomies*, but also sets of predicates that form *scales*. Indeed, the sets in (23) are both examples of the cohyponymy relation, but only the second has entailment relations between the predicates.

(23) a. {*green*, *white*, *red*, *blue*, *pink*, ...}
b. {*warm*, *hot*, *boiling*}

(23a) is a *taxonomy* in the sense of being a classification hierarchy forming a tree structure, where for any two given sisters, neither entails the other.

(24)

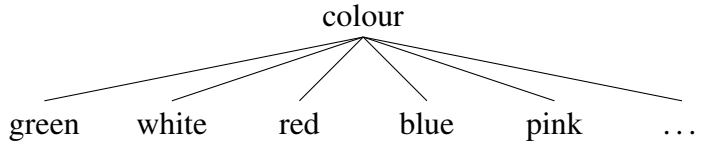

This is generally the case with the sorts of examples we have been discussing:

(25) a. GENRES: {*comedy*, *tragedy*, *epic*, ...}
b. UTENSILS: {*fork*, *spoon*, *knife*, ...}
c. VEHICLES: {*train*, *plane*, *boat*, ...}
d. ...

I am not proposing that sets like those in (25) are part of grammar or the lexicon. It suffices to say that they are the result of world knowledge about the concepts referred to by the predicates, or perhaps a notion of relevance (Geurts 2010). I will return to this question in chapter 4.

In contrast, (23b) is a *scale* because the members of the set are entirely ordered by entailment.

(26) *warm* $\supset$ *hot* $\supset$ *boiling*

Scales can be thought of as cohyponyms, as in (27) (overlooking for simplicity the fact that there are cold temperatures too). Cann (2011:460) calls them 'quasi-hyponyms' of *temperature* because *temperature* is of a different syntactic category (quasi-hyponyms are "predicate-denoting expres-

sions like adjectives which often seem to relate to (abstract) nouns as superordinates rather than some other adjective").

(27)

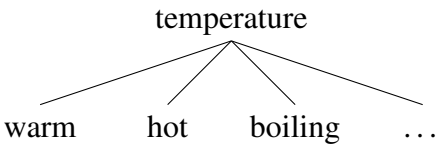

But calling all the predicates in (23a) and (23b) 'cohyponyms' is of limited profitability, given that the logical relations of (27) are different from (24). The cohyponyms in (24) have no entailment relation to one another,[13] whereas the 'cohyponyms' in (27) do (assuming that *warm* lexically means '*at least* hot enough to count as warm,' and similarly for the other predicates). In this thesis, I will continue simply referring to 'cohyponymy' with the intent of referring to the taxonomic relations of (24), and not the scale relations of (27). Cohyponyms are 'taxonomic predicates' rather than 'scalar predicates.'

This difference is important for understanding cohyponymic exclusivity, which is observable with taxonomic predicates, as we have just seen, but not scalar ones:

(28) Some of the warm plates are hot.

These examples serve as an important minimal pair with pseudo-repetitions like (29) that have been the basis of much of my discussion so far.

(29) #Some of the green flags are white.

Indeed, (28) shows that there is nothing inherently formally wrong with pseudo-repetitions. The problem with examples like (29) is due to the choice of predicates, not the form of the sentence.

Beyond scalar predicates narrowly defined, we also do not observe the contradiction effect of cohyponymic exclusivity between hyponyms and their hypernym:

(30) Some of the red flags are (in fact) scarlet.

(30) means that some of the entirely red flags are entirely scarlet; scarlet is the shade of red that they have. It does not mean that the flags have a non-scarlet red part and a scarlet part.

On last note: in this thesis, I will not be too concerned with the exact definition of cohyponymy even among taxonomic predicates. On their basic definition, cohyponyms are sisters in a conceptual taxonomy. There are two points about this that I will not take a stance on. First, a priori, it's not clear whether taxonomies are fixed mental objects fed to language from the conceptual module, or whether it is possible for ad hoc taxonomies to be created on the fly in conversation. On the former view, there are taxonomies of (for example) colours, genres, utensils, and so on; on the latter view, there could be strange taxonomies like 'the kinds of things that I can put in this bottle' or 'the things that are cold enough to help me heal my burn.' It will suffice to usually assume that

[13] At this point in the chapter, I am discussing cohyponyms as having a mutual-exclusivity entailment (two cohyponyms A and B lexically or conceptually exclude one another). I will later claim that the meaning of cohyponyms is underlyingly weak. This means that cohyponyms are actually logically independent of one another, other than all entailing the same hypernym.

taxonomies are fixed, given that the examples in this thesis are usually given without discourse contexts; I will raise the issue when it matters, mostly in chapter 4.

The second point about cohyponyms is how seriously one wants to take the notion of *sisterhood*. On a strict understanding of cohyponymy, for instance, *cat* and *dog* might be cohyponyms as direct daughters of *mammal*, but *cat* and *fish* would not be cohyponyms (at least for human beings who classify *fish* as a sister of *mammal* and not particular species of mammals):

(31)

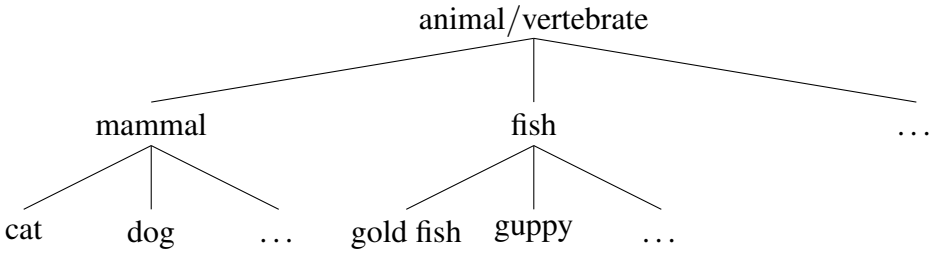

That is, strictly speaking, *cat* and *fish* are not cohyponyms. But they are both members of the same taxonomy such that they share a hypernym (*animal*) and neither asymmetrically entails the other. This could be said to be 'good enough' for them to count as cohyponyms. I will not be too concerned about the strict technical sense of the term 'cohyponym' in this thesis, and some of my examples will use terms that are not strictly speaking cohyponyms.

2.2.4 Some prior literature on exclusivity

We have just seen that many cohyponyms are interpreted as mutually exclusive in a way that is 'arbitrary'—that is, not open to explanation from world knowledge alone. What, if not world knowledge, underlies cohyponymic exclusivity? In this section, we will consider two slightly different proposals that exist in the literature. Neither of these proposals was explicitly formulated to deal with cohyponymic exclusivity as such, but they could explain the data we have seen so far. The first proposal claims that concepts from a given domain are inherently incompatible as a fact of discrete categorization 'overriding' world knowledge; exemplars are forcibly categorized as belonging exclusively to one concept or another. The second proposal is quite similar, but moves the burden of explanation from a non-linguistic conceptual module to the lexicon (i.e., part of language). The claim is that the lexicon is structured such that lexical items delimit one another, so that a given area of conceptual space can be covered by only one lexical item (the assumption being that concepts underlie predicates' meanings). I now turn to these proposals in turn.

Exclusivity from partitioned conceptual geometries

Gärdenfors (2000) lays out a theory of conceptual domains as involving geometries (see also van Fraassen 1967; Lambert & van Fraassen 1970; Stalnaker 1981 for similar 'geometric' theories of concepts). According to this proposal, concepts exist in a geometric space, where they are associated with regions of the space. The space is partitioned: distinct regions cannot overlap, which means that concepts within a given space are inherently mutually exclusive. From this perspective, cohyponymic exclusivity holds regardless of world knowledge because the partitioning of conceptual space is imposed even over a continuous array of exemplars.

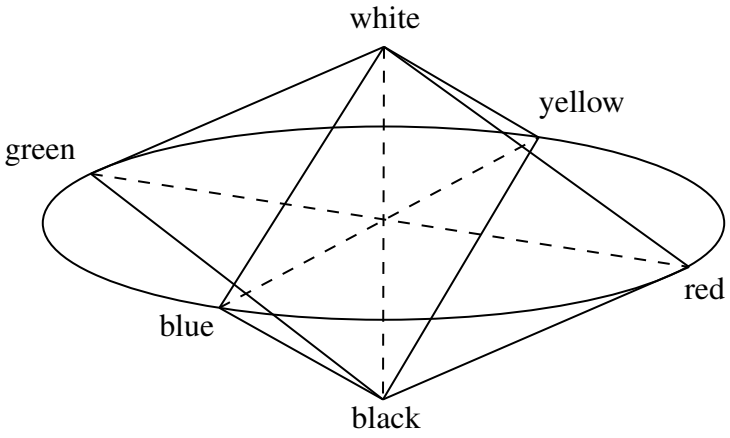

Figure 2.1: *The conceptual space associated with colours is three-dimensional (Sivik & Taft 1994:148).*

Let's see how this works by building up a conceptual geometry in the abstract. First, we need dimensions to create an *n*-dimensional space. A real example discussed by Gärdenfors (2000) is the space for colours. This space is three-dimensional, with the dimensions corresponding to values for brightness, hue, and saturation (Figure 2.1). I will abstract away from this and use a two-dimensional space for simplicity.

The space created by the dimensions is populated with exemplars. These are represented as individual points with a coordinate for each dimension, providing them a place in the space. In the early-stage building of a conceptual space, exemplars are associated with a concept on a case-by-case basis. In the following diagram, I have a hypothetical two-dimensional space populated with exemplars assigned to various concepts through the coloration of the points (this is for presentation; they are not intended to represent concepts corresponding to colours).[14]

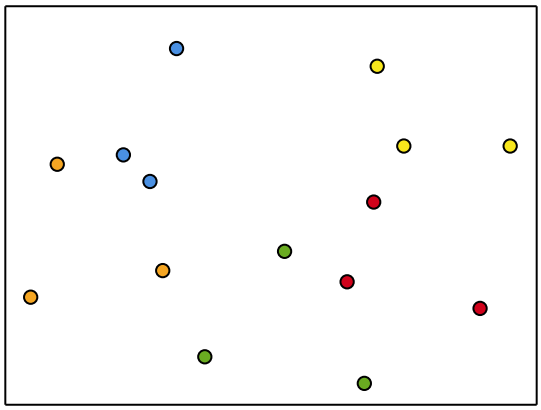

From these exemplars, we mentally abstract away from them by creating prototypes (see e.g. Rosch & Mervis 1975; Rosch 1978; Hampton 2006 on prototypes in concepts, and Margolis & Laurence 2021 for critical discussion). The prototype is a point in the space whose coordinates correspond

[14]See Gärdenfors (2000:88, 124) for diagrams similar to the following three figures. I made the Voronoi partition in the third figure using Alex Beutel's online generator (`http://alexbeutel.com/webgl/voronoi.html`, accessed Nov. 4, 2021).

to the mean value for each dimension of all the exemplars for the concept that the prototype represents. A prototype is an idealization, which does not need to correspond to an actually-existing exemplar. In the following figure, I represent prototypes with large crossed circles.

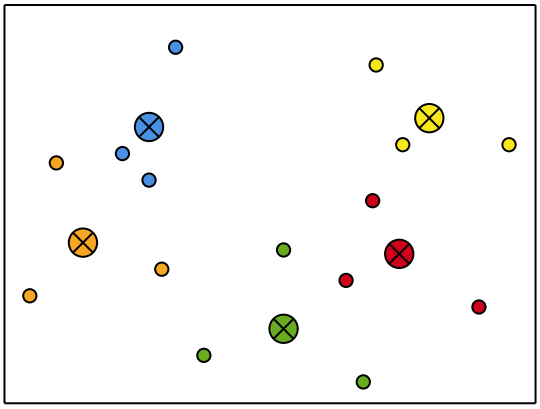

Now that the space is populated with abstract prototypes corresponding to concepts, Gärdenfors (2000) posits that the space is partitioned, such that each prototype is associated with a region comprising all the points (not just the exemplars, but all the points of the space) that are closer to it than to any other prototype. This is a 'Voronoi partition' of the space.

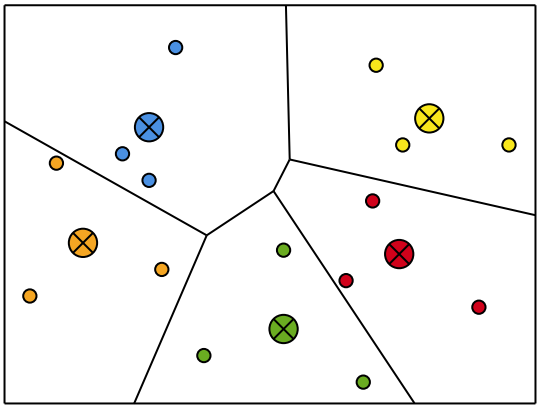

A concept, then, is a cell in a Voronoi partition of a conceptual space, and a 'conceptual domain' is the entire space.

This theory has several clear advantages. It captures prototype effects (according to how close an exemplar is to the prototype), gradient similarity between different concepts (if cell B is between cells A and C, concept A is more similar to B than to C), and the possibility for human beings to actively categorize new exemplars (by checking their coordinates in a space and seeing in which cell they land).

To see how this works in explaining cohyponymic exclusivity, and in particular the persistence of this phenomenon even when world knowledge could not explain it, consider (32), repeated from (11d).

(32) #This fork is a spoon.

Pretend for simplicity of presentation that utensils are in a one-dimensional geometry, where the unique dimension corresponds to the length of their prongs. A true spoon is at zero cm, a normal fork at several cm, and a spork is somewhere in between. This is a simplifying assumption for ease of presentation, of course; in particular, knives have nowhere to go in this toy geometry.

We have immediately run into a complication, namely the question of whether the concept of *spork* is active in the space. This is a more general complication for Gärdenfors' (2000) theory. For example, in order to capture the hyponymy relation, we need hyponyms and hypernyms not to be active at the same time (Peter Gärdenfors, p.c.). If *dog* and *poodle* were both active in a space, then poodles would not be dogs, because the space would be partitioned between dogs and poodles.

To avoid becoming stuck on this particular point, I will simply consider for (32) both the possibility that *spork* is active in the space, and the possibility that it is not. If *spork* is active, we have a space with at least three cells:

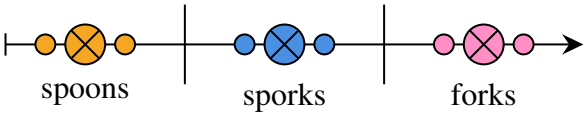

If *spork* is not active, there are only two cells, and the middle-ground exemplars that would have been categorized as sporks are now categorized as either forks or spoons, depending on which prototype they are closest to:[15]

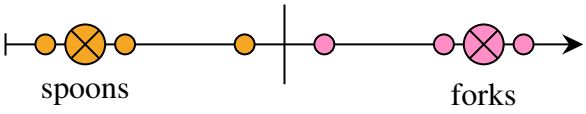

Either way, the world knowledge that there are exemplars in the middle-ground between forks and spoons is 'overridden' by the categorization process. The middle-ground exemplars are either categorized as sporks (which are neither spoons nor forks), or they are variously pushed into the nearest other concept, whether that means they are categorized as spoons (which are not forks), or as forks (which are not spoons). In this way, Gärdenfors' (2000) theory predicts cohyponymic exclusivity straightforwardly.

Exclusivity as a fact of the lexicon

Another proposal in the literature, while less directly aimed at deriving something like cohyponymic exclusivity, can explain this exclusivity as a fact of the lexicon. The idea is very similar to Gärdenfors'. For Gärdenfors, the conceptual space and the lexicon act essentially as a single partitioned geometry. The approach that we will now consider claims that the lexicon in fact imposes a partition over and above a non-partitioned conceptual space. This is, indeed, the basis of early 20th-century structuralist linguistics.

de Saussure (2011[1916]) viewed the lexicon as forming a 'structure' (in Gärdenfors' terms, a geometry) in which the meaning of lexical items is delimited by other adjacent lexical items. As

[15]For simplicity, the figure overlooks movement of the prototypes toward the centre following the inclusion of middle-ground exemplars in its cell.

he wrote, "the value of each term results solely from the simultaneous presence of others" ('solely' is an exaggeration; he assumed some sort of conceptual space associating with lexical items), as in his famous diagram (de Saussure 2011[1916]:114–115):

In de Saussure's diagram, what is important for us is not the distinction between signified and signifier, but the idea that lexical items form a structure, with each lexical item delimiting the others. As the arrows are meant to indicate, they are in a geometric relation of some sort to one another.

To motivate this, de Saussure concerns himself entirely with crosslinguistic comparisons. His examples begin with content vocabulary, in particular predicates approaching a cohyponymic relation. While they are not quite there, his discussion would translate over to cohyponyms. The famous example that de Saussure (2011[1916]:115) gives is that in French, there is a single word, *mouton*, for sheep and their meat, whereas English has a word for each concept, namely *sheep* and *mutton*.[16] As he explains, "the difference in value between *sheep* and *mutton* is due to the fact that *sheep* has beside it a second term while the French word does not" (de Saussure 2011[1916]:116). The area in conceptual space corresponding to sheep and mutton is partitioned in two in English (due to there being two lexical items) but not in French. de Saussure's view is that there is a single 'universal' conceptual space unaffected by a language's lexicon; what changes language to language is which part of the space is associated with which lexical items. Conceptual space is continuous and lacking in boundaries, but the lexicon superimposes a discrete partition over it. Hence, on this view, we could claim that cohyponyms are judged as contradictory in spite of world knowledge for a reason virtually identical to what Gärdenfors (2000) proposes, only moving the partitioning from the conceptual space itself to the lexicon.

A significant distinction arises between this theory and Gärdenfors' for logical vocabulary, which de Saussure proposes to treat in the same way as content vocabulary. Here, in contrast to his discussion of content vocabulary, de Saussure's discussion becomes quite aligned with modern formal semantic theory. Perhaps most clearly, one of his example is grammatical number, where "the value of a French plural does not coincide with that of a Sanskrit plural" because Sanskrit has a dual whereas French does not (de Saussure 2011[1916]:116). That is, the meaning of the plural is delimited by non-plural number(s), but the amount of non-plural number(s) a language has is an area of crosslinguistic variation. As such, the plural does not have quite the same meaning from language to language. This notion of the plural being 'delimited' by other numbers is just another way of saying that the plural is *strengthened* by them. de Saussure could just have well have written that there is a quantity implicature, or similarly an exhaustivity effect, whereby the semantically unmarked plural is strengthened to mean 'not singular' in French and 'neither singular nor dual' in Sanskrit, resulting in different meanings (a sum of two is plural in French, but not Sanskrit). In fact, this is exactly the theory argued for by Sauerland et al. (2005): on these authors' view, the plural is semantically unmarked, and its meaning ('not one' in English) is the result of strengthening

[16] Another similar example he gives (de Saussure 2011[1916]:116) is the French verb *louer* 'rent / rent out' (pay for or receive payment for) vs. the English verbs just given or German *mieten* and *vermieten*.

vis-à-vis non-plural number(s).[17]

In this specific sense, de Saussure's structuralism is surprisingly aligned with modern formal semantic theory. In his discussion of grammatical vocabulary like number and tense, de Saussure can be reinterpreted as positing an exhaustivity effect. On the other hand, de Saussure is quite out of step with modern theorizing in collapsing the strengthening of grammatical vocabulary with effects around content vocabulary like *sheep* and *mutton* or *rent* and *rent out*. The closest modern parallel to this comes from Horn scales like {*warm*, *hot*, *boiling*}, where a predicate (*warm*) is strengthened by stronger scalemates. But Horn scales touch on a rather limited proper subset of content vocabulary; de Saussure's proposal is meant as a general claim about the lexicon.

I will be claiming in this chapter that exhaustivity is, as a matter of fact, generally involved in the meaning of predicates, not just Horn scales. de Saussure (2011[1916]) was right to collapse the strengthening of grammatical vocabulary with the strengthening of predicates. On the other hand, he was wrong to view this as a fact of the lexicon. In this way, this chapter can be read as a defense of the spirit of structuralist linguistics, but with the machinery moved from the lexicon to the grammar.

2.3 Cohyponymic exclusivity as a fact of grammar

So far, we have seen that cohyponyms are regularly interpreted as mutually exclusive of another, even when world knowledge cannot realistically be taken to underlie this exclusivity. We then briefly overviewed two theories that can capture this exclusivity. Gärdenfors (2000) could point to the Voronoi partitioning of conceptual space, whereas de Saussure (2011[1916]) would blame a structure in the lexicon that exists over and above the conceptual space.

Despite moving the explanation for cohyponymic exclusivity from a non-linguistic conceptual module to the lexicon, de Saussure's (2011[1916]) proposal has an important point in common with Gärdenfors' (2000). The two theories converge in predicting that this exclusivity is impossible to remove linguistically in any sentence that entails that an individual should be in the set denoted by both cohyponyms, because cohyponymic exclusivity is a fact of categorization or of the lexicon. It is baked into lexical items; no individual could be in the sets denoted by two different cohyponyms. The only way for two cohyponyms to be predicated of the same individual would be through expressions that are non-intersective in one way or another. This could be achieved through non-intersective conjunction (33a) or possibly via non-intersective adjectives like *fake* (33b).

(33) a. The buildings are big and small.
≈ some of the buildings are big and the other ones are small
b. a fake gun

[17] Another example from logical vocabulary that de Saussure (2011[1916]:117) gives is the difference between languages like French or English that have future tense or something like it, and other languages (e.g., Proto-Germanic) with no future: "to say that the future [in Proto-Germanic] is expressed by the present is wrong, for the value of the present is not the same in Germanic as in languages that have a future along with the present." That is, the meaning of the present is delimited by the future in English-type languages, but not in Proto-Germanic–type languages, since they do not in fact have a future. This is also highly reminiscent of Sauerland's (2002) argument that the present tense is vacuous at least in English, with its intuited meaning arising due to pragmatic competition with the past tense (de Saussure would add the future tense).

I will refer to such expressions as 'non-intersective expressions'—expressions whose effect is that double-predications might be true and felicitous due to an individual not needing to fall in the sets denoted by the predicates. We now turn to seeing that this prediction from both theories does not hold up. The conclusion will be that cohyponymic exclusivity cannot be baked into cohyponyms' lexical or conceptual meaning. It must come from somewhere else.

2.3.1 Set-intersection in double-predication and the weak conceptual-lexical meaning of cohyponyms

Consider again the examples in (34), repeated from (11).

(34) a. #The white flag is green.
b. #Some comedies are tragedies.

Is it really the case that, if we put aside expressions with non-intersective composition like some instances of conjunctions and adjectives like *fake*, the cohyponymic exclusivity these examples display is non-removable?

In fact, language comes with several lexical items capable of removing the contradiction. I will focus on conjunction (35) and additive particles (36), but there are others (37). We just saw that some conjunctions are non-intersective, so I hasten to add that I will be commenting on (35) and motivating that the conjunction there is intersective.

(35) a. The flag is white **and** green.
b. Tragicomedies are plays that are both comedies **and** tragedies.

(36) a. The white flag is **also** green.
b. Tragicomedies are comedies that are **also** tragedies.

(37) Some comedies are **simultaneously** tragedies.

Examples like (36) and (37) may require a bit of work to explain why one predicate is in the subject position and the other one is in the VP, but scenarios that make this natural are not hard to come by. For example:

(38) a. SCENARIO: *We are at a plant that specializes in recycling cloth; pieces of cloth must be sorted by colour. There is a pile of flags, most of which are entirely white, but a few of which are both white and green. The boss tells a worker that they need to remove all the green parts from the otherwise white flags:*
b. Some of the white flags are **#(also)** green, so I want you to cut off the green parts.

(39) a. SCENARIO: *There are two flags on a flagpole, one of which is purple, the other is green and white. From your position, you can only see the white part of the green and white one. From my position, I can see both colours.*
b. You: Oh no, the white flag is on fire!
Me: This is very upsetting. The white flag is (in fact) **#(also)** green, and green is my favourite colour.

I put aside for the rest of this thesis the sort of pragmatic work that is helpful to bring out the

acceptability of sentences involving one predicate in the subject and another in the VP (i.e., pseudo-repetitions).

With that aside, we go back to the basic observation about (35)–(37). Specifically, it is not clear how it could be maintained in light of these examples that the cohyponyms under discussion are inherently conceptually or lexically incompatible. Crucially, none of these involve non-intersective expressions. I take this up for *and* (35) and *also* (36) in the following subsections. As for (37), it is perhaps a little bold to claim that there is no non-intersective expression here without a formal proposal for the meaning of *simultaneously*. But it seems intuitive that *simultaneously* should be intersective in some sense: I would think that *simultaneously* asserts that two events or states hold at the same time ($e_1 \wedge e_2$, at the same time), without modifying the internal nature of these events or states.

Controlling for an intersective conjunction

The conjunction examples in (35) are crafted to control for an underlyingly intersective conjunction. I will discuss conjunction in depth in chapter 5, to which the interested reader is referred. For now, suffice to say that both examples are necessarily intersective, but for different reasons. In (35a), the conjunction is intersective because there is never a non-intersective *and* with atomic subjects (Winter 2001, Paillé 2021). In (35b), the lexical item *both*, at least in the syntax given in the example, marks intersective conjunction (Schwarzschild 1996, Paillé 2021, a.o.).

Additives are always intersective

As for (36), it has never been claimed that there is a non-intersective *also*. Indeed, as we saw in chapter 1, *also*'s presupposition is anaphoric to a proposition in the discourse context (Kripke 2009[1990]). It does not modify the content of these previous propositions. For instance, adopting Heim's (1992) use of indexation to capture anaphoricity and assuming an assignment function g where $[1 \rightarrow \lambda w.\ \mathsf{sing}(a)(w)]$ (40a), the truth conditions in (40b) hold for the sentence *Ben also sang*.

(40) a. [Aisha sang]$_1$. Ben also$_1$ sang.

b. $[\![\text{also}_1\ [\text{Ben}_F\ \text{sang}]]\!]^g = = \begin{cases} 1 \text{ if } \mathsf{sing}(b), \\ \# \text{ if } \neg\mathsf{sing}(a), \\ 0 \text{ otherwise} \end{cases}$

This is equivalent to (41), highlighting the intersective nature of the additive's effect:

(41) $[\![\text{also}_1\ [\text{Ben}_F\ \text{sang}]]\!]^g = \begin{cases} 1 \text{ if } \mathsf{sing}(a) \wedge \mathsf{sing}(b), \\ 0 \text{ if } \mathsf{sing}(a) \wedge \neg\mathsf{sing}(b), \\ \# \text{ otherwise} \end{cases}$

There is no claim in the literature that a non-intersective *also* exists, and it is not clear what such a lexical entry for *also* would look like. To begin having a sense of what a non-intersective *also* would look like, we can take a cue from an example with non-intersective conjunction:

(42) The parents are 28 and 30 years old.

Presumably the existence of a non-intersective *also* would mean that (43) would also be acceptable, then:

(43) #The parents are 28 years old. They are also 30 years old.

Clearly, this is not the case. The effect is no better sentence-internally:

(44) #Some 28-year-old parents are also {30 years old, 30-year-olds}.

Therefore, the data in (36) clearly show that an intersective semantic expression, namely *also*, makes possible the pseudo-repetition of cohyponyms. Thus, it is possible for the cohyponyms under discussion to be predicated of a single individual through set-intersection.

Interim conclusion: many cohyponyms are underlyingly weak

In light of the possibility to pseudo-repeat cohyponyms with set-intersective linguistic material like *and* and *also*, we must conclude that the cohyponyms showing this behaviour are not, in fact, underlyingly mutually exclusive.[18] Cohyponymic exclusivity (18) is a linguistic mirage, rather than a fact of cohyponyms as concepts or lexical items. In light of this, there must be a grammatical process that obscures cohyponyms' underlying consistency in basic sentences.[19] This misleadingly gives conceptual space or the lexicon the *appearance* of being neatly divided. (45) provides a summary of the claim.

(45)

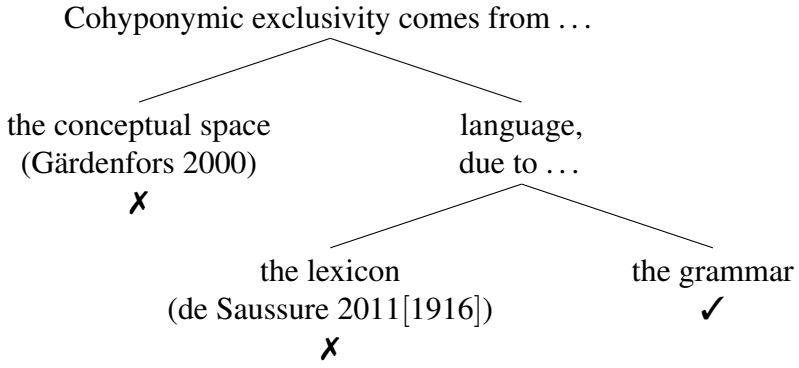

A different approach to the fact that the set-intersection of cohyponyms is not (necessarily) empty would be to claim that data like (46) motivate that the predicates *comedy* and *tragedy*, rather than interacting with grammar, are *lexically ambiguous*.[20] On this view, *comedy* is lexically ambiguous between referring only to bona-fide comedies (46a), or to both comedies and tragicomedies (46b) (and likewise for *tragedy*).

(46) a. #This comedy is a tragedy.

[18]To be clear, my discussion of 'cohyponyms' is meant to refer to predicates like *comedy* and *tragedy* that are not underlyingly incompatible, but are intuited as compatible in basic sentences. It is also perfectly possible for cohyponyms to be mutually exclusive due to world knowledge (e.g., *sheep* and *horse*) or possibly other factors.

[19]By 'basic sentences,' I mean any sentence where cohyponyms are not joined by an additive, a conjunction, or other similar material like *simultaneously*.

[20]I thank Dimitrios Skordos for raising this possibility.

b. This comedy is also a tragedy.

On such an alternative approach, however, it's not clear why (46a) is infelicitous; speakers would be able to choose the weak lexical meaning for *comedy* and *tragedy*. I therefore assume that these predicates are lexically non-ambiguous, and—as stated above—something grammatical causes them to become apparently incompatible in (46a).

Before moving on, let me comment briefly on the meaning of a particular set of cohyponyms, namely colour adjectives. Consider the meaning of the following sentences:

(47) a. #The white flag is green.
b. (i) The flag is white and green.
(ii) The white flag is also green.

The problem with (47a) is that both colour terms are true of all parts of the subject (the flag),[21] clashing with our conceptualization of the world, according to which a given surface may only have one colour. When we solve this problem with *and* or *also* as in (47b), what changes is not the particular combination of brightness, hue, and saturation that *white* and *green* denote; these sentences do not mean that the entire flag is of a single colour situated in between prototypical whiteness and greenness. Rather, what changes is that the *universal* meaning of the colour terms in (47a) is now interpreted more or less *existentially*.[22] In (47b) *white* means 'partially white,' and *green* means 'partially green.' That is, what we learn from the data with *and* and *also* is not something about the conceptual structure of colours; rather, it is that their lexical meaning is underlyingly existential (Harnish 1976, Levinson 1983):

(48) $[\![\text{green}]\!] = \lambda x.\exists y[y \sqsubseteq x \wedge \text{green}(y)]$.

I will often emphasize colours' existential meaning by writing them in the meta-language as '$(\text{colour})_{\exists}$.' Colour terms are the main empirical focus of chapter 5. In general, even if the (non-) consistency of colour terms is due to quantification over parts rather than the existence of overlap in conceptual space (as is the case with *comedy* and *tragedy*, for example), I will spend the thesis collapsing the colour-term data with other data. This is because they pattern together empirically in terms of their interaction with *and*, *also*, and so on.

2.3.2 Exclusivity as a product of exhaustivity

We have just observed that the conceptual/lexical meaning of cohyponyms is underlyingly weak, as demonstrated by their behaviour with intersective material like *and* and *also*. This raises the question of where exactly cohyponymic exclusivity comes from.

Exhaustifying cohyponyms

If intersective linguistic material can lift the exclusivity effect, presumably it is something grammatical that causes cohyponymic exclusivity in the first place, taking underlyingly compatible

[21] To be precise: *white* actually holds of the x that is a flag, and *green* holds of the unique x that is (entirely) white and a flag.

[22] In fact, the meaning is still stronger than a mere existential: *white* and *green* must still jointly cover all parts of the flag. We will return to this observation many times.

expressions and making them incompatible with one another. Descriptively, if it is accepted that cohyponymic exclusivity is a product of grammar rather than conceptual space or the structure of the lexicon, what we have at hand is a strengthening effect. The core (conceptual) meaning of a cohyponym like *comedy* does not exclude tragicomedies, but its strengthened meaning does, and likewise with *tragedy*; hence the contradiction in (49).

(49) #This comedy is a tragedy.

To posit the right grammatical mechanism as the culprit behind the strengthening effect underlying cohyponymic exclusivity, we must find a way to strengthen predicates' meaning that predicts that the strengthening can be circumvented through conjunction or additive particles, as in (50).

(50) a. This is a comedy and a tragedy.
b. This comedy is #(also) a tragedy.

Let's focus on additives. As we saw in chapter 1, and indeed as part of the basic premise of this thesis, additives have been argued to interact with exhaustivity in an important way. Indeed, I follow Bade (2016) (cf. Krifka 1998; Sæbø 2004; Aravind & Hackl 2017) in taking additives to be obligatory when an unwanted exhaustivity effect would arise without them.

Recall that Bade considers two theories of why additives could be obligatory: either because they somehow circumvent an unwanted exhaustivity effect that would otherwise arise, or because it is obligatory to use them when their presupposition is met (due to *Maximize Presupposition*). Bade distinguishes empirically between these two theories using data like (51).

(51) a. Aisha sang. Ben #(also) sang.
b. Aisha sang. Ben didn't (also) sing.

It is true that the obligatory additive in (51a) could in principle be the result of *Maximize Presupposition*, but this claim would wrongly predict that the additive would be equally obligatory in (51b). Under this account, the negation would have no effect on whether *also* is obligatory, because presuppositions project past negation. Clearly, this is not the case.

Hence, Bade favours an account where the additive is obligatory in (51a) because otherwise, the second sentence would contradict the first, due to an exhaustivity effect:

(52) $[\![\text{Exh}_{\text{ALT}}\ [\text{Ben}_F\ \text{sang}]]\!] = 1$ iff Ben sang $\wedge$ Aisha didn't sing.

In light of this, I propose to explain cohyponymic exclusivity by claiming that cohyponyms form a set of alternatives and are obligatorily exhaustified. As a first approximation, this looks like (53), where Exh scopes above a predicate and takes as its domain a set of alternatives generated by replacing the predicate with its cohyponyms.[23]

(53) a. Exh_{ALT} [this is a comedy].
b. $\text{ALT} = \{\text{this is a comedy, this is a tragedy, this is an epic, } \dots\}$
c. $[\![(53a)]\!] = 1$ iff this is a comedy $\wedge$ this is not a tragedy $\wedge$ this is not an epic $\wedge \dots$

This, then, is the first important theoretical claim of this thesis: exhaustivity is standardly involved

[23] Note that I will substantially modify the syntax of this proposal in section 2.4 and adjust slightly the nature of the alternatives in chapter 4.

in the intuited meaning of cohyponyms. In addition to being a new domain in which to observe exhaustivity effects, this proposal is also a substantial addition to the way such predicates are usually treated in formal semantics. The norm (e.g., Heim & Kratzer 1998) is to adopt the idea that predicates simply denote sets. The observed meanings of such predicates are primitive: *cat* means 'cat,' *comedy* means 'comedy,' and so on. I am proposing that this is not the case: at the very least, there is some structure to these predicates' meanings, in that *comedy* (in basic sentences, due to exhaustivity) does not only mean 'comedy' but also 'not a tragedy,' and so on. Claiming that *comedy* refers to the set of comedies, and that this is all there is to say about this predicate, would be like claiming that *warm* refers to everything that is above a certain temperature, without incorporating into the grammar the idea that *warm* is also strengthened by stronger predicates like *hot*.

Cohyponymic strengthening vs. contrastive focus

The claim in (53) is *not* simply that predicates can be contrastively focused so as to exclude other predicates. The possibility of contrastively focusing predicates is not a new claim by any means. My claim is that cohyponyms *always* undergo this exhaustification, not just in contexts where they are contrastively focused. This is why, in (53), I write simply 'comedy' and not 'COMEDY' or 'comedy$_F$.' Let's unpack this a bit.

Simplifying, there are two classes of expressions in the domain of strenghtening and alternatives: those expressions that always bear alternatives and do not need to be focused in order to do so (e.g., *some*), and those that only bear alternatives in particular conditions (e.g., *Aisha*).

(54) a. Ben ate some of the apples. ⇒ no focus on 'some'
$\rightsquigarrow$ Ben did not eat all of the apples

b. (i) Aisha$_F$ ate the apples. ⇒ focus on 'Aisha'
$\rightsquigarrow$ Ben did not eat the apples

(ii) Aisha [ate the apples]$_F$. ⇒ no focus on 'Aisha'
$\not\rightsquigarrow$ Ben did not eat the apples

In (53), given that the predicate *comedy* is new information and indeed the predicate of the sentence, there is little way to tell whether focus-marking plays a role in cohyponymic exclusivity here. That is, in (53), it may be that *comedy* is F-marked for independent reasons, rather than the F-marking being of importance to cohyponymic exclusivity. But in fact, there is reason to think that cohyponymic exclusivity is calculated *on top of* normal focus effects, so we want to distinguish between normal focus effects on predicates (which involves F-marking/focus intonation) and the computation of cohyponymic exclusivity.

To see this, start by assuming that colour terms, as cohyponyms, behave just like *comedy* does in (53):

(55) $[\![\text{Exh}_{\text{ALT}}\ [\text{the flag is green}]]\!]$
$= 1$ iff the flag is green$_\exists \wedge$ the flag is not white$_\exists \wedge$ the flag is not red$_\exists \wedge \ldots$

The flag is green semantically means that it it is 'only green$_\exists$' (55), which is pragmatically strengthened to mean that it is entirely green (we assume that all parts of the flag have a colour). Now notice that the strong meaning of colour terms (where they are obligatorily predicated of all material

parts of their argument) is still intuited in contexts where a colour term is in a contrastive-focus relationship with some other predicate that does *not* come from the domain of colours:

(56) This car is not ITALIAN$_F$, it's GREEN$_F$. (Michael Wagner, p.c.)

Two things simultaneously hold in (56): *green* means 'all green,' so it must be exhaustified vis-à-vis other colour terms; but *green* and *Italian* are also alternatives, and in this sentence—as opposed to (55)—we clearly need contrastive focus intonation on *green*.

In principle, it could be that there is a single exhaustivity effect on the second clause in (56) simultaneously excluding *Italian* and other colour terms. The set of alternatives would be {it's Italian, it's green, it's red, it's white, it's ...}. This would make it puzzling why there is clear focus intonation on *green* in (56) but not (55); but even more empirically troubling, a problem would arise in contexts like (57) (inspired by Katzir (2013, 2014)).

(57) The car is not required to be ITALIAN$_F$, it's required to be GREEN$_F$.

Intriguingly, (57) has and hasn't the following entailments/inferences:

(58) a. $\rightsquigarrow$ The car is not required to be Italian.
b. $\not\rightsquigarrow$ The car is required not to be Italian.
c. $\rightsquigarrow$ The car is required not to be red$_\exists$.
d. $\rightsquigarrow$ The car is required not to be white$_\exists$.
e. $\rightsquigarrow$ (and so on for all other colours)

As can be discerned from the difference in how *Italian* and *red* scope vis-à-vis the negation and the universal *required*, the relation between *green*$_\exists$ and *Italian* is different from the relation between *green*$_\exists$ and *red*$_\exists$. This can be captured if there are *two* Exh operators, one associated with the alternatives {Italian, green} and the other being in charge of cohyponymic exclusivity, based around the alternatives {green, red, white, ...}.[24]

(59) a. $\text{Exh}_{\text{ALT-2}}$ [it's required $\text{Exh}_{\text{ALT-1}}$ [t to be green]].
b. (i) $\text{ALT-1} = \{t \text{ to be green}, t \text{ to be red}, t \text{ to be white}, \dots\}$
(ii) $\text{ALT-2} = \left\{\begin{array}{l}\text{it's required Exh}_{\text{ALT-1}}\ [t \text{ to be green}],\\ \text{it's required Exh}_{\text{ALT-1}}\ [t \text{ to be Italian}]\end{array}\right\}$
c. $[\![(59a)]\!] = 1$ iff it's required to be $\left(\begin{array}{c}\text{green}_\exists\ \&\\ \text{not red}_\exists\ \&\\ \text{not white}_\exists\ \&\\ \text{not} \dots\end{array}\right) \wedge$ it's not required to be Italian.

Thus, we want to distinguish between the strengthening of *green* vis-à-vis its cohyponyms, and the strengthening of *green* vis-à-vis the contextually provided alternative *Italian*.

In effect, if cohyponymic exclusivity was understood as a run-of-the-mill focus effect on predicates, (57) (and by extrapolation, (56)) would have to be understood as there being 'double focus' on *green*. Combined with the observation that sentences with cohyponyms do not really require focus intonation, I conclude tentatively that the formalism in (53), where the cohyponym in Exh's

[24] In (59), it is not clear what the status of the embedded $\text{Exh}_{\text{ALT-1}}$ is when it comes to determining the alternatives of $\text{Exh}_{\text{ALT-2}}$. I simply ignore it in the truth conditions, but this might be wrong.

prejacent is not presented as focused in any way, is in fact the right idea—at least as far as cohyponymic exclusivity is concerned. Predicates bear their cohyponyms as alternatives in the same way that *some* bears *all* as an alternative, rather than in the way that expressions like *Aisha* can bear alternatives when focused (54b). Cohyponymic exclusivity occurs as a matter of course, and without requiring focus intonation; and predicates are exhaustified vis-à-vis their cohyponymic alternatives *in addition to* (and *separately from*) other focus relations they may enter in. In this sense, there is something very lexical-like about the strengthening of cohyponyms; notice in particular about (59) that the creation of cohyponymic exclusivity must be dealt with *prior* to the other exhaustivity effect (it must be *more local* to the predicate), a point to which we return in section 2.4.

Before moving on, let me also comment specifically on the quantificational force of colour terms. Like other cohyponyms, my proposal is that the strong meaning of colour terms is derived from the exhaustification of the asserted colour term (which is lexically existential) with all other colour terms as excluded alternatives. This is shown in (60), repeated from (55).

(60) $[\![\text{Exh}_{\text{ALT}}\ [\text{the flag is green}]]\!]$
$= 1$ iff the flag is $\text{green}_{\exists} \wedge$ the flag is not $\text{white}_{\exists} \wedge$ the flag is not $\text{red}_{\exists} \wedge \ldots$

The logic here is worth making explicit. As stated, the meaning derived is only of the form 'the flag is only partially of such and such colour.' This does not actually make the existential meaning of the colour terms universal. While much current work uses exhaustivity to turn various types of existentials into universals (Bowler 2014; Bar-Lev & Margulis 2014; Meyer 2016; Oikonomou 2016a,b; Bassi & Bar-Lev 2018; Francis 2019, 2020; Staniszewski 2020), in (60), colour terms' lexical meaning remains existential and non-asserted colours are excluded. Thus, I rely on world knowledge kicking in to strengthen the meaning even more: we know that all areas of a surface have some colour, so 'partially green, and no other colour' in (60) is strengthened pragmatically to 'green all over.' This is well-motivated empirically: sentences with an overt *only* also come to mean that a colour is true of all parts of its argument (61). In (61), *only* excludes other colours and the result is again the meaning that the flag is entirely green, by virtue of being at least partially green and having no other colour.

(61) The flag is only green.

As a final note on presentation, I will generally not continue writing '…' in my alternative and truth conditions for sentences with colour terms (meant to indicate that I have, obviously, not listed out all basic colour terms in English). I will just pick a few colour terms (usually white, green, and red), leaving implicit that the other colour terms are also alternatives.

The claim that the meaning of words like *comedy* or *green* involves exhaustivity comes with some interesting theoretical questions and consequences. I will turn to these soon. First, however, let us trace back our steps a little bit, and ask whether the claims I have made so far really hold up. In particular, I claimed that cohyponyms are underlyingly consistent, and that it is only a grammatical exhaustivity process that yields the observed exclusivity. This was based in part on data from additive particles. In the next subsection, I ask whether it might not be better to turn this claim on its head, and say that cohyponymic exclusivity is the result of the plain meaning of cohyponyms (that is, they are in fact underlyingly strong), and additives somehow manage to weaken their normal strong meanings. I will claim that this is not the case, and that the claims I

have made so far are indeed correct.

2.3.3 Might *also* weaken predicates' meanings?

In section 2.3.1, I argued that data like (62), show us that the underlying conceptual or lexical meaning of predicates like *white* and *green* and *comedy* and *tragedy* is weak, but involves strengthening in basic sentences.

(62) a. #This comedy is a tragedy.
b. This comedy is also a tragedy.

My stance can be summed up as follows, with reference to the examples in (62).

(63) **The hypothesis pursued in this thesis: Exh strengthens cohyponyms**
The non-contradictory sentence (62b) demonstrates that cohyponyms' lexical meanings are weak, and (62a) shows us that the lexical meaning of cohyponyms is regularly strengthened in language.

This can be represented graphically as in the following figure, where *g* is a grammatical process (exhaustification) that takes predicates from a non-partitioned conceptual space *A*, and creates cohyponymic exclusivity (i.e. the *appearance* that the space does not contain overlap), as in *B*:

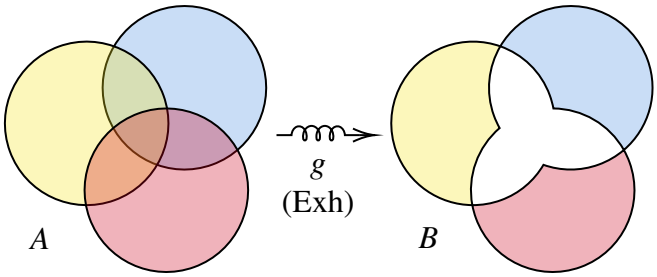

In this section, we consider what an analysis of the data in (62) would look like, if we wanted to maintain that the predicates are truly underlyingly contradictory (due to their lexical meaning or the underlying conceptual structure). This would mean in effect that *also* can weaken meaning: it would *remove* a contradiction by letting additional exemplars be included in the reference of one or both predicates. (64) spells this out, again with reference to (62).

(64) **A counter-hypothesis: *also* weakens cohyponyms**
(62a) shows that cohyponyms are lexically/conceptually inconsistent, and (62b) shows that additive particles can weaken predicates.

On this view, *also* signals that at least one of the predicates denotes a superset of its normal denotations. To help the reader compare the two competing hypotheses at a glance, I sum up this alternative hypothesis graphically as well:

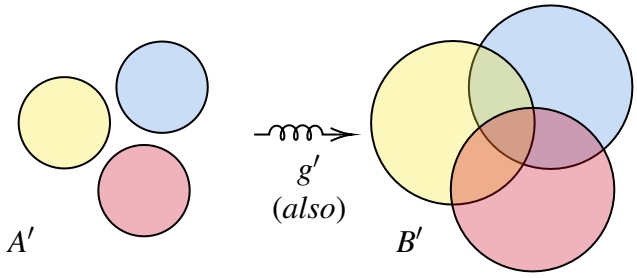

Under (64), predicates that are in a cohyponymy relationship with one another, such as *comedy* and *tragedy*, are underlyingly strong, in the sense of being mutually incompatible. There is no overlap in the conceptual space (A'). However, a grammatical process g'—which, under this hypothesis, corresponds to the presence of *also* rather than Exh—expands the reference of these predicates. This can result in the appearance of overlap (B').

In my opinion, there is a serious lack of conceptual appeal to this alternative proposal. The typical use of *also* is as in discourses like (65), repeated from (40a).

(65) Aisha sang. Ben also sang.

On my actual hypothesis, the behaviour of *also* in the data with cohyponyms is fully collapsed with its behaviour in (65). In both cases, it avoids an unwanted exhaustification effect (how exactly it does this is the topic of chapter 3). On the alternative hypothesis in (64), *also* does two essentially unrelated things. This is clearly less desirable.[25] A second conceptual hurdle for (64) is that on this view, *also* would effectively be reaching into the predicates (so to speak) to change their meaning. This is a surprising thing for a sentential adjunct to do. One would have expected the job of broadening predicates' denotation to be taken up by something with a syntactically closer relationship to the predicate, in particular a modifier internal to the predicate's XP (e.g. a suffix like *-ish*, so which we return below, or an adjective like *fake*, as in (33b)).

Nonetheless, for the sake of discussion, let us put aside these conceptual qualms, and ask in earnest whether *also* can broaden the meaning of predicates as under (64). Let me emphasize that (64) is similar to my actual claim (63) in important ways. My claim is not that *also* weakens as a direct consequence of its lexical meaning, but—following Bade (2016)—it *does* have a de-exhaustification effect of some sort. In this sense, it actually does some kind of weakening—more accurately, it has an *anti-strengthening* property. Thus, teasing apart (64) from my actual proposal (63) is not trivial.[26]

[25] Another possible related issue with the hypothesis in (64) is that *also* adds meaning, namely an additive presupposition. Thus, in order to maintain the hypothesis in (64), one has to claim that *also* is in fact ambiguous between the particle that we usually observe (the 'normal,' strengthening *also*) and the weakening *also* in (62). This is not impossible, but odd, and one would want evidence of such an *also* beyond data like (62). However, I will in fact be claiming in chapter 3 that *also* does 'weaken' meaning by making Exh operators exclude fewer alternatives (in addition to strengthening it through its addition of a presupposition). As such, this criticism does not really hold.

[26] At least, it is not trivial on the premise of this section, which is to focus exclusively on the 'anti-exclusivity' function of *also*, while putting aside that other expressions like *and* or *simultaneously* have this function too. In principle, it could be that I am wrong to analyze the *and* in the examples given as intersective, and I have not taken a particularly elaborate stance on the meaning of *simultaneously*. To be careful about the data, each of these expressions is worth taking one by one.

In what follows, I show on empirical ground that the alternative hypothesis (64) is not appealing. I will control for a non-overlapping conceptual space by turning to predicates that we know on independent grounds do not overlap. We will see that *also* does not behave as predicted by (64). Further, we will also see that there is in fact an expression in English which *does* expand the meaning of predicates, namely *-ish*. This will let us use *-ish* to make sure that the tests we are putting *also* through really are expected to work, under hypothesis (64). The conclusion will be that *also* does not truly broaden predicates' meanings (while *-ish* does).

We start with a description of the suffix *-ish*, which looks to me like the best candidate for a real instantiation of what (64) suggests *also* might do.

-ish weakens predicates

In this section, I show that *-ish* does exactly what the hypothesis in (64) suggests *also* might do: it selects a predicate and expands the set of referents of which it is true. I will focus exclusively on the *-ish* found on words or phrases (66a), not the *-ish* that modifies entire clauses (66b).

(66) a. green-ish
b. This is a decent painting. I like it, ish.

Sugawara (2017) analyzes *-ish* as taking a gradable predicate P and yielding the meaning that the subject comes close to meeting the standard to count as P, but ultimately falls short.[27] Thus, (67) means Aisha is just slightly too short to count as tall.

(67) Aisha is tall-ish.

Sugawara formalizes this as in (68) (where s(P) is the standard of P and d_c is a "contextually provided expectation value").

(68) $[\![\text{-ish}]\!] = \lambda P_{\langle d,\, et\rangle}.\ \lambda x_e.\ \max\{d|P(d)(x)=1\} < s(P) \wedge (s(P) - \max\{d|P(d)(x)=1\} < d_c)$

However, this analysis is insufficiently general. Indeed, Sugawara (2017) notes *-ish* can occur with non-gradable elements, including adverbs, common nouns, names, and numerals. The following examples are all from Sugawara (2017) except for (69d-ii).

(69) a. We have to leave now-ish/at two o'clock-ish.
b. This seemingly toy-ish Ferrari is in fact a cellphone.
c. six-ish
d. (i) Gibson-girl-ish hair
(ii) Stocks, in fact, come off their best four days since September, as traders shifted focus from negative European headlines to mostly positive U.S. economic reports. "We're beginning to see something of this **Santa Clausish** rally. This is

[27] Specifically, Sugawara (2017) writes that *-ish* takes gradable predicates and yields non-gradable ones (e.g., something can be 'tall-ish' but not 'more tall-ish'). It is of type $\langle\langle d,et\rangle, \langle e,t\rangle\rangle$. She finds that this does not work for nouns, which can in fact be the mirror image of this: *-ish* takes non-gradable nouns and yields a gradable adjective. For instance, *-ish* can take *child* and yield the gradable adjective *childish* (something can be 'more childish'). But I suggest that the *-ish* in *childish* is not the same as the one in *tall-ish*; in fact, Sugawara (2017) notes that this use of *-ish* appeared at a different stage in the history of English. Besides, as far as I can tell, someone can be ?*childish-ish* but not **tallish-ish*, suggesting the two *-ish* suffixes on *child* here are homophonous but different morphemes.

> what tends to happen now. We're up six of the last seven days, not by all that much and historically this is the 'Santa Claus' period of time which extends into the New Year," said Dan Greenhaus, chief global strategist at BTIG.[28]

Intuitively, the meaning of *-ish* has more to do with approximation than degree. 'Santa Claus-ish,' for instance, is paraphrasable as 'Santa Claus-like'; it does not mean that the rally comes close, but fails to meet, the degree for being Santa Claus (or something like this). Likewise, in (69d-i), the meaning we want to capture is paraphrasable as 'hair approximating the hair of a Gibson girl,' not 'hair that doesn't quite meet the standard for being Gibson girl hair.'

In centering on degree semantics, Sugawara's (2017) proposal mistakes an epiphenomenon for the true underlying meaning of *-ish*. That is, if *-ish* is about approximation or comparison, then when it takes a gradable predicate like *tall*, the resulting meaning ('approximating *tall*') happens, in this case, to be equivalent to meaning that 'x doesn't quite meet the standard for being tall.' Something that meets the standard for being tall would simply be described as *tall*; something that comes close to meeting it can be described as approximating being tall (that is, being *tallish*); something that does not come close to meeting the standard does not approximate being tall, and therefore is not *tallish*.[29] As such, what these examples show is that *-ish*, rather than referring specifically to degree, weakens or otherwise shifts the meaning of its complement P, so that the extension of P-*ish* includes individuals that only approximate being P.

There is a question about whether *-ish* broadens (weakens) or 'shifts' the meaning of its complement. That is, if x and y are truly green and z is almost green, is the extension of *greenish* $\{x, y, z\}$ or just $\{z\}$? The former is a weakening effect, the latter is a 'shift' (the simultaneous inclusion of the not-truly-green z and exclusion of individuals denoted by *green*). That is, are green things greenish too? As far as linguistic intuitions are concerned, the answer is obviously no:

(70) #x is greenish. *(false in the above toy scenario)*

But I think there are reasons to think this is not due to the lexical meaning of *-ish*, but to an exhaustivity effect. The idea is that, in my toy scenario, the linguistic intuition that only z is greenish arises because *greenish* competes with *green*, and therefore comes to mean *not green*. The weirdness of (70) is captured because (70) involves negation of the stronger alternative 'x is green' when x is in fact green.

There is empirical motivation for this analysis. If we are dealing with an exhaustivity effect, it should preferably disappear in downward-entailing (DE) environments. This seems right:

(71) a. If the flag is pinkish, you can start driving.
b. If we leave in five minutes-ish, I'll give you candy.
c. Every tallish student will fail this class.
d. Every diligent-ish student will pass this class.

It is possible (indeed preferred) to interpret these sentences as meaning that a truly pink flag would

[28] `https://www.cnbc.com/id/45779919`, accessed December 17, 2020.

[29] Lower-bound gradable adjectives are where Sugawara's theory might be appealing: she notes that these can't take *-ish*, e.g. upper-bound *cleanish* vs. lower-bound **dirty-ish*. But this could fall out from these adjectival roots' status as total and partial (Yoon 1996). Since *clean* is total, it is impossible (or at least odd) to be 'almost dirty': either something is entirely clean (in which case there is no room to distinguish between *clean* and *dirty-ish*), or it is not clean (in which case it is *dirty*, not *dirty-ish*).

mean you can start driving, that I'll give you candy even if we leave in exactly five minutes, that all students that are at least close to being tall (including truly tall ones) will fail the class, and that all diligent and diligent-approximating students will pass the class. The examples (71b) and (71d) are ones where pragmatic or world-knowledge factors encourage the non-strengthened reading: it would be silly to promise candy to someone if you leave in around five minutes but not if you leave in exactly five minutes, and it would likewise be an unusual class indeed if all students approximating diligence are bound to pass, but not the students who are truly diligent. If *-ish* shifted rather than weakened its complement, (71b) and (71d) would have the odd meanings just identified (!). Crucially, (71a) and (71c) are *not* sentences where there is pragmatic motivation for any particular reading of *-ish*. For instance, there would be nothing stranger about a scenario where one can only start driving if the flag is off-pink, than a scenario where one can only start driving if the flag is pink or pink-like—'go!' is culturally associated with green lights, not pink flags. But the 'weak' (rather than 'shifted') meanings are clearly available, showing that they are not just the result of world-knowledge expectations in these examples.

In sum, the lexical meaning of *-ish* broadens the meaning of the predicate it affixes to. This is precisely what the hypothesis under consideration (64) claims that *also* does in sentences like (72).

(72) This comedy is #(also) a tragedy.

The hypothesis in (64) would therefore lead us to expect that *-ish* and *also* should behave in just about the same way. I now turn to showing that this is not the case, and in exactly the way that is expected if *-ish* weakens predicates but *also* does not.

An experiment: comparing *-ish* and *also* in contradictory copular sentences

If the hypothesis in (64) is correct—meaning that *also* in (62b) is removing the contradiction in (62a)/(72) by weakening the meaning of one or both of the predicates—then *also* is semantically identical to *-ish*, at least if we control for the exhaustivity effect associated with *-ish* (which leads *-ish* to be intuited as shifting rather than weakening meanings). Thus, we should be able to create semantically identical minimal pairs out of copular sentences using *also* and *-ish*. We will control for the exhaustivity effect associated with *-ish* by creating a kind of example where the exhaustivity simply does not matter for the judgment.[30]

The experiment is to create a sentence that will be felicitous with *-ish*, but where we only expect *also* to create felicity if the hypothesis in (64), according to which *also* weakens predicates' underlyingly strong meanings, is correct. We will do this by relying on cases where infelicity in basic sentences results clearly from world knowledge, rather than something murkier like conceptual geometry, the lexicon, or indeed exhaustivity on predicates. In light of this, consider (73), keeping in mind that real tigers, being cats, are not marsupials.

(73) #This tiger is a marsupial.

Now consider that so-called 'Tasmanian tigers' are tiger-like marsupials—but not real tigers (the fact that they have the word 'tiger' in their name is not a relevant part of the experiment). (73) is

[30]In fact, whatever exhaustivity effect is associated with *-ish* and alternatives created by structurally removing *-ish* should also come about with *also* given the possibility of structurally removing *also*. As such, it is not impossible under the hypothesis in (64) that *-ish* and *also* would be semantically identical all the way to the exhaustivity effect.

still not an acceptable sentence if it is spoken of a Tasmanian tiger (unless someone uses the word *tiger* to refer to Tasmanian tigers as part of their idiolect). Now consider the judgment with *-ish* in (74), spoken of a Tasmanian tiger.

(74) This tiger-ish animal is a marsupial.

This strikes me as a perfectly sound way of salvaging (73), without requiring any caveat about the term 'tiger' being used idiolectically to refer to Tasmanian tigers.

The acceptability of (74) is unsurprising in light of the proposed meaning of *-ish*. Tasmanian tigers are not part of the set of animals denoted by *tiger*, but they do approximate tigers (due to their stripes; hence, indeed, their name). If *-ish* takes a predicate and makes it denote a superset of its normal extension based on a principle of similarity or approximation, it is unsurprising that Tasmanian tigers could be referred to as *tiger-ish*.

The ground is now set to turn to *also*, and see if it can fix (73) in the same way as *-ish* can in (74). If, as claimed by the hypothesis in (64), *also* takes underlyingly strong predicates and weakens them, it should behave like *-ish* in (74). On the other hand, if my actual claim (63) is correct, and *also* does the reverse (it takes underlyingly weak predicates, which would be strengthened through exhaustification in basic sentences, and 'de-exhaustifies' them so that their underlyingly weak meaning can come to the surface), we do not expect *also* to improve (73). After all, in this particular case, we know that *tiger* is underlyingly strong in the sense of excluding Tasmanian tigers, because this exclusion of Tasmanian tigers comes about from the world knowledge that so-called Tasmanian tigers merely resemble tigers, without actually being tigers; the exclusion of Tasmanian tigers does not come from whatever mechanism strengthens *comedy*, *train*, or *fork* (in my view, exhaustification).

Here, then, is our test:

(75) #This tiger is also a marsupial.

The prediction of the hypothesis under consideration (64) does not go through: *also* cannot fix the contradiction in (73) in the way that *-ish* can in (74). This is unsurprising if *-ish* makes its complement denote a superset of the real tigers, whereas *also* does not change the predicates' underlying meanings. The additive *also*, I claim, can only remove or weaken the *grammatical strengthening* of predicates, but in (73) the contradiction does not come about due to strengthening, but due to the *plain* meaning of the predicates and our world knowledge about them.

Note that, in (75), it happens to be that *also* would need to broaden the meaning of the predicate in the subject rather than the predicate in the *v*P to avoid a contradiction. But, putting aside how exactly this happens (see chapter 3), we know that *also* can in fact circumvent the strengthening of subject-internal predicates, as in sentences like (76) where not only *tragedy* but also *comedy* must have a weak meaning for the sentence to be non-contradictory. That is, both predicates must be compatible with the other as a result of the presence of *also*.

(76) This comedy is also a tragedy.

To be sure, there *is* a way to use *also* to assert of a single individual two predicates like *tiger* and *marsupial* with an empty (real-world) intersection: by embedding the entire phrase under *like*, which effectively plays the same role as *-ish*:

(77) a. #A Tasmanian tiger is a tiger that is a marsupial.
b. #A Tasmanian tiger is a tiger that is **also** a marsupial.
c. #A Tasmanian tiger is **like** a tiger that is a marsupial.
d. A Tasmanian tiger is **like** a tiger that {is, was} **also** a marsupial.

Intuitively, the idea is that, if there were individuals in the intersection of tigers and marsupials, Tasmanian tigers would approximate such individuals. Other similar examples are given in (78)–(79).

(78) a. #A platypus is a duck that is a beaver.
b. #A platypus is a duck that is **also** a beaver.
c. #A platypus is **like** a duck that is a beaver.
d. A platypus is **like** a duck that {is, was} **also** a beaver.

(79) a. #A labradoodle is a labrador that is a poodle.
b. %A labradoodle is a labrador that is **also** a poodle.
c. #A labradoodle is **like** a labrador that is a poodle.
d. A labradoodle is **like** a labrador that {is, was} **also** a poodle.

Note that even with *like*, *also* is absolutely necessary in (77)–(79), so that these predicates pattern very similarly to the other cohyponyms we have focused on (literary genres, colours, and so on). Even if the set-intersection of *duck* and *beaver* is empty, the additive must ensure that neither of these predicates is exhaustified to negate the other, in order to even consider comparing the platypus to what an individual in that set-intersection of ducks and beavers would be.

Other examples of *also* failing to create consistency

Finally, it is also possible to see that *also* cannot create consistency out of underlyingly incompatible predicates based on predicates that are inconsistent in all worlds, rather than due to knowledge about the real world:

(80) a. (i) #Some triangles are **also** squares.
(ii) #This shape is a triangle **and** a square.
b. (i) #Some visible signs are **also** invisible.
(ii) #This sign is visible **and** invisible.
c. (i) #John is a bachelor who is **also** married.
(ii) #John is both a bachelor **and** married.

While this is an important point, it is not clear that the alternative hypothesis in (64) makes any particular prediction here given that the predicates in (80) are incompatible in all worlds (they are logically or mathematically contradictory).

2.3.4 Interim summary

Some cohyponyms happen to refer to concepts that are inherently incompatible (81a) or at least incompatible in the real world (81b). These can be understood as coming from conceptual spaces that are genuinely partitioned, descriptively speaking.

(81) a. #This triangle is (also) a square.
b. #This beaver is (also) a duck.

On the other hand, many cohyponyms that are intuited as mutually exclusive in basic sentences actually come from domains where predicates underlyingly have non-empty intersections. A particular exemplar can belong to more than one concept. For instance, a spork has the properties of both forks and spoons, and can therefore be described as a *fork* that is *also* a *spoon*. The evidence for this comes from formally intersective conjunction (82), additive particles (83), as well as other expressions like *simultaneously* (84).

(82) a. This is both a car and a boat.
b. This is both an epic and a comedy.
c. This is both a fork and a spoon.

(83) a. This car is #(also) a boat.
b. This epic is #(also) a comedy.
c. This fork is #(also) a spoon.

(84) a. This car is #(simultaneously) a boat.
b. This epic is #(simultaneously) a comedy.
c. This fork is #(simultaneously) a spoon.

Evidence for the formally intersective status of the conjunctions in (82) will be given in chapter 5. In the present chapter, I dwelt on *also*, and in particular considered the idea that the *also* observed in (83) could perhaps be non-intersective. My actual proposal is that *also* is always intersective, and in (83), its ability to make some cohyponyms mutually compatible shows us that the cohyponyms refer to concepts that overlap. The alternative, according to which *also* takes cohyponyms corresponding to concepts that do *not* underlyingly overlap, but weakens their meaning much like *-ish* does, does not hold water empirically. If *also* could weaken predicates, we would expect it to remove contradictions even when we know with certainty (due to independent factors like world knowledge) that the concepts referred to by the cohyponyms do *not* overlap. This is not the case:

(85) a. #A Tasmanian tiger is a tiger that is also a marsupial.
b. #A platypus is a duck that is also a beaver.

It is not that language is unable to weaken predicates per se; *-ish* does precisely this (86). (85) and (86) do not contain perfect minimal pairs only because *-ish* creates adjectives, requiring some (irrelevant, I think) syntactic changes.

(86) a. A Tasmanian tiger is a tiger-ish animal that is a marsupial.
b. A platypus is a duck-ish, beaver-ish animal.

Thus, (83) truly shows that the sets denoted by cohyponyms like *car* and *boat* or *fork* and *spoon* have a non-empty intersection.

If cohyponyms like *fork* and *spoon* are indeed non-contradictory at the conceptual level (that is, they intersect), the question becomes what creates the illicitness of (87).

(87) #This fork is a spoon.

I suggested that the strengthening of cohyponyms is the result of exhaustivity. Thus, at first blush, a sentence like (88a) has the LF in (88b). It comes with alternatives (88c) generated by replacing a predicate with its cohyponyms. The resulting truth conditions (88d) are such that the asserted predicate has become semantically incompatible with its cohyponyms: there may be a non-empty intersection between cohyponyms at the level of lexical–conceptual meaning, but language strengthens the predicates so that they exclude the individuals in the intersection.

(88) a. This is a fork.
b. Exh_{ALT} [this is a fork].
c. $\text{ALT} = \{\text{this is a fork, this is a spoon, this is a knife}\}$
d. $[\![\text{Exh}_{\text{ALT}}\ [\text{this is a fork}]]\!]$
$= 1$ iff this is a fork & this is not a spoon & this is not a knife.

This is not an effect stemming from contrastive focus of predicates; contrastive focus operates *on top of* this exhaustification:

(89) The car is not required to be ITALIAN, it's required to be GREEN.

We saw that (89) involves the local strengthening of *green* vis-à-vis other colour terms separately from the strengthening of *green* vis-à-vis its contextual, contrastive alternative *Italian*.

My proposal exemplified by (88) is in effect a 'neo-structuralist' theory of language. Like de Saussure (2011[1916]), my claim is that strengthening is at play not just in grammatical vocabulary like the meaning of plurals (Sauerland et al. 2005), but that *the same strengthening mechanism* (exhaustivity—for me, but evidently not for de Saussure) is involved in the meaning of *predicates*. This goes for all predicates, not just predicates that form entailment scales.

My suggestion that cohyponymic exclusivity is the result of exhaustivity requires two claims about Exh. The first is that Exh is obligatory with cohyponyms; otherwise, cohyponymic exclusivity would be optional and non-contradictory meanings would be available with pseudo-repetitions. The second is more novel and surprising: the Exh that cohyponyms require is necessarily syntactically quite close to the cohyponym. I turn to showing this in section 2.4. I will refer to these twin properties of the Exh creating cohyponymic exclusivity as *controlled exhaustivity*. Cohyponyms 'control' Exh in requiring its presence in the sentence, and, moreover, dictating (in approximate terms) its syntactic locus.

In later chapters, I will fine-tune the proposal in (88) in two ways. In chapter 6, I will argue, with some extrapolation, that (88b/d) is actually not correct, because Exh is too syntactically far from the cohyponym it associates with. Second, in in chapter 4, I will claim that the alternatives are not actually created from a given set of cohyponyms. To properly understand the nature of the alternatives, I will introduce the notion of 'predicational jurisdictions,' as a replacement for the notion of cohyponymy that has so far been of central empirical and theoretical importance.

2.4 Cohyponyms and the syntactic distribution of exhaustivity

The meaning of cohyponyms involves exhaustivity. This proposal is interesting in its own right for our understanding of predicates and their relationship to concepts; it is far from trivial that the meaning of predicates in basic sentences (i.e., without conjunction, additives, or the like) is not a

one-to-one match with the concepts referred to by the predicates. This is a new claim, presumably with general consequences for work in psychology and cognitive science that uses natural-language data to research concepts (see e.g. Margolis & Laurence 2021).

What we turn to now, however, is a consequence of my claim of a rather different sort. Indeed, if the exhaustivity account of cohyponymic exclusivity is accepted, it comes with an important consequence not just for the meaning of predicates, but also for our understanding of exhaustivity as a grammatical phenomenon.

As described by Chierchia et al. (2012), Exh is essentially unconstrained. It is an optional operator (an adjunct) which language users can choose to put at any *t*-type node in a sentence. Call this *free exhaustivity*.

(90) **Free exhaustivity**:
   a. Exh is optional.
   b. Exh is syntactically free (although dispreferred in DE contexts).[31]

Among other things, this means that there is no direct formal link between Exh and the expression(s) in its prejacent that contribute to forming the set of alternatives that Exh takes. That is, while Exh operates over a set of alternatives, which are shaped by alternative-triggering expressions, these expressions have no grammatical link to Exh. Alternative-triggering expressions do not require the presence of an Exh operator, and if one is present, they do not directly govern where it appears.

In the examples in (91), for example, there is an alternative-triggering expression, namely *or*. The lexical meaning of *or* is inclusive, but *or* is regularly strengthened to being exclusive, as a result of the negation of an alternative with *and*: 'P $\vee$ Q' is strengthened to 'P $\vee$ Q, but not P $\wedge$ Q.' However, the Exh operator that performs this strengthening can optionally appear within the same clause as *or* (91a) or not (91b). In fact, in (91b), it is not even clear if there is an Exh at all: if there was one, it would have no discernable effect on the meaning of *or*, since it is in a DE environment.

(91) a. If [Exh_{ALT} [you take salad **or** dessert]], you pay $20; but if you take both there is a surcharge.
   b. If [you take salad **or** dessert], you'll be real full.
   (Chierchia et al. 2012:2306)

Thus, there is surely validity to the notion of free exhaustivity. For some alternative-triggering expressions like *or*, it is entirely correct to claim that there is no direct grammatical link between them and Exh. They may or may not be exhaustified, and if they are, they have no say on where the Exh operator is located in the syntax.

What I will show in this section is that the exhaustification of cohyponyms does not behave in this way. Indeed, Exh with cohyponyms is both obligatory and necessarily local to the predicate. While I will not give a formal characterization of the locality requirement until chapter 6, I will informally discuss Exh as being 'ultra-local' because some examples will suggest that even a locality requirement like 'Exh must be in the same minimal clause as the cohyponym' is too weak (in chapter 6, I will in fact suggest that Exh must be in the XP of the cohyponym). In what follows, I will make the case for the obligatory presence of Exh and its locality constraint in turn.

[31] See Fox & Spector 2018 for elaboration on this.

2.4.1 Exh is obligatory with cohyponyms

When a weak alternative-triggering expression like *or* is located in a DE context, Exh has no visible effect, and it could perhaps not be present at all. But there are reasons to think that Exh is optional even in upward-entailing contexts—that is, where its presence *would* have a noticeable effect. Consider the discourse in (92). While the first sentence S1 is initially parsed as meaning that Aisha ate some but not all the apples, upon the second sentence S2, S1 is apparently freely reanalyzed as not involving exhaustification. Indeed, (92) is not intuited as involving the speaker correcting themselves; S2 simply provides additional information. It is an elaboration, not a correction.

(92) Aisha ate some apples. In fact, she ate all of them.

Thus, at least in the right discourse conditions, *Aisha ate some apples* can optionally not involve the strengthening of *some* to *some but not all*.

Now consider pairs of cohyponyms in a similar discourse:

(93) a. #The flag is white. In fact, it's green.
b. #This play is a comedy. In fact, it's a tragedy.

The judgment for the discourses in (93) is completely different from (92). The S2s in (93) do not serve as an elaboration of the S1s, but in fact as a self-correction (hence the #). (93a) does not successfully convey that the flag is partly white and partly green, and (93b) does not convey that the play is a tragicomedy.

This is not a trivial finding; language could have worked otherwise. The predicates in the S1s and S2s in (93) are all underlyingly mutually compatible (whether due to the existential meaning of colour terms, or the conceptual overlap between literary genres), and only become incompatible due to the presence of Exh. If Exh was optional, the discourses in (93) would be acceptable, with the S2 specifying that the individual flag or play is not just in the set denoted by *white* or *comedy*, but more specifically, in the intersection of that set and the set denoted by *green* or *tragedy*.

To be sure, the discourses in (92) and (93) are not completely parallel. In (92), S2 asymmetrically entails S1 (*all* is stronger than *some*). Thus, it could be claimed that, while members of the scale $\{$*some*, *most*, *all*$\}$ must not be exhaustified in a given sentence, they must be exhaustified *somewhere* in the discourse. However one wants to formalize this, the intuition would be that (92) only allows S1 to be de-exhaustified due to the presence of a stronger scale-mate appearing in S2, which is itself exhaustified. Indeed, it is difficult to get S2 not to be exhaustified; in (94), *most* still means *not all*.

(94) Aisha ate some apples. In fact, she ate most of them.
$\rightsquigarrow$ Aisha ate most but not all of the apples

In a sense, it is overstating things to claim that Exh is truly optional; its non-appearance (perhaps more accurately: its post-hoc removal) from S1 in discourses like (92) and (94) is only licensed if subsequent material would be contradictory with an exhaustified version of the sentence.

In contrast, the S2s in (93) do not entail the S1s; the non-exhaustified versions of the sentences are logically independent of one another. What is more, if the S2s in (93) are exhaustified, the two sentences are irreparably contradictory. Removing Exh from S1 cannot improve the discourse because the exhaustified S2 still contradicts it.

Yet, if the idea for (92) and (94) is that the S1s are de-exhaustified in order to avoid the creation of a contradiction with later material, it is really not clear why this de-exhaustification could not take place in both S1 *and* S2 in (93). Indeed, even if this putative contradiction-avoidance principle would only apply to the S1s in (92) and (94), what would stop it from applying to both the S1s and the S2s in (93)? While my examples are not perfect minimal pairs, the basic point stands: it is not possible to avoid exhaustifying cohyponyms (93), even if this creates a contradiction.

In fact, the examples in this subsection are just bisentential equivalents of the basic pseudo-repetition examples we have seen, like (95).

(95) a. #The white flag is green.
b. #This comedy is a tragedy.

No contradiction would be intuited in (95) if Exh was optional.

The claim that Exh *is* or *can be* obligatory is not new (e.g., Magri 2009, Chierchia 2013, Bade 2016, Bar-Lev 2018). But in what follows, I show that cohyponyms not only require Exh, but also dictate its syntactic position.

2.4.2 Exh must be ultra-local to cohyponyms

In this section, we will observe that cohyponyms require Exh to be in close syntactic proximity to them. In particular, we will see that in all cases where a difference in meaning is predicted to be intuited according to whether Exh takes scope globally or locally to the cohyponym, the meanings that would be obtained from a global Exh are not intuited. Most of the examples involve mono-clausal sentences, so 'locally' has stronger meaning than just 'in the same clause.' Empirically, the observation in all cases is that cohyponyms are interpreted as incompatible regardless of the syntactic structure or semantic environment they are in (unless there is conjunction, an additive, etc., as already discussed). The main exception to this is sentential negation, which I will discuss both in this section and in chapter 6.

I will show that all the intuited meanings can be obtained if Exh is stipulated to be necessarily very close to the cohyponym. For our purposes, it will suffice to understand 'very close' as meaning 'in the XP of the cohyponym' (e.g., anywhere in the AP of an adjective); we will return to this in chapter 6. I will refer to this as an *ultra-local* Exh, and I will schematize this as in (96) for colour terms (putting aside the question of whether Exh is necessarily ultra-local in this particular example):

(96) $[\![\text{The flag is } [_{\text{AP}}\ \text{Exh}_{\text{ALT}}\ \text{green}]]\!] = 1$ iff the flag is $[\text{green}_{\exists}\ \&\ \text{not red}_{\exists}\ \&\ \text{not white}_{\exists}]$.

(96) has an Exh operator taking a colour term, and nothing else, as its prejacent. But of course, Exh, as defined in chapter 1, is a propositional operator. To have Exh take a proposition, we can claim that Exh's prejacent is actually composed not just of the colour term, but also some trace; the trace and the colour terms jointly yield a proposition. This is what Heim & Kratzer (1998:§8.5) suggest as a possibility for adjectives as well based on different empirical data. They suggest that APs can have a semantically vacuous PRO subject that moves out of Spec-AP to create a trace. Under this view, the syntax of (96) is actually (97); Exh's prejacent is propositional, even though Exh is internal to the AP.

(97) The flag is $[_{AP}$ PRO λ_1 $[_{AP}$ Exh_{ALT} t_1 green$]]$.

While assuming this structure throughout this thesis, I will abbreviate it as in (96) for ease of presentation.[32]

We now turn to various kinds of empirical evidence suggesting that Exh is necessarily ultra-local to cohyponyms.

Predicating two cohyponyms of a single referent

I start with precisely the kinds of examples that motivated cohyponymic exclusivity in the first place, namely pseudo-repetitions. The fact that language does not allow pseudo-repetition *even in a single sentence* (98) is in fact unanticipated for the Exh account of cohyponymic exclusivity.

(98) a. #The green flag is white.
≈ the entirely green flag is entirely white
b. #This comedy is a tragedy.
≈ this non-tragic comedy is a non-comedic tragedy

If Exh could take scope anywhere, the sentences in (98) would in fact be non-contradictory. It would be possible for Exh to scope globally, as shown in (99) for (98a). From this global position, Exh's prejacent would entail both the whiteness and greenness of the flag. Since Exh does not exclude alternatives that are entailed by its prejacent (it only excludes non-weaker alternatives), neither colour would be excluded. Other colours would be excluded, but not green or white. Thus, what would result is the non-contradictory meaning in (99a), rather than something like (99b), as needed to have a contradiction. For ease of exposition, (99) only shows the predicative adjective *white* as having alternatives, but I believe that *green* does too. We will return to this both shortly below and in chapter 3; nothing hinges on this for our present purposes.

(99) $[\![Exh_{ALT}$ [The green flag is white]$]\!]$

a. $= 1$ iff the green$_\exists$ flag is white$_\exists \wedge \neg$[the green$_\exists$ flag is red$_\exists$].
$\Rightarrow$ no contradiction

b. $\neq 1$ iff $\begin{cases} \text{the green}_\exists \text{ flag is white}_\exists \wedge \\ \neg[\text{the green}_\exists \text{ flag is green}_\exists] \wedge \Rightarrow \text{contradiction} \\ \neg[\text{the green}_\exists \text{ flag is red}_\exists] \end{cases}$

This example involves a definite subject, which triggers a uniqueness presupposition; this is not a relevant factor to my argumentation. The contradictory meaning of pseudo-repetitions is observed with other types of subjects as well (as discussed in section 2.2.2), and prior to exhaustification, the uniqueness presupposition is merely that there is a unique *at least partly* green ('green$_\exists$') flag. This is not incompatible with that flag also having other colours, so it cannot be blamed for the intuition of a contradiction in (98a).

In order for the cohyponyms in sentence-internal pseudo-repetitions to be strengthened irrespective of one another, Exh must be syntactically constrained. It must appear locally to each colour term, so as not to take the other one in its scope:

[32] A second possibility is to allow Exh to take non-propositional elements—specifically, predicates. A type-flexible Exh would rely on a generalized notion of entailment (cf. Fox & Spector 2018:7 fn. 12).

(100) ⟦The [Exh$_{\text{ALT}}$ green] flag is [Exh$_{\text{ALT}}$ white]⟧
= 1 iff the [green$_{\exists}$ & not white$_{\exists}$ & not red$_{\exists}$] flag is [white$_{\exists}$ & not green$_{\exists}$ & not red$_{\exists}$].
≈ the entirely green flag is entirely white
⇒ contradiction

Based on (98) alone, we would conclude that, as long as one or the other cohyponym is strengthened irrespective of the other, we would obtain a contradiction. That is, (100) is not the only possible LF that would derive a contradiction; it could be that one of the colour terms has a local Exh, while the other is associated with a global Exh or no Exh at all. Naturally, if it was not the case that both cohyponyms have an ultra-local Exh, we would need to explain why they behave asymmetrically.

However, there is in fact data suggesting empirically that both cohyponyms must be exhaustified ultra-locally (that is, without taking the other into account). Consider the following paradigm, in which I explicitly weaken one or both colour terms through the adverb *partially*:[33]

(101) a. The partially white flag is #(also) green.
b. The white flag is #(also) partially green.[34]
c. The partially white flag is (also) partially green.

The fact that (101a) requires *also* shows that the predicational adjective *green* is exhaustified independently of the attributive *partially white*: if *green* had no Exh or had an Exh which scoped above *partially white*, no contradiction would arise. Likewise, the fact that (101b) requires *also* shows that the attributive adjective *white* is exhaustified independently of *partially green*, for the same reason. Finally, (101c) is an important control, where we observe that *also* is no longer required if both adjectives are explicitly made weak (the intuition for the sentence is that the presence of *also* is preferred, but the sentence without *also* is not contradictory like (101a–b)). Thus, both attributive and predicational cohyponymic adjectives can be shown to be exhaustified independently of the other. If this is the case in (101a–b), it's not clear on what grounds this would *not* be the case in sentences without *partially* like (98)/(100).

Cohyponyms co-occurring with another scope-bearing element

Another way to observe the locality requirement on Exh is to place a cohyponym in the same clause as some other scope-bearing element. The goal is to see whether Exh can scope above both the cohyponym and the other element, or if it must scope above only the cohyponym. To test this, let's consider examples with a universal quantifier:

(102) a. Every flag is green.
≈ every flag is entirely green

[33]The fact that (101c) does not require *also* emphasizes once again that there is nothing inherently wrong with pseudo-repetitions; when they are infelicitous, it is due to the meaning of cohyponyms, rather than the verb *to be*, something about information structure, or the like.

[34]The judgment for this example is not entirely sharp, because it is easy to think of the sentence as meaning that the background of the flag is entirely white, and it has some green parts *over* this background. See chapter 5 for discussion of background and foreground with colour terms. Crucially, the judgment given for (101b) does hold for a flag that has a fairly large green part that is not conceptualized as being superimposed on a white background (e.g., a half-white, half-green flag).

b. Every play in this classroom is a tragedy.
≈ every play in this classroom is a non-comedic tragedy

Yet, global Exh, scoping above both *every* and the cohyponym, yields (103) for (102a):

(103) $[\![\text{Exh}_{\text{ALT}}\ [\text{every flag is green}]]\!] = 1$ iff $\begin{cases} \text{every flag is green}_{\exists} \wedge \\ \neg[\text{every flag is white}_{\exists}] \wedge \\ \neg[\text{every flag is red}_{\exists}] \end{cases}$

The truth conditions in (103) would be met if (for example) all flags were only half green, with some flags half red and the rest half white. This clearly goes against the intuited meaning. On the other hand, the right truth conditions are obtained by having Exh scope below *every*—that is, ultra-locally to the cohyponym:

(104) $[\![\text{Every flag is } [\text{Exh}_{\text{ALT}}\ \text{green}]]\!]$
$= 1$ iff every flag is $[\text{green}_{\exists}$ & not $\text{white}_{\exists}$ & not $\text{red}_{\exists}]$.
≈ every flag is entirely green

Cohyponyms with a scalar expression

In a similar vein, we can also observe the locality requirement on Exh by considering cases where a cohyponym co-exists with another alternative-triggering expression, such as an existential quantifier (105). Such elements are usually also exhaustified.

(105) a. Some flags are green.
≈ some but not all flags are entirely green
b. Some plays are tragedies.
≈ some but not all plays are non-comedic tragedies

As we will see, if Exh could scope above both the cohyponym and the other alternative-triggering expression *some*, we would obtain the wrong results.

When a single Exh operator has a prejacent with more than one alternative-triggering expression, I assume that the alternatives it takes comprise all the possible sentences that can be obtained by replacing one or more alternative-triggering expression with one of its alternatives (Sauerland 2004). Thus, for (105a), the alternatives are in (106).

(106) $\text{ALT} = \left\{ \begin{array}{c} \text{some flags are green,} \\ \text{some flags are white,} \\ \text{some flags are red,} \\ \text{all flags are green,} \\ \text{all flags are white,} \\ \text{all flags are red} \end{array} \right\}$

The question is which of these alternatives Exh actually excludes. It cannot exclude all of them. If it did, we would obtain the truth conditions in (107).

(107) $[\![\text{Exh}_{\text{ALT}}\ [\text{some flags are green}]]\!] = 1$ iff $\begin{cases} \text{some flags are green}_{\exists} \wedge \\ \neg[\text{some flags are white}_{\exists}] \wedge \\ \neg[\text{some flags are red}_{\exists}] \wedge \\ \neg[\text{all flags are green}_{\exists}] \wedge \\ \neg[\text{all flags are white}_{\exists}] \wedge \\ \neg[\text{all flags are red}_{\exists}] \end{cases}$

This is certainly not the meaning we are after. (107) means that there are flags that are partly green, not all flags are partly green, and no flags are of any colour other than green. Thus, according to (107), there must be at least one colourless flag. Again, it means that not all flags have a green part, and no flag has a part that is any colour other than green.

(107) will have to be salvaged by appealing to Innocent Exclusion (see chapter 1). It is not entirely obvious that Innocent Exclusion can kick in to ensure that some alternatives in (106) would not be excluded. After all, the meaning in (107) is not a *logical* contradiction: from a purely logical perspective, there could be colourless flags. But (107) is certainly incompatible with world knowledge in some sense; presumably, we conceptualize of all surfaces as having a colour (perhaps with transparency counting as a colour). Either way, there is something akin to a 'contradiction' in (107), due to the exclusion of the subset of alternatives in (108).

(108) {some flags are white, some flags are red, all flags are green}

Appealing to Innocent Exclusion, we can try to salvage (107) by simply not excluding these alternatives.[35] Since these alternatives are not innocently excludable, they are not excluded at all, leaving us with the truth conditions in (109) instead of (107).

(109) $[\![\text{Exh}_{\text{ALT}}\ [\text{some flags are green}]]\!] = 1$ iff $\begin{cases} \text{some flags are green}_{\exists} \wedge \\ \neg[\text{all flags are white}_{\exists}] \wedge \\ \neg[\text{all flags are red}_{\exists}] \end{cases}$

While this no longer entails that some flags lack a colour entirely, this is not a good result either. First of all, (109) has the same problem as (103): it only means that some flags are *partly* green, not *entirely* green. What is more, it does not strengthen *some* to mean 'not all.' The result is that some and maybe all flags are partly (maybe entirely) green, rather than some but not all flags being entirely green. Thus, our attempt in (109) at modifying the unwanted result of (107) via Innocent Exclusion has failed.[36]

In contrast, we can obtain the right truth conditions by having two Exh operators in the sentence, one that is immediately above the cohyponym and its alternatives created by replacing the cohyponym with its peers, and another for the non-cohyponymic scalar element (here *some*). While the Exh operator associated with the cohyponym is obligatory and necessarily local, the exhaustification of *some* is an instance of *free exhaustivity*: this Exh is optional and able to scope anywhere. This free Exh operator takes alternatives created by replacing the non-cohyponymic scalar element with its scalemates, but not replacing the cohyponym with anything (as if the cohyponym has been

[35] I will return to this in chapter 4, where I will in fact follow Magri (2009) in assuming that Exh does not, in fact, take any sort of world knowledge into account. Nothing hinges on this for now.

[36] Nothing improves if Exh also has the Innocent Inclusion property. With Innocent Inclusion, Exh would end up *asserting* the alternatives in (108). The resulting meaning would be that for every colour, there is at least one flag that is at least partly of that colour; and, moreover, *all* flags are at least partly green.

rendered 'inactive' by having already been used by the lower Exh operator). Thus, we have the LF in (110a), which has two Exh operators; these take the two distinct sets of alternatives in (110b), creating the truth conditions in (110c).

(110) a. $\text{Exh}_{\text{ALT-2}}$ [some flags are [$\text{Exh}_{\text{ALT-1}}$ green]].

b. (i) $\text{ALT-1} = \{\text{green, white, red}\}$

(ii) $\text{ALT-2} = \{\text{some flags are [Exh}_{\text{ALT}}\text{ green], all flags are [Exh}_{\text{ALT}}\text{ green]}\}$

c. $[\![(110a)]\!] = 1$ iff $\left\{ \begin{array}{l} \text{some flags are } \left(\begin{array}{c} \text{green}_{\exists}\ \& \\ \text{not white}_{\exists}\ \& \\ \text{not red}_{\exists} \end{array} \right) \wedge \\ \neg[\text{all flags are } \left(\begin{array}{c} \text{green}_{\exists}\ \& \\ \text{not white}_{\exists}\ \& \\ \text{not red}_{\exists} \end{array} \right)] \end{array} \right.$

This is the desired meaning. To reiterate, the upshot is that the Exh associated with the cohyponym (here *green*) and its set of alternatives has to be local to the cohyponym, scoping below other alternative-triggering expressions (here *some*). The other Exh (the global one) in the sentence (110a) is an independent operator due to the presence of *some*, which has nothing to do with the strengthening of cohyponyms or the notion of necessarily ultra-local exhaustivity being developed in this section.

Cohyponyms in downward-entailing environments

Exhaustivity normally optionally disappears in DE contexts, depending on whether Exh is located above or below the DE operator. Returning to an example we have already seen a few times, disjunction in a DE context can optionally be interpreted as exclusive or inclusive.

(111) a. If you take salad or dessert, you'll be really full.

b. If you take salad or dessert, you pay $20; but if you take both there is a surcharge. (Chierchia et al. 2012:2306)

On the other hand, cohyponyms in most DE contexts are still interpreted as mutually exclusive. To show this thoroughly, I will go through all the environments provided by Chierchia (2004:§2.2) as contexts where implicatures standardly disappear. Chierchia's examples are based empirically on contexts where *or* is readily intuited as inclusive.[37]

While cohyponymic exclusivity is maintained in *most* DE environments, there are some cases where cohyponymic exclusivity disappears (one where this is clearly the case, and a few where the data are less clear). The latter are specifically *negative* DE environments. As such, the environments in which cohyponymic exclusivity disappears is not well characterized as 'DE environments' generally; rather, it is specifically negative environments. We will return to negative environments in chapter 6, where I will claim that these can be understood as still involving exhaustification if we adopt the presuppositional Exh of Bassi et al. (2021); as such, while they will appear like an empirical hurdle for the claim that Exh is always obligatory and ultra-local with cohyponyms, this is not actually the case.

[37] In addition to the examples below, Chierchia (2004:54) also has 'irrealis' environments, where he relies on Italian data with main-clause subjunctives. I will stick to English data, so I put this particular environment aside.

In this subsection, I discuss non-negative DE contexts like the antecedents of conditionals, which virtually all display cohyponymic exclusivity (there are a few unclear cases we will get to); we will move on to negative environments in the following subsections. The first non-negative DE environment we consider is the antecedent of conditionals; we already saw in (111a) that *or* can be inclusive in this context. Cohyponyms retain their strong interpretations, however.

(112) a. If the flag you wave is green, Aisha will sing the national anthem.
$\Rightarrow$ green $\approx$ all green
b. If you read a comedy, the hero lives.
$\not\rightsquigarrow$ if you read a comedy or tragicomedy, the hero lives

On the old-fashioned definition of comedies and tragedies as defined largely in terms of whether the protagonist lives or dies, (112b) provides a particularly compelling datapoint: a tragicomedy would not necessarily involve the protagonist living, given that it includes elements of tragedies too. But (112b) is clearly good.

One empirical puzzle about conditionals is the observation that on the sentence in (112a), it seems that Aisha would not know what to do if the flag was only *partly* green (rather than knowing that it is not the case that she must sing), in the same way as she would not know what to do if someone said (113) and you only raised half the flags.[38]

(113) If you raise the flags, Aisha will sing.
$\rightsquigarrow$ if you wave all the flags, Aisha will sing.

This could come about because conditionals are often strengthened to biconditionals in natural language (e.g., Geis & Zwicky 1971, Horn 2000, von Fintel 2001, Herburger 2015). (112a) means not only that Aisha will sing if the flag is green, but also that she will not sing if the flag is not green. As we are about to see in the next subsection, *not green* means 'not green at all' rather than 'not entirely green.' On the right theory of the strengthening of *if* to *iff*, the biconditional only refers to cases where the flag is all green (with the Exh associating with *green* situated below *if*) or cases where it is not green at all (the negation), remaining mum on partly-green cases.[39]

A second environment given by Chierchia (2004) as preferably not involving exhaustification is the restrictor of *every*:

(114) Every student who wrote a squib or made a classroom presentation got extra credit.
(Chierchia 2004:50)

Cohyponyms in this context retain their exclusivity.[40]

[38] As I will discuss in chapter 5, the universal meaning of colour terms in positive sentences, together with the fact that (as discussed below in the current chapter) they are interpreted as negated existentials in negative sentences, is describable as a homogeneity effect. Many descriptions of homogeneity claim that non-homogeneous cases (in the case of colour terms, objects which are only partly of a given colour) give rise to neither the logical value 1 nor 0; they are undefined. On this approach, the intuition that Aisha would not know what to do if the flag was partly green would presumably be understood as emerging from the antecedent lacking a truth value in such a scenario.

[39] It may seem strange to claim that the biconditional makes reference to an antecedent of the form *the flag is not green* with no Exh. In fact, there could very well be an Exh there, if we adopt the presuppositional Exh of Bassi et al. (2021); see chapter 6.

[40] Recall that the examples above with *every* (102) had a cohyponym in the scope of *every*, not its restrictor, so we have effectively found that cohyponymic exclusivity is observed in both the scope and restrictor of *every*.

(115) a. Every flag that is green was made in Turkey.
⇒ green ≈ all green
b. Every comedy that Aisha wrote was badly received.
⇏ every comedy or tragicomedy that Aisha wrote was badly received

Third, Chierchia (2004:53) gives generic (non-episodic) statements as DE environments, noting that "the behavior of generics can be reduced to that of the universal quantifier (and of conditionals)."

(116) A linguist or a philosopher doesn't easily give in. (Chierchia 2004:52)

Generic statements do not license a weaker meaning for cohyponyms:

(117) a. A green flag doesn't easily break.
⇒ green ≈ all green
b. A tragedy entails the death of the protagonist.
⇏ a tragedy or tragicomedy entails the death of the protagonist

Fourth, Chierchia shows that implicatures normally disappear in comparatives:

(118) Theo is taller than Bill or John. (Chierchia 2004:53)

This is not the case for cohyponymic exclusivity.

(119) a. The Canadian flag is bigger than any green flag.
⇒ green ≈ all green
b. Aisha's poem is longer than any comedy on this shelf.
⇏ Aisha's poem is longer than any comedy or tragicomedy on this shelf

The same goes for the complements of verbs of comparison like *prefer*.

(120) I prefer Theo to John or Bill. (Chierchia 2004:53)

(121) a. Aisha prefers Canadian flags to green flags.
⇒ green ≈ all green
b. Aisha prefers newspaper articles to comedies.
⇏ Aisha prefers newspaper articles to comedies or tragicomedies

Fifth, Chierchia lists *before*- and *without*-clauses.[41] *Before*-clauses invite a weak reading for *or* (122) but not for cohyponyms (123).

(122) John arrived before Paul or Bill. (Chierchia 2004:53)

(123) a. The Canadian flag was flown before the green flag.
⇒ green ≈ all green
b. Aisha wrote a tragedy before she wrote a comedy.
⇏ Aisha wrote a tragedy before she wrote a comedy or tragicomedy

[41] While at least *without*-clauses are clearly negative-flavoured, I present them in this subsection because Chierchia lists them together with *before*-clauses and because, as far as cohyponymic exclusivity is concerned, they pattern with positive DE environments anyway.

The same goes for *without*-clauses.

(124) John will come without pen or notepads. (Chierchia 2004:53)

(125) a. Aisha will come without a green flag.
$\Rightarrow$ green $\approx$ all green
b. Aisha will pass her English exam without having read a comedy.
$\not\rightarrow$ Aisha will pass her English exam without having read a comedy or tragicomedy

Chierchia (2004:54) also notes that imperatives might be another DE environment:

(126) Get me Paul or Bill. (Chierchia 2004:54)

These do not license weak readings of cohyponyms either:[42]

(127) a. Get me a green flag.
$\Rightarrow$ green $\approx$ all green
b. Write me a comedy.
$\not\rightarrow$ write me a comedy or tragicomedy

There are a few positive DE environments where the judgment is less sharp for cohyponyms, but still strongly lean toward the generalization that DE environments do not remove cohyponymic exclusivity. First, Chierchia lists questions as an environment where implicatures disappear:

(128) A: Did John or Paul arrive? (Chierchia 2004:54)
B: (i) Yes, they both did.
(ii) #No, they both did.

Comedy-type cohyponyms (i.e., taxonomic predicates whose lexical entries do not involve part-quantification) remain strong in questions:

(129) A: Is this a comedy?
B: (i) #Yes, it's a tragicomedy.
(ii) No, it's a tragicomedy.

Colour-type cohyponyms are less clear. (130) tests the adjectival *green* in both predicative and attributive position. In attributive position, they seem to remain strong; the judgment is less sharp in predicative position.

(130) a. A: Is this a green flag?
B: (i) #Yes, it's green and white.
(ii) No, it's green and white.
b. A: Is the flag green?
B: (i) ??Yes, it's green and white.
(ii) No, it's green and white.

[42] I occasionally find (127b) less sharp than (127a); this is probably because the line between comedies and tragicomedies is not particularly clear or objective. If a professor says (127b) and a student writes a tragicomedy, the professor might accept this, but this would be lenient behaviour on their part.

Another unclear environment is under modals of permission. Chierchia notes that implicatures normally disappear here.

(131) It is permitted/legal to smoke or drink. (Chierchia 2004:54)

Again, *comedy*-type cohyponyms still display cohyponymic exclusivity here:

(132) It is permitted/legal to watch a comedy.
$\not\rightsquigarrow$ it is permitted/legal to watch a comedy or tragicomedy.

But colour terms are less sharp:

(133) It is permitted/legal to wave green flags.
$\Rightarrow$ green $\approx$ all green (?)

I suspect this is a so-called non-maximality effect (see chapter 5), whereby universal quantification can be interpreted as existential in certain pragmatic contexts. On this view, speakers who interpret *green* as weak in (133) do so because the first idea that comes to mind is that a rule or law would ban the colour green on flags, rather than specifically banning all-green flags. If the ban is on the colour green on flags, it follows that what matters pragmatically is whether a flag has any green on it. Speakers for whom the first scenario to come to mind is a ban on entirely green flags (perhaps due to some symbolism associated with such a flag) presumably intuit *green* universally in (133).

In sum, putting aside the small handful of less clear cases, cohyponymic exclusivity is intuited in all non-negative DE environments.

Cohyponyms under sentential negation

Negative DE environments are the main systematic exception to cohyponymic exclusivity (in sentences without conjunction or an additive, of course). The most important datapoint consists of sentential negation, where cohyponymic exclusivity is straightforwardly not intuited. Other negative DE environments involve less clear judgments (but generally lean toward displaying cohyponymic exclusivity), which I essentially leave as an open question for this thesis. The data in this section initially appear to be a major exception to the generalization that cohyponyms require an Exh operator in their immediate vicinity (or perhaps shows us that in some cases, 'their immediate vicinity' is defined so as to allow a DE operator to intervene between them and Exh). I reiterate that I will argue in chapter 6 that this can just as well be understood as showing that Exh is presuppositional: by the definition of Exh given by Bassi et al. (2021), it is expected to be vacuous in negative (not DE, but specifically *negative*) environments. If so, the data in the current subsection, while interesting, do not constitute a counterexample to the claim that there is always an Exh operator taking scope more or less immediately above cohyponyms. Until chapter 6, however, I will keep assuming the standard definition of Exh, and therefore write as if there was simply no Exh in such environments.

Sentential negation is the main and clearest exception to the rule that cohyponyms remain strong even in DE contexts. Let's take colours and genres one by one, starting with colour terms. Under negation, colour terms do not have the universal meaning we have been noting elsewhere:

(134) This flag is not green.

$\not\approx$ green $\approx$ all green

The intuition for (134) is that the sentence is true if flag is not green at all, rather than in any scenario where the flag is not entirely green. Descriptively, what we observe with colour terms can be labelled as a homogeneity (all-or-nothing) effect (Löbner 2000): *x is green* means it is all green, whereas *x is not green* means it is not green at all, and neither sentence is true if *x* is only partly green. I will engage with the homogeneity literature in chapter 5.

On our current assumptions, the only way to obtain the meaning in (134) is to claim that there is simply no Exh, whether above or below the negation.[43] If there was an Exh scoping below negation, the meaning of (134) would be the very weak meaning just identified as not intuited; the sentence would be true as long as the flag is not entirely green.

(135) $[\![\text{not } [\text{Exh}_{\text{ALT}} \text{ [the flag is green]}]]\!] = 1$ iff $\neg[\mathsf{green}_{\exists}(f) \wedge \neg\mathsf{white}_{\exists}(f) \wedge \neg\mathsf{red}_{\exists}(f)]$.
$\approx$ the flag is not only/entirely green

Just to confirm my claim that these are not the right truth conditions, consider a sentence like (136). The flag of Ireland is a tricolour with a green part, but (136) is not a true sentence.

(136) #The flag of Ireland is not green.

As for a parse with an Exh above sentential negation, this would produce a meaning very clearly not intuited.

(137) $[\![\text{Exh}_{\text{ALT}} \text{ [not [the flag is green]]}]\!] = 1$ iff $\neg\mathsf{green}_{\exists}(f) \wedge \neg\neg\mathsf{white}_{\exists}(f) \wedge \neg\neg\mathsf{red}_{\exists}(f)$.
$\equiv \neg\mathsf{green}_{\exists}(f) \wedge \mathsf{white}_{\exists}(f) \wedge \mathsf{red}_{\exists}(f)$.
$\approx$ the flag is all colours except for green

This meaning ('green$_{\exists}$ is the only colour that the flag is not') is straightforwardly not intuited.

It is not immediately clear if the same goes for other cohyponyms. I do not have a strong intuition about whether (138) would be true of a tragicomedy.

(138) This play is not a comedy.
(? if about a tragicomedy)

On the other hand, it is clear that (139) would be true of a spork (that is, sporks are not forks).

(139) This utensil is not a fork.

As such, *fork* is strong in (139) despite being below *not*.

On the other hand, in cases where no portmanteau exists for conjoined concepts (140), unlike (138) and (139) (where we have *tragicomedy* and *spork*), the judgment returns to paralleling colour terms. That is, (140) behaves like colour terms, in that *x is not municipal* means *x is not municipal at all* (hence the infelicity of (140b)), just like *x is not green* in (134) means *x is not green at all*. This contrasts with (139) where *x is not a fork* does not mean that *x* does not have any of the qualities associated with forks (beyond being part of the domain of utensils), and as such is not

[43] Alternatively, one could say that there is an Exh but, rather than taking the entire set of colour terms as its domain, it only has *green*—perhaps due to the pragmatic salience of polarity in negative sentences. I do not pursue this idea further.

paraphrasable as *x is not a fork at all*.

(140) a. SCENARIO: a public transit project is co-funded by the municipal, provincial, and federal levels of government.
b. #The funding for this transit project is not municipal.

The same goes for other predicates where portmanteaus (like *spork* or *tragicomedy*) do not exist for areas of conceptual overlap:

(141) a. This is not a car.
⇒ false about a half-car, half-boat vehicle
b. This morpheme is not derivational.
⇒ false about a portmanteau morpheme that is both derivational and inflectional

There is probably something important to learn here about the way that the existence of portmanteaus can interfere with judgments about cohyponymic exclusivity, but I must leave this for future research. For now, the conclusion is that at least *in some cases*, such as colour terms, sentential negation means that no Exh is observed at all; whether *fork*-type cohyponyms constitute a true counterexample to this depends on a better understanding of the role played by portmanteaus.

Sentential negation, then, is (in some cases) an exception to the rule, as a DE context where the underlyingly weak meaning of cohyponyms is observable. To be sure, it is still possible to exhaustify them. With normal (free) exhaustivity effects, focus intonation is a way to mark out embedded exhaustivity (see e.g. Horn 1985 and Bassi et al. 2021):

(142) If you eat SOME of the apples, Aisha will smile. (. . . But she wants you to leave some for her!)

In negative sentences, we have seen that non-focused cohyponyms are not exhaustified (143a). But focus intonation on the cohyponym does create a strengthening effect (143b).

(143) a. #The flag of Ireland is not green, it's green, white, and orange.
b. The flag of Ireland is not GREEN, it's GREEN, WHITE, AND ORANGE.

Green is non-exhaustified and existential in (143a), so the first clause means that the flag is not green at all. This is false and indeed contradicted by the second clause in the discourse. On the other hand, *green* in (143b) must be strengthened, so that the first clause is paraphrasable as 'the flag of Ireland is not entirely green,' which is both true and compatible with the second clause in the discourse. Thus, focus intonation can be used to signal exhaustivity embedded underneath the negation: the first clause of (143b) is as in (144), repeated from (135).

(144) $[\![\text{not } [\text{Exh}_{\text{ALT}} \text{ [the flag is green}_F]]]\!] = 1$ iff $\neg[\mathsf{green}_{\exists}(f) \wedge \neg\mathsf{white}_{\exists}(f) \wedge \neg\mathsf{red}_{\exists}(f)]$.
$\approx$ the flag is not only/entirely green

Sentential negation, so far, is the only DE context observed in which (*i*) cohyponyms' meaning is not exhaustified in the basic case, but (*ii*) focus intonation can signal the presence of an embedded Exh operator below negation, paralleling the behaviour of free exhaustivity.

Cohyponyms in other negative DE environments

Beyond sentential negation, there are two classes of DE environments that are 'negative' in some sense. The first class is 'negative-flavoured' clause-embedding verbs, and the second is the restrictor and scope of the quantifier *no*. In these negative DE environments, I find the judgments rather unclear, although in most cases I lean toward interpreting the sentences as displaying cohyponymic exclusivity. Hopefully, future work will show more clearly what to do with these, and whether to collapse them with sentential negation (forming a natural class of negative environments) or positive DE environments (forming the class of DE environments other than *not*).

Let's first consider the complements of what Chierchia (2004:51) calls "clause-embedding verbs that have a negative coloring of some sort." These include dubitatives, negative factives, negative propositional attitude verbs, and predicates of minimum requirement. Dubitatives are verbs like *doubt*:

(145) John doubts that Paul or Bill are in that room. (Chierchia 2004:51)

I haven't been able to get any particular sharp intuitions from myself or others on cohyponyms in such contexts, at least for colour terms:

(146) Aisha doubts that the flag in the box is green.
$\Rightarrow$ green $\approx$ all green? partly green?

The judgment with *comedy* is still an exclusive one, however:

(147) Aisha doubts that you (can) read an entire comedy.
$\not\rightsquigarrow$ Aisha doubts that you (can) read an entire comedy or tragicomedy

Like dubitatives, negative factive predicates like *regret* or *be sorry* see free implicatures disappear:

(148) John regrets that Paul or Bill are in the room. (Chierchia 2004:51)

These contexts seem to pattern like positive DE environments in displaying cohyponymic exclusivity:

(149) a. Aisha regrets that the flag she made is green.
$\Rightarrow$ green $\approx$ all green
b. Aisha regrets that she wrote a comedy.
$\not\rightsquigarrow$ she regrets that she wrote a comedy or tragicomedy

In fact, (149b) is quite strongly infelicitous if Aisha wrote a tragicomedy and not a comedy.

Third, we have negative propositional attitude predicates; Chierchia gives *fear* and *complain* as examples.

(150) John fears that Paul or Bill might not come. (Chierchia 2004:51)

Cohyponyms are still strong in this context.

(151) a. Aisha fears that the flag will be green.
$\Rightarrow$ green $\approx$ all green

b. Aisha fears that the play will be a tragedy.
$\not\rightarrow$ she fears it will be a tragedy or a tragicomedy

Predicates of minimum requirement include *be enough* and *suffice*:

(152) It's enough to know Italian or French (to be admitted to the program).
(Chierchia 2004:51)

These too do not license a weak meaning for cohyponyms.

(153) a. It's enough for the flag to be green.
$\Rightarrow$ green $\approx$ all green
b. It's enough for the play to be a comedy.
$\not\rightarrow$ it's enough for the play to be a comedy or a tragicomedy

Naturally, *be enough* requires some sort of scale pertaining to the satisfaction of some desire or requirement; so both examples in (153) require some pragmatic work to get the cohyponyms to be ordered along a scale. A scenario for (153a) could be that there is a shortage of most inks other than green, so that green flags are considered cheap in comparison to other colours. A scenario for (153b) could simply be that the speaker considers comedies to be of lesser value than some other genres.

In sum, most of the examples of negative-flavoured verbs behave like positive DE environments, in displaying cohyponymic exclusivity. However, at least one of the judgments—with colour terms under *doubt*—is quite unclear.

The second class of negative DE environments—still following the examples given by Chierchia (2004)—consists of the restrictor or the scope of *no*:

(154) a. No student with an incomplete or a failing grade is in good standing.
b. No student who missed class will take the exam or contact the advisor.
(Chierchia 2004:49)

The data with *comedy*-type cohyponyms is about as sharp as in other DE contexts:

(155) a. No play here is a comedy.
$\not\rightarrow$ no play here is a tragicomedy
b. No comedy on this shelf has ever been performed.
$\not\rightarrow$ no tragicomedy on this shelf has ever been performed

Not all speakers I consulted find the data with colour terms to convincingly display cohyponymic exclusivity, however.

(156) a. No flag here is green.
$\Rightarrow$ green $\approx$ all green (?)
b. No flag that is green has Aisha's allegiance.
$\Rightarrow$ green $\approx$ all green? partly green?

My personal hunch is that *green* here does mean 'all green,' and I think some reverse-engineering suggests that this might be right. We just saw that sentential negation requires focus intonation

to allow the strong meaning of cohyponyms in DE contexts. Intuitively, this is because there is a contrast to be signalled: since Exh is not obligatory with cohyponyms under sentential negation (and is in fact strongly dispreferred), its optional presence is signalled through focus intonation.

(157) a. The flag is not green.
$\rightsquigarrow$ the flag is not green at all
b. The flag is not GREEN.
$\rightsquigarrow$ the flag is not entirely green

With *no*, however, this contrast is not found; the sentences in (158) are equally acceptable.

(158) a. No flag here is green, but some are green and white.
b. No flag here is GREEN, but some are green and white.

This suggests that Exh is always obligatory and local to the cohyponym in the scope of *no*; while focus intonation is not disallowed, it does not signal a difference in meaning (there is no contrast), as it does under sentential negation. The same goes for cohyponyms in the restrictor of *no*:

(159) a. No green flag is flying high, but some GREEN AND WHITE flags are.
b. No GREEN flag is flying high, but some GREEN AND WHITE flags are.

As such, I tentatively conclude that both the scope and restrictor of *no* are DE contexts where cohyponymic exclusivity is maintained, much like the other DE contexts I identified and in contrast to sentential negation.

In summary, putting aside the clear exception of sentential negation, and the less clear exceptions of some negative clause-embedding verbs (in particular *doubt*), cohyponyms in DE contexts are intuited as mutually exclusive. This can be captured if Exh is obligatory with and necessarily local to the cohyponyms, as in (160) for the antecedent of a conditional.

(160) $[\![$If the flag is [Exh_{ALT} green], Aisha will sing the national anthem$]\!]$
$=$ 1 iff Aisha will sing the national anthem if the flag is [$\text{green}_{\exists}$ & not $\text{white}_{\exists}$ & not $\text{red}_{\exists}$].
$\approx$ Aisha will sing the national anthem if the flag is entirely green

Interim summary: Exh is ultra-local to cohyponyms

We have gone through four classes of data: sentences with sentence-internal contradictions, sentences with a scope-bearing element like *every*, sentences with an alternative-triggering expression like *some* other than the cohyponym itself, and finally sentences with DE operators. Of course, these four types of data are not entirely logically independent of one another, and to a great extent merely serve to make the same point in empirically different ways. What we found is that in all cases—except for one particular DE environment, viz. sentential negation—cohyponymic exclusivity holds. This shows that the strengthening of cohyponyms does not behave as a run-of-the-mill 'free' exhaustivity effect. A free Exh should never be forced to occur in an embedded context, and indeed should be dispreferred in DE environments.

As we went along, I also showed that it is perfectly possible to generate the meanings observed. The exhaustivity account of cohyponymic exclusivity is not to be rejected as a result of these data. If Exh is claimed to not only be obligatory with cohyponyms, but also to necessarily be very local

to them, cohyponymic exclusivity is expected to remain even in the environments discussed here. The following examples are repeated from the previous subsections; they show how cohyponymic exclusivity can be generated through ultra-local Exh operators.

(161) ⟦The [Exh_{ALT} green] flag is [Exh_{ALT} white]⟧
= 1 iff the [green$_\exists$ & not white$_\exists$ & not red$_\exists$] flag is [white$_\exists$ & not green$_\exists$ & not red$_\exists$].
≈ the entirely green flag is entirely white
⇒ contradiction

(162) ⟦Every flag is [Exh_{ALT} green]⟧
= 1 iff every flag is [green$_\exists$ & not white$_\exists$ & not red$_\exists$].
≈ every flag is entirely green

(163) ⟦Exh_{ALT} [some flags are [Exh_{ALT} green]]⟧

$$= 1 \text{ iff } \begin{cases} \text{some flags are} \begin{pmatrix} \text{green}_\exists \ \& \\ \text{not white}_\exists \ \& \\ \text{not red}_\exists] \end{pmatrix} \ \wedge \\ \neg[\text{all flags are} \begin{pmatrix} \text{green}_\exists \ \& \\ \text{not white}_\exists \ \& \\ \text{not red}_\exists] \end{pmatrix}] \end{cases}$$

≈ some but not all flags are entirely green

(164) ⟦If the flag is [Exh_{ALT} white], the battle ends⟧
= 1 iff the battle ends if the flag is [white$_\exists$ & not green$_\exists$ & not red$_\exists$].
≈ if the flag is entirely white, the battle ends

We have already seen that the exhaustivity account of cohyponymic exclusivity involves Exh being *obligatory*; as I mentioned, this is not unheard of in the literature on exhaustivity. But now we have something quite new: the obligatory Exh that is present with cohyponyms must be local to them. What is more, this locality requirement is even more stringent than merely requiring that Exh should be in the same clause as its associated cohyponym: examples (161)–(163) are monoclausal, but Exh must still take scope closer to the cohyponym than the top of the clause. This is a real puzzle for the syntactic distribution of exhaustivity. Finally, the fact that Exh is both obligatory with cohyponyms and necessarily local to them gives its effect a quasi-lexical flavour, clearly making it quite easy to miss as an exhaustification effect.

The contribution of this section is limited to making note of the puzzle of Exh's ultra-locality; I will return to proposing ways of explaining this in chapter 6. There, I will discuss ways to define what exactly it means for Exh to be obligatorily 'local' to the cohyponym: what exactly is the nature of the locality requirement, and how strict is it? For now, I only want to mention two more empirical points. In (161)–(164), I show Exh as taking scope immediately above its associated cohyponym, and refer to this as an 'ultra-local' Exh, because the locality requirement is more strict than simply needing Exh to be in the same clause as the cohyponym.[44] However, we have also seen that cohyponymic exclusivity disappears in two environments (other than sentential negation): when cohyponyms are conjoined, and with additive particles. For conjunction, the simplest

[44]The literature on exhaustivity often uses the word 'local' to simply mean 'not global,' i.e. embedded in any way. I use the term 'ultra-local' to emphasize the positive nature of the closeness requirement between Exh and its associated cohyponym; it is not just that Exh must be embedded, but that it must be close to the cohyponym.

hypothesis is that Exh is not completely restrained in its syntactic distribution with cohyponyms, such that it can take an entire conjunction of cohyponyms as its prejacent (165). If its prejacent entails two cohyponyms, neither will be strengthened irrespective of the other.

(165) ⟦Every flag here is [Exh_{ALT} [green and white]]⟧
= 1 iff every flag here is $\text{green}_{\exists}$ & $\text{white}_{\exists}$ & not $\text{red}_{\exists}$.

This will be discussed in chapter 6. The other empirical point requiring mention at this moment is the disappearance of cohyponymic exclusivity with additives:

(166) The white flag is also green.

If (166) had an ultra-local Exh on each cohyponym, or even on only one, the sentence would be contradictory. The additive must remove, interact with, or modify the Exh on each cohyponym in (166). How exactly this comes about will be the focus of chapter 3.

2.4.3 Interim conclusion: a first sketch of controlled exhaustivity

We have just seen that the hypothesis that cohyponymic exclusivity is the result of exhaustivity teaches us something interesting about Exh, which shows novel behaviour around cohyponyms. Exh is not just obligatory with them, but also obligatorily local to them. Call this *controlled exhaustivity*: cohyponyms *control* Exh in that they require it and dictate its syntactic position.[45]

(167) **Controlled exhaustivity**: Some alternative-triggering expressions ε 'control' Exh, meaning that
a. there must be an Exh operator scoping above ε; and
b. the Exh operator must scope locally above ε.

This formulation raises a few questions, to be discussed at various points in this thesis, including this very section.

First of all, there is the question of which expressions ε refers to (i.e., what is the set of 'controllers' of Exh). Given the discussion so far, it would seem obvious that only and all cohyponyms are controllers. However, this will be challenged in various ways in chapter 4. Likewise, the requirement (167b) is mum on the nature of the locality requirement; this will be discussed in chapter 6. Another question pertains to the requirement in (167a), which only states that there must be an Exh above the controller, but does not state that the alternatives that this Exh takes should be generated specifically by replacing the controller with its alternatives. In the examples we have seen so far, this is certainly the idea: in (161)–(164), the idea is that the Exh immediately above *green* has alternatives created by replacing *green* with other colours, for example. Data from chapter 4 will complicate this picture, however. Thus, for now, I have written (167a) without referring explicitly to controllers' alternatives.

The notion of controlled exhaustivity in (167) is a claim that certain expressions ε have a close relationship with Exh, of some sort. This suggestion should be understood as a net *addition* to the amount of Exh operators one normally posits in a sentence. The Exh operators that are controlled

[45] This should not be confused with the syntactic notion of control. In the syntactic literature, if α controls β, α must scope above β. In my notion of controlled exhaustivity, if α controls β (β an Exh operator), β scopes above α.

by cohyponyms exist alongside free exhaustivity effects (as discussed with *some* in section 2.4.2). I am not claiming that the other Exh operators are constrained; non-controller alternative-triggering expressions like *some* can be exhaustified at a distance (or not at all) even when a cohyponym in the same sentence puts constraints on its Exh operator. To see this, consider again a sentence with two alternative-triggering expressions, one a controller (e.g., *green*) and one not (e.g., *some*) (168a), and in particular consider it embedded under a DE operator (e.g., *regret*) (168b).

(168) a. Some flags are green.
b. Aisha regrets that some flags in Canada are green.

We already discussed (168a) as requiring two Exh operators, one scoping below *some* and associated with a set of alternatives generated by replacing *green* with its cohyponyms, and one scoping above the entire sentence and associated with alternatives generated by replacing *some* with its scalemates. I showed this in (163). Turning to (168b), however, this example has a DE context. On the most salient parse of the sentence, the alternatives lent by *some* are not excluded (because this would lead to global weakening), but the alternatives lent by *green* are in fact excluded. Indeed, the intuited meaning is that Aisha regrets that any flags in Canada are entirely green. For *some*, there is either no Exh or a global one whose effect is vacuous:[46]

(169) ($\text{Exh}_{\text{ALT-2}}$) [Aisha regrets that some flags in Canada are [$\text{Exh}_{\text{ALT-1}}$ green]].
a. $\text{ALT-1} = \{\text{green, red, white, } \dots\}$
b. $\text{ALT-2} = \left\{ \begin{array}{c} \text{Aisha regrets that some flags in Canada are [Exh}_{\text{ALT-1}}\text{ green],} \\ \text{Aisha regrets that all flags in Canada are [Exh}_{\text{ALT-1}}\text{ green]} \end{array} \right\}$

My point is to emphasize that, while *green* in this example puts special constraints on Exh, this should not be taken to mean that there cannot be other, independent Exh operators behaving along the lines of the description by Chierchia et al. (2012) (i.e., as free exhaustivity, neither obligatory nor syntactically constrained). The Exh associated with *green* must occur below the DE operator, but this does not mean that the Exh associated with *some* must also occur below the DE operator.

2.5 Conclusion

This chapter has laid the groundwork for a new view on the relation between the lexical-conceptual meanings of lexical items and the meanings we intuit when they compose in sentences. I have shown that predicates (at least nouns and adjectives) are systematically strengthened in language, and this occurs even in the absence of entailment scales like {*warm*, *hot*, *boiling*}. We also learned that this exhaustification is necessarily very local to the predicates in question. Indeed, this is why sentence-internal contradictions like (170) are intuited as contradictions; if Exh could be situated anywhere, speakers would choose to interpret these sentences with a global Exh, and the offending predicates would not be strengthened so as to exclude one another.

(170) a. #This comedy is a tragedy.
b. #The white flag is green.

[46] (169) should really also have an Exh operator on *flag*, since it is a predicate and it presumably has cohyponyms.

Predicates are systematically strengthened in sentences, but in ways that effectively refuse to take the rest of the sentence into account.

Let's recap what this means in concrete terms. Consider a simple sentence like (171).

(171) The inflectional morpheme in the word disappeared.

Standard discussions of exhaustivity would not clearly posit any Exh operators here, except perhaps at the top of the sentence as part of a general condition that sentences are exhaustified answers to questions under discussions (QUDs) (e.g., Aravind & Hackl 2017). In contrast, in addition to such QUD-related exhaustivity effects (which I put aside), the conclusion of this chapter is that (171) has at least the Exh operators in (172).

(172) [The [[[$\text{Exh}_{\text{ALT-1}}$ inflectional] [$\text{Exh}_{\text{ALT-2}}$ morpheme]]] [in the [$\text{Exh}_{\text{ALT-3}}$ word]]] disappeared.

We have not yet discussed how local exactly the Exh operators have to be, but (172) is at least an approximation. In (172), each taxonomic predicate is ultra-locally exhaustified.

In the rest of this thesis, I continue building up a theory of controlled exhaustivity, fine-tune some components of the analysis, and defend the analysis against theories that offer competing approaches for a small subset of the data (specifically colour terms). Since much of the argumentation comes from the interaction of additives and Exh, we begin in chapter 3 with a theory of how exactly these interact; I simply assumed in the present chapter that additives manage to solve the problems caused by Exh operators, without yet saying how. In chapter 4, I will turn to the nature of the prejacent and alternatives of controlled Exh. I will show that controlled Exh is observable not just with simplex nouns and adjectives, but also some phrases, namely PPs and some clausal constituents (perhaps VP). I will also elaborate on the claim that the relation between two predicates that must hold for them to be alternatives for the purposes of controlled Exh is the cohyponymy relation. I will try to deal with a broader range of data (including the PPs and VPs just mentioned) by claiming that cohyponymy is a special case of a broader relation. I will claim that, in fact, phrases are alternatives for controlled Exh if they contribute the same kind of information to a given sentence (they have the same 'jurisdiction').

In chapter 5, I focus exclusively on the data we have seen with colour terms, e.g.:

(173) a. The white flag is #(also) green.
b. the white #(and) green flag

In the present chapter, I have treated colour terms as no different from other predicates like *truck*, *dog*, *comedy*, or *inflectional*. While colour terms are not lexically mutually exclusive (they have existential quantificational force over the subatomic parts of their argument), they are made mutually exlusive through controlled exhaustivity. However, there exists a rich literature on so-called 'homogeneity effects,' of which colour adjectives' behaviour is often cited as an example, and where the theories developed treat the strong meanings of colour terms as involving universal quantification, rather than existential quantification together with the exclusion of other colour terms (as is the case in my view). I will therefore turn to the homogeneity literature in chapter 5, and defend the view that colour adjectives' behaviour is best captured through controlled exhaustivity. Finally, chapter 6 will both give a characterization of the locality horizon on controlled Exh, and derive it (as well as the obligatoriness of controlled Exh) through an Agree relation between Exh and

derivational morphology: derivational morphemes have a [*u*Exh] feature requiring the presence of Exh within their projection.

Chapter 3

How Exh and *also* interact

3.1 Introduction

Part of the premise of this thesis has been to use additive particles (*also*, *too*) to uncover exhaustivity effects, as made possible by the findings of Bade (2016) and others. In chapter 1, I accepted Bade's claim that additive particles are obligatory when an unwanted exhaustivity effect would arise without them. As a result, additive particles become a tool by which to probe for exhaustification. So far, however, I have put aside an important question in this pursuit: I have not discussed exactly how additives manage to circumvent problems arising from exhaustification. This is the question I turn to in this chapter.

One of the main concerns in the literature on additive particles is why they are sometimes obligatory, as in (1), for example.

(1) Q: Who sang?
A: Aisha sang. Ben #(also) sang.

Bade (2014, 2016) and Aravind & Hackl (2017) take (1) to involve an obligatory exhaustivity effect that the additive manages to circumvent. On this approach, the problematic meaning that would arise without *also* comes from the structure and meaning in (2), where both sentences are strengthened so as to exclude the other.

(2) Exh_{ALT} [Aisha$_F$ sang]. Exh_{ALT} [Ben$_F$ sang].
a. $[\![\text{Exh}_{\text{ALT}}\ [\text{Aisha}_F\ \text{sang}]]\!] = 1$ iff Aisha sang $\wedge$ Ben did not sing $\wedge$ Carrie did not sing.
b. $[\![\text{Exh}_{\text{ALT}}\ [\text{Ben}_F\ \text{sang}]]\!] = 1$ iff Ben sang $\wedge$ Aisha did not sing $\wedge$ Carrie did not sing.
$\Rightarrow$ contradiction

We will walk through this in more detail in due time. The question this chapter asks about this hypothesis is the following: How exactly does *also* interact with each sentence in (2) (call them S1 and S2) to make them mutually compatible? That is, assuming that Bade (2016) is right in taking (1) to involve an additive successfully circumventing an unwanted obligatory exhaustivity effect, how does *also* do this? There are two parts to this question, if (2) is the right pair of LFs for the answer in (1) without the additive. We need to know how S2, with *also*, is compatible with the plain (non-exhaustified) meaning of S1; and we need to know how, by some 'spooky action at a

distance' (Einstein 1971), the addition of an additive in S2 makes S1 compatible with the plain meaning of S2.

I will start my argumentation by following Aravind & Hackl (2017) in taking S2 to be non-exclusive of S1 because *also* scopes below Exh. The idea is that *also* simply adds the meaning of S1 as an entailment to the prejacent of the Exh in S2 (henceforth Exh-S2), so that Exh-S2 does not end up excluding the meaning of S1. Then, I will turn to the question of how S1 does not contradict S2; after all, as drawn in (2), it too is exhaustified. Two theories have been proposed in the literature. For Bade (2016), S1's compatibility with S2 arises from Exh-S1 being entirely removed from S1 following the utterance of S2 (3). This is a *post-hoc de-exhaustification* approach to the compatibility of S1 with S2.

(3) ~~Exh_{ALT}~~ [Aisha sang]$_1$. Exh_{ALT} [also$_1$ [Ben sang]].

For Aravind & Hackl (2017), on the other hand, S1 is compatible with S2 because the domain of Exh-S1 is restricted (4). S2 is simply not an alternative to Exh-S1, and as such cannot be excluded by it.

(4) $\text{Exh}_{\{\mathsf{sang}(a),\, \sout{\mathsf{sang}(b)},\, \mathsf{sang}(c)\}}$ [Aisha sang]$_1$. $\text{Exh}_{\{\mathsf{sang}(a), \mathsf{sang}(b), \mathsf{sang}(c)\}}$ [also$_1$ [Ben sang]].

These approaches are not clearly teased apart empirically. Indeed, on either proposal, the answer in (1) taken as a whole means that only Aisha and Ben sang, due to Exh-S2 excluding alternatives about other people ('Carrie sang'). To tease apart the theories, I therefore turn to *negative* S2s like in (5), where additives are optional.

(5) a. The flag is white. It is not (also) green.
b. Carrie married Aisha. She did not (also) marry Ben.

My argumentation will centre on a subtle difference in meaning in either sentence in (5), corresponding to the presence of the additive. With *also*, S2 is intuited as informative, while without it, it is intuited as redundant. I will show that this can be understood without creating any problems if the presence of *also* correlates with the domain of Exh-S1 being restricted, as suggested by Aravind & Hackl (2017); but if *also* results in Exh-S1 disappearing altogether, the fact that such discourses taken as a whole are still understood exhaustively will no longer be captured.

At this point, it will be established that S1 is compatible with the meaning of S2 because, even though S1 is exhaustified, S2 is not an alternative for Exh-S1. I will therefore turn to the question of *how* the domain of Exh-S1 comes to be restricted. For Aravind & Hackl (2017), this is due to the domain of Exh-S1 being inherited from the QUD that S1 answers; the domain of answers for QUD-S1 is a proper subset of that of QUD-S2. I will show that this is not the case. In particular, 'clause-internal' additivity effects—single clauses where an additive's presupposition refers anaphorically to the subject of that very same clause—involving exhaustivity on two constituents (X1 and X2) cannot be analyzed as involving two different QUDs. These sentences, such as those in (6), have been called 'pseudo-repetitions' in chapter 2, because they involve an expression from the subject being 'repeated differently' (so to speak) in the predicate of the sentence.

(6) a. The white flag is #(also) green.
b. The person who married Aisha #(also) married Ben.

These data lead to the conclusion that the domain of Exh-S1 is not restricted by the QUD or by some freely applied pragmatic principle, but rather is the direct result of *also* itself. The presence of the additive is necessary for the domain of Exh-S1 to be restricted. In sum, in this chapter, I first give evidence in favour of the domain-restriction theory of Aravind & Hackl (2017), then also modify the theory to accommodate new data with clause-internal additivity.

Having established that additive particles restrict the domain of Exh-X1 (that is, the Exh operator on their antecedent), I will conclude the chapter by pointing that the analysis initially proposed for additives' interaction with Exh-X2 (according to which additives scope below Exh-X2) is no longer particularly well motivated. I will then go further by showing empirically that it does not work for the monoclausal additivity effects discussed in chapter 2. In particular, I will take examples like (7) to show that that *also* cannot always scope below Exh-X2.

(7) The partially white flag is #(also) green.

These operators' interaction must therefore involve *also* restricting the alternatives of Exh-X2, rather than scoping underneath it.

This chapter is organized as follows. In section 3.2, I provide background on the interaction between Exh and additives, and suggest that the fact that S2 does not contradict the plain meaning of S1 is the result of the relative scope of Exh and *also* (as just stated, doubt will eventually be cast on this suggestion). I also lay out the two competing theories for how the presence of *also* in S2 correlates with S1 not being so strengthened as to contradict the plain meaning of S2. Then, in section 3.3, I turn to negative S2s, arguing that they favour the theory that *also* in S2 correlates with the domain of Exh-S1 being restricted. Section 3.4 seeks to establish the mechanism through which the domain of Exh-S1 is restricted, arguing that it is the direct result of the presence of the additive in S2. I make a tentative proposal about how exactly additives manage to restrict the domain of Exh-X1, and point out that, if additives can restrict the domain of Exh-X1, we may as well have them interact with Exh-X2 in the same way. I conclude the section by giving empirical evidence that restricting the domain of Exh-X2, rather than scoping below it, is exactly what additives do. Finally, section 3.5 concludes, and an appendix briefly discusses what would change if we adopted the 'specific existential' definition of additive particles rather than the 'anaphoric' definition.

3.2 Using *also* to make S1 and S2 mutually compatible

Let's start with simple discourses of the kind usually discussed in the obligatory-additive literature.

(8) a. Q: Who sang?
A: Aisha sang. Ben #(also) sang.
b. Q: What did Aisha and Ben do?
A: Aisha sang. Ben #(also) sang.

We have two discourses in (8), but the relevant observation about them is the same either way: in the answer, S2 requires an additive. This goes regardless of whether *also*'s associate (Ben) is focused as new material (8a) or not (8b). If we want our analysis to rely in one way or another on the semantics of focus, we can claim that *Aisha* and *Ben* in (8b) are contrastive topics, and therefore focused (Büring 2016); see in particular the discussion by Sæbø (2004) on the status of

these constituents as contrastive topics. Given that *also* is obligatory regardless of the QUD (8), I will not continue putting an overt QUD on my examples in this chapter.

One caveat about the judgment, in particular for (8a), is that so-called 'list intonation' (see e.g. Steindel Burdin & Tyler 2018) removes the need for the additive:

(9) Q: Who sang?
A: Aisha sang . . . Ben sang . . . That's it, I think.

I take this to be due to meaning contributed by the list intonation (descriptively, a sort of 'de-exhaustification'), and put this aside for this chapter. The sentences in this chapter should be read as intonationally unmarked declaratives.

The discourses in (8) imitate the kind of example conventionally used in the literature on obligatory additives. It is worth noting that they differ from the obligatory-additive effects discussed previously in this thesis in two ways. First, they involve focused individuals (contrastive topics) rather than predicates which are not contrastively focused (or focused at all; see chapter 2). Second, they are biclausal (and indeed bisentential) discourses, whereas I have been discussing monosentential and usually monoclausal data.[1] The kind of data discussed can be summed up as in (10), where (10a) looks like the standard discussion of obligatory additives in the literature, while (10b) is the kind of example I focused on in chapter 2.

(10) a. Aisha sang. Ben #(also) sang.
b. The white flag is #(also) green.

In this chapter, I take as a starting point not the previous contents of this thesis (examples like (10b)), but rather the existing literature on additives (examples like (10a)). However, examples like (10b) will eventually come into play as well, and play a central role in some of the argumentation.

As discussed in chapter 1, Bade (2014, 2016) and Aravind & Hackl (2017) argue that the additive in discourses like (8) is obligatory because an unwanted grammatical exhaustivity effect (Chierchia et al. 2012) would otherwise arise. Without *also*, each sentence is exhaustified in a way that leads to a contradiction (11b).

(11) a. Aisha sang. Ben #(also) sang.
b. (i) $[\![\text{Exh}_{\text{ALT}}\ [\text{Aisha}_F\ \text{sang}]]\!] = 1$ iff $\mathsf{sing}(a) \wedge \neg\mathsf{sing}(b) \wedge \neg\mathsf{sing}(c)$.
(ii) $[\![\text{Exh}_{\text{ALT}}\ [\text{Ben}_F\ \text{sang}]]\!] = 1$ iff $\mathsf{sing}(b) \wedge \neg\mathsf{sing}(a) \wedge \neg\mathsf{sing}(c)$.

Again, in (11), Aisha and Ben are focused due to being either new material (if the QUD is *Who sang?*) or contrastive topics (if the QUD is *What did Aisha and Ben do?*), explaining why they both bear alternatives. For (11b), and for the rest of this chapter, the domain of individuals D, and therefore the alternatives for the individuals Aisha and Ben (on the simplest theory, at least if the QUD is *Who sang?* rather than *What did Aisha and Ben do?*), is $\{$Aisha, Ben, Carrie$\}$.

There are two sides to accounting for how *also* avoids the unwanted outcome in (11b). First, we must explain how *also* makes S2 non-contradictory with the plain (non-exhaustified) meaning

[1] Another difference between my discussion and the standard discussion of obligatory additives in the literature is that I have used *also* rather than *too*. This is not an important distinction because, as far as I can tell, all the claims made in the literature on obligatory additives using *too* would hold just as well if the examples had used *also*. See chapter 7 for more discussion of *too* and *also*.

of S1; that is, with *also*, S2 does not mean 'only Ben sang.' Second, we must explain why S1 does not contradict the plain meaning of S2; at least when S2 has an additive, S1 does not mean 'only Aisha sang.' But there is also another component of the meaning of (11a) which must be captured. Indeed, the discourse as a whole is intuited as meaning that *only* Aisha and Ben sang.[2] Thus, while S1 and S2 must be weak enough not to contradict one another, at least one of them must be strong enough to exclude other alternatives, namely 'Carrie sang.'

In the rest of this section, I take each sentence (S1 and S2) and their interaction with Exh in turn, starting with S2.

3.2.1 Why S2 is compatible with the plain meaning of S1: *also* scopes below Exh

What is *also*'s role in making S2 compatible with the plain meaning of S1? In this section, I propose as a starting point that this has to do with the relative scope of *also* and Exh in S2. This closely follows the proposal by Aravind & Hackl (2017).[3] In section 3.4.4, I will point out that some of the subsequent discussion in this chapter will in fact raise doubts about how well motivated the present proposal is. Still, nothing goes wrong if we take it as a starting point.

The proposal in this section, in a nutshell, is that additive particles add an entailment to the sentences they occur in; therefore, if they occur in the prejacent of that sentence's Exh, Exh will not exclude any alternatives that are entailed by the additive. To see this, we must first have a theory of the meaning of additives. As discussed in chapter 1, Kripke (2009[1990]) shows that additives' presupposition is not just existential: (12) requires a salient proposition in the common ground of the form *x is having dinner tonight* (where $x \neq$ Aisha). It is not enough for conversational participants to know that others are dining too, and as such, it cannot be uttered out of the blue. For (12) to be felicitous, the conversational participants must have a particular individual in mind who is having dinner that night in addition to Aisha.

(12) Aisha is also having dinner tonight.

Heim (1992) deals with additives' anaphoricity through indexation. (13) follows the spirit of her proposal in having *also* co-indexed with some proposition $g(i)$ from the prejacent's alternatives.

(13) $[\![\text{also}_i]\!]^g = \lambda\,\text{ALT}_{\langle st,t\rangle}.\lambda p.\lambda w : g(i) \in \text{ALT} \wedge g(i)(w) \wedge g(i) \neq p.\ p(w).$

As I discussed in chapter 1, other authors deal with Kripke's observation not through an 'anaphoric' additive bearing indexation, but a 'specific existential' additive; I return to this possibility in an appendix.

[2] At least, this is the case when the QUD is *Who sang?* If the QUD is *What did Aisha and Ben do?*, there is no longer any inference about other people, like Carrie. But this is just as well, because with this QUD, presumably the only alternatives are Aisha and Ben, so we would not expect an Exh operator on either sentence to add any entailments about Carrie.

[3] My discussion differs slightly from Aravind & Hackl's. We both claim that *also* strengthens Exh's prejacent, but while I will use trivalent semantics to explain this, Aravind & Hackl (2017) rely on a stipulation that *also* both presupposes and asserts its antecedent. I will show that there is in fact no need to break from convention in claiming that additives assert anything.

With (13), and assuming an assignment function g where $[1 \rightarrow \lambda w.\ \mathsf{sing}(a)(w)]$, the truth conditions in (14b) hold for the S2 in (14a). Given that the sentence is only defined if Aisha sang, it is required that Aisha sang in order for the S2 to be true, even if only Ben is overtly mentioned in the sentence and only Ben is 'asserted' to have sung. As such, with this indexation, S2 entails that Aisha sang.

(14) a. [Aisha sang]$_1$. Ben also$_1$ sang.

b. $[\![\text{also}_1\ [\text{Ben}_F\ \text{sang}]]\!]^g = \begin{cases} 1 \text{ if } \mathsf{sing}(a) \wedge \mathsf{sing}(b), \\ 0 \text{ if } \mathsf{sing}(a) \wedge \neg\mathsf{sing}(b), \\ \#\ \text{otherwise} \end{cases}$

My proposal is that, in discourses like (14a), all of linguistic material in (14b)—which entails the plain meaning of S1—is exhaustified.[4] That is, Exh scopes above *also*.

(15) $\text{Exh}_{\text{ALT}}\ [\text{also}_1\ [\text{Ben}_F\ \text{sang}]]$

Since Exh's prejacent in (15) entails that both Ben and Aisha sang, the sentence is not strengthened to negate that Aisha sang. It still means that other people like Carrie did not sing, accounting for the discourse's inference that only Aisha and Ben sang.[5]

Note that there is no need to stipulate that Exh must outscope *also*. If *also* outscoped Exh, the discourse would be contradictory and *also*'s presupposition would not be met (it cannot be that two different people are both the only ones who sang). Speakers simply choose the semantically congruent parse where Exh outscopes *also*.

Let's see a bit more formally how this works. The alternatives for Exh in (15) are obtained by replacing *Ben* with focus alternatives. It is not clear *a priori* whether the alternatives Exh takes in (15) must include *also*$_1$ (16a), or whether *also*$_1$ can be pruned (16b). With the alternatives in (16a), (15) does not exclude that Aisha sang because there is no alternative referring to Aisha. On the other hand, with the alternatives in (16b), while we *do* have the alternative *Aisha sang*, it it is not excluded by virtue of being entailed by Exh's prejacent. Note that in either case, *Aisha also*$_1$ *sang* is not a viable alternative due to being semantically ill-formed (*also* requires its antecedent and focus associate to be different), although if it was an alternative, it would cause no problem beyond its own ill-formedness. On the presence of 'also$_1$' in the alternatives, recall from chapter 1 that I assume with Katzir (2007) that alternatives are syntactic objects.

(16) a. $\text{ALT} = \{\text{Ben also}_1\ \text{sang, Carrie also}_1\ \text{sang}\}$

b. $\text{ALT} = \{\text{Ben sang, Aisha sang, Carrie sang, Ben also}_1\ \text{sang, Carrie also}_1\ \text{sang}\}$

Thus, due to both the entailment of Exh's prejacent and the ill-formedness of *Aisha also*$_1$ *sang*, (15) entails that both Ben and Aisha, but no one else, sang. With *also*, S2 does not contradict the plain meaning of S1.

This analysis is both theoretically and empirically well-motivated. Theoretically, the analysis

[4] See Krifka 1992 for discussion of two focus-sensitive operators co-occurring in a sentence, in a way formally akin to (15).

[5] In fact, if we adopt Innocent Inclusion (see chapter 1), the alternative 'Aisha sang' will be included (since it cannot be excluded). As such, the fact that Aisha sang will be part of the assertion. In the rest of the chapter, I will keep treating Exh as only excluding alternatives. It does not seem to fit with intuitions that 'Aisha sang' should be anything more than presupposed.

introduces nothing new to the semantics of additive particles and exhaustivity beyond what is already widely accepted in the literature: that additive particles are anaphoric rather than existential, and that Exh does not exclude alternatives entailed by its prejacent. Empirically, the analysis captures that discourses like (11a)/(14a) mean that *only* Aisha and Ben sang. This comes out from the meaning of S2 alone (Exh-S2 excludes non-entailed alternatives like *Carry sang*), removing any burden of explaining this inference from S1.

At the same time, once we will arrive at section 3.4.4, both these motivations will be weakened or gone. Empirically, the inference that no one other than Aisha or Ben sang will in fact be covered by S1, making it unnecessary for S2 to also have this entailment. Theoretically, we will see that accounting for S1 must involve claiming that *also* does more than simply adding an entailment: it must also prune alternatives from Exh-S1, raising the question of why it couldn't also be pruning alternatives for Exh-S2. Nonetheless, we begin with the analysis in this section as our starting point.

Before moving on, let me compare this analysis to previous work in the literature on additive particles. My proposed analysis differs from Aravind & Hackl's (2017) minimally; they also use the scope of *also* vis-à-vis Exh to explain why S2 is not strengthened so as to contradict the plain meaning of S1. However, as pointed out in footnote 3, these authors do not rely on the fact that S2 with *also* entails S1 due to its presupposition; rather, they suggest to modify the lexical meaning of *also* so that it *asserts* the proposition it is co-indexed with:

(17) a. $[\![\text{too}_i\ p]\!]^{g,w}$ is defined iff (Aravind & Hackl 2017:185)
(i) g(i) is true at w
(ii) $\exists q \neq p \in [\![p]\!]^f$ such that $q = \text{g(i)}$
b. If defined, then, $[\![\text{too}_i\ p]\!]^{g,w} = [\![p \wedge q]\!]^{g,w}$

This is not necessary; Exh-S2's non-exclusion of S1 can be captured without modifying the lexical meaning of *also*. It is also not aligned with the intuitive meaning of additives, to the extent that we have intuitions about what is presupposed and asserted. The second analysis in the literature is Bade's (2016), who claims that *also* results in (the possibility of) the removal of Exh-S2. Empirically, this does not immediately capture that the entire discourses under consideration have an exhaustivity inference, and theoretically, it requires making a claim about additives and Exh that do not fall out for free from their independently known properties (but see section 3.4.4).

3.2.2 Why S1 is compatible with S2: de-exhaustification or domain restriction?

The scope of *also* vis-à-vis Exh explains why S2 does not contradict the plain meaning of S1 in (11a)/(18). But why is S1's meaning not contradictory of S2 ('only Aisha sang')?

(18) Q: Who sang?
A: Aisha sang. Ben also sang.

Two theories have been proposed. For Bade (2016), Exh is entirely removed from S1, as a 'post-hoc' effect arising upon utterance of S2. Call this *(post-hoc) de-exhaustification*.

(19) Exh_{ALT} [Aisha$_F$ sang]. $\xrightarrow[\textit{post-hoc de-exhaustification}]{}$ Aisha$_F$ sang.

For Aravind & Hackl (2017), on the other hand, S1 is always exhaustified; they assume that all sentences answer some QUD and all answers are exhaustified. Rather, S1 does not contradict S2 because Ben is not included in the alternatives of Exh-S1 due to domain restriction:

(20) a. $\text{Exh}_{\text{ALT-S1}}$ [Aisha$_F$ sang]$_1$. $\text{Exh}_{\text{ALT-S2}}$ [also$_1$ [Ben$_F$ sang]].
b. (i) ALT-S1 = {Aisha sang, ~~Ben sang~~, Carrie sang}
(ii) ALT-S2 = {Aisha sang, Carrie sang, Ben sang}

Aravind & Hackl's (2017) theory raises the question of why *also* is necessary in S2; couldn't the domain of Exh in S2 be pruned to just {~~Aisha sang~~, Ben sang, Carrie sang}, in which case *also* would not be necessary to avoid excluding the alternative *Aisha sang*? Perhaps this is impossible simply because the speaker has already mentioned Aisha at this point. While Ben can be pruned from S1 (the speaker is putting Ben aside at that point in the discourse), by S2, Aisha has just been discussed and cannot immediately be pruned without some overt restriction like in (21)—where, indeed, *also* is not needed (and indeed infelicitous).

(21) Aisha sang. Other than Aisha, Ben (#also) sang.

This suggests that in (20), *also* is necessary in S2 because Aisha cannot be pruned from ALT-S2.

With the data at hand, and on the assumption that the theory proposed in section 3.2.1 for S2 is correct, it is not clear how to differentiate between these proposals. They both capture that the discourse as a whole entails that only Aisha and Ben sang. The fact that Carrie did not sing comes only from S2 for the post-hoc de-exhaustification theory of S1, but from both S1 and S2 for the domain-restriction theory of S1. I now turn to new data that only the domain-restriction theory can capture.

3.3 Negated S2s as evidence for domain restriction in S1

In this section, I show that the meaning of negated S2s constitutes evidence for the domain-restriction account of S1's compatibility with S2. The argumentation has two steps. The first is to find an independent way to remove Exh-S2, in order to study Exh-S1 without the complicating factor of inferences that come from Exh-S2. I will show that sentential negation does exactly this (cf. e.g. Bassi et al. 2021). Then, the second step is to observe that negated S2s can take *also* in addition to their negation. The presence of *also* in the non-exhaustified S2 comes with a change in meaning, concerning whether S2 is redundant or informative in the discourse. On the surface, this can be captured as long as the addition of *also* in a negated S2 ensures that the positive S1 does not entail the negative S2, which both the post-hoc de-exhaustification theory and the domain-restriction theory can accomplish. But on the de-exhaustification account, there would be no Exh at all in the discourse: none in S1 due to post-hoc de-exhaustification by *also*, and none in S2 due to it being negative. This would fail to capture that the discourse as a whole is still intuited exhaustively, as I will show. Thus, the domain-restriction account is empirically preferable.

3.3.1 Our starting point: removing Exh-S2 through negation

As stated in section 3.2.2, both the de-exhaustification theory and the domain-restriction theory of S1's compatibility with S2 make similar predictions for data like (11a)/(18). Both theories predict that (11a)/(18), taken as a whole, mean that only Aisha and Ben sang; but they predict it due to different parts of the discourse. The inference that Carrie did not sing arises only from Exh-S2 for the de-exhaustification account, while it arises from both Exh-S1 and Exh-S2 on the domain-restriction account. In this section, we begin teasing apart these theories by pointing out that there is an independent way to remove the Exh from S2, thus ensuring that the only sentence in the discourse to have an Exh is S1. If the only Exh is in S1, then any exhaustivity inferences we observe must be caused by Exh-S1.

This mechanism is sentential negation. Negation, as in (22), results in there being no Exh at all in the sentence it occurs in.

(22) Aisha sang. Ben didn't sing.

S2 in (22) has neither an Exh above negation (23) nor below (24), both of which would yield unattested meanings. While both (23) and (24) are compatible with S1, (23) is too strong, and (24) is too weak (it does not entail that Ben did not sing).

(23) $[\![\text{Exh}_{\text{ALT}}\ [\text{not}\ [\text{Ben}_F\ \text{sang}]]]\!] = 1$ iff $\neg\mathsf{sing}(b) \wedge \neg\neg\mathsf{sing}(a) \wedge \neg\neg\mathsf{sing}(c)$.
$\equiv \neg\mathsf{sing}(b) \wedge \mathsf{sing}(a) \wedge \mathsf{sing}(c)$.
$\approx$ everyone sang except Ben.

(24) $[\![\text{not}\ [\text{Exh}_{\text{ALT}}\ [\text{Ben}_F\ \text{sang}]]]\!] = 1$ iff $\neg[\mathsf{sing}(b) \wedge \neg\mathsf{sing}(a) \wedge \neg\mathsf{sing}(c)]$.
$\approx$ it is not the case that only Ben sang.

Thus, negated S2s are not exhaustified.[6]

Going forward, we will use negated S2s to ensure that there is no Exh on S2, thereby studying Exh-S1 in a controlled environment. Specifically, the lack of Exh under negation makes possible the following experiment. If we find a discourse with a negative S2 that has the inference that Carrie did not sing, this must come from S1, not S2. For instance, in (22), there such an inference, and it must stem from Exh-S1, as in (25). In (25), in order to avoid unnecessarily taking a stance on whether 'Ben sang' is part of the domain of Exh-S1, I put this possible alternative in parentheses.

(25) $[\![\text{Exh}_{\text{ALT}}\ [\text{Aisha}_F\ \text{sang}]]\!] = 1$ iff $\mathsf{sing}(a) \wedge (\neg\mathsf{sing}(b)?) \wedge \neg\mathsf{sing}(c)$.

But (22) and (25) do not tell us anything about the two theories under investigation. First, it is not immediately clear whether the postulated de-exhaustification or domain-restriction of these theories are intended to arise freely in any discourse, or only in the presence of an additive particle in a subsequent sentence (more on this in section 3.4). It is clear that the domain-restriction account can handle the exhaustivity inference of (22): it could claim there is no Exh on S2, and the inference that Carrie did not sing comes from Exh-S1, whose domain may or may not be restricted so as not to include Ben (the discourse entails that Ben did not sing either way, due to the plain meaning of S2). As for the post-hoc de-exhaustification account, it could be claimed that (22) has the

[6] At least, they are not exhaustified with these sorts of alternatives, and/or they are not exhaustified with an Exh operator as defined in chapter 1. See chapter 6 and Bassi et al. 2021 for more discussion of negation.

inference that Carrie did not sing because there is no *also* in S2 that could remove Exh-S1. S1 therefore looks exactly as in (25). In sum, because these theories are made for discourses with an additive in S2, neither predicts anything in particular for (22), and both can be understood in a way that is compatible with the data.

We now turn to what happens when S2 is given an additive in addition to negation. The presence of an additive in S2 modifies this sentence's discourse contribution. This must be captured through some change in the meaning of S1, and I will show that only the domain-restriction theory can do this properly.

3.3.2 Domain restriction in S1

In this section, I show that when S2 is negated, the meaning of discourses like (22) changes according to whether there is an additive in addition to the negation. To account for this, we will need to claim that additives in S2 co-occur with domain restriction of Exh-S1, rather than the removal of Exh-S1. However, the data are not without complications; in particular, the strength of the relevant judgments changes according to the example in ways I do not fully understand. For this reason, I will initially depart from examples like (22) in order to build my case around another type of exhaustivity effect, namely the strong (universal) meaning of colour terms (discussed in chapter 2, and again later on in chapter 5), before going back to examples like (22).

Informativity and redundancy with colour predicates

Let's start with a short recap of some of the argumentation in chapter 2 about colour terms and in what sense they involve strengthening. In basic sentences (26), colour terms are interpreted as modifying all parts of their argument.

(26) The flag is green.
$\approx$ the flag is entirely green

Following Harnish (1976) and Levinson (1983), I argued in chapter 2 that this is the result of strengthening. On this view, colour terms are lexically weak:

(27) $[\![\text{green}]\!] = \lambda x.\ \exists y[y \sqsubseteq x \wedge \text{green}(y)]$.

Motivation for colour terms' lexical weakness comes from sentences like (28) (see the discussion in chapter 2 on additives, and chapter 5 on conjunction).

(28) a. The flag is green and white.
b. A: The flag is white.
B: Yes, but it's #(also) green.

The universal inference observed with colour terms in sentences like (26) can be accounted for if colour predicates are strengthened vis-à-vis other colour predicates—roughly as in (29).[7] Pretend

[7](29) is good enough for our purposes, but the syntactic placement of Exh is not necessarily correct. Indeed, I argued in chapter 2 that the Exh operator with predicates like colour terms is constrained by some locality requirement. Depending on the nature of this requirement, (29) might have the Exh operator syntactically too far from the colour term. We return to this locality requirement in chapter 6.

the only colours are green, white, and red.

(29) $[\![\text{Exh}_{\text{ALT}}\ [\text{the flag is green}]]\!] = 1$ iff $\mathsf{green}(f) \wedge \neg\mathsf{white}(f) \wedge \neg\mathsf{red}(f)$.

Given that all parts of the flag must have a colour, if the flag is at least partially green and not partially of any other colour, it must be entirely green. The exhaustivity effect interacts with *also* in the expected way (28b), and disappears entirely with sentential negation (30), where what is negated is the plain existential meaning from (27). This patterns just like the exhaustivity effect observed with focused individuals, as shown in section 3.3.1.

(30) The flag is not green.
$\approx$ the flag is not green at all

As such, we have in colour terms an exhaustivity effect which disappears under negation. Discourses with colour predicates behave as expected: with a positive S1, if the S2 is positive too, *also* is required; but it is not required if S2 is negative.

(31) a. The flag is white. It is #(also) green.
b. The flag is white. It is not green.

Interestingly, it is possible to add an additive in the negative S2 (cf. Bade 2016 on optional additives with negation). However, this comes with a subtle change in meaning. Indeed, while the S2 is intuited as entirely redundant information without *also* (32a), this is not the case with *also* (32b). In this case, S2 is intuited as contributing something relevant. It is as if, with *also* in S2, S1 is intuited as having left it open whether the flag is partially green; thus, upon the utterance of S2, it is an informative assertion to negate that the flag is partially green.

(32) a. The flag is white. It is not green. (S2 is *redundant*)
b. The flag is white. It is not also green. (S2 is *informative*)

The judgment is even stronger in monoclausal additivity effects (the 'pseudo-repetitions' of chapter 2; more on these in section 3.4). With these, rather than a somewhat informal notion of redundancy and informativity, we can characterize the difference in meaning brought about by *also* as the stark difference between an outright tautology and a contingent statement.

(33) a. The white flag is not green. (*tautologous*)
b. The white flag is not also green. (*contingent*)

I will focus in this section on bisentential discourses like (32).

This effect is curious: what exactly is *also* contributing to create this subtle difference in meaning? Indeed, in section 3.2.1, I analyzed *also*'s contribution as simply being to add an entailment below Exh-S2. But the negated S2s in (32) have no Exh for *also* to scope below. What is more, the presupposition of the additive in S2 is met/redundant in the context of S1, and is therefore not predicted to make S2 any more or less informative. As such, the presence of *also* must correlate with something happening in S1 rather than S2. The difference in whether a sentence is intuited as redundant or informative at a point in discourse must arise from whether that sentence is entailed by the prior discourse. Thus, for (32), it must be that *also* is affecting whether S2's meaning is entailed by S1.

In which way, then, does *also* affect S1: by removing Exh-S1, or restricting the domain of Exh-S1? These are the two theories we are considering for how S1s in discourses with *positive* S2s come not to be so strengthened as to contradict the meaning of S2; on the simplest theory, an additive in S2s would interact with Exh-S1 in the same way whether S2 is positive or negative. As such, let's consider both theories now, on the assumption that whichever is better suited for these data with negative S2s is also better suited for the data with positive S2s.

We start with the de-exhaustification theory (34):

(34) Exh_{ALT} [the flag is white]. $\xrightarrow[\textit{post-hoc de-exhaustification}]{}$ The flag is white.

This account immediately runs into an important empirical problem. On this view, both S1 and S2 in (32b) are non-exhaustified, due to *also* and negation, respectively. But this misses something important: the discourse means that the flag is not of any other colour. Intuitively, the flag under discussion is entirely white, so it must be that alternatives such as *the flag is red* are excluded (recall that colour terms are lexically existential). The inference that the flag is not red cannot come from S2, which lacks an Exh altogether due to the negation. Hence, we must strengthen S1 to mean that the flag does not have other colours like red; we need an Exh on S1.

At the same time, S1 in (32b) cannot entail that the flag is not partly green. If it did, S2 would be entailed and thus redundant; in effect, there would be no intuited difference between (32a) and (32b). Hence, S1 in (32b) must be strengthened enough to mean that the flag is not red (or other colours like blue, if we are considering all basic colour terms as alternatives), but not so much as to mean that the flag is not green. We can do this by pruning the domain of Exh-S1:

(35) a. Exh_{ALT} [The flag is white].
b. $\text{ALT} = \{$The flag is white, The flag is red, ~~The flag is green~~$\}$
c. $[\![(35a)]\!] = 1$ iff $\mathsf{white}_{\exists}(f) \wedge \neg \mathsf{red}_{\exists}(f)$.

The truth conditions in (35c) entail that the flag is partly white, possibly partly green, but certainly no other colour. With S2 entailing that the flag is in fact not partly green, the discourse ends up meaning that the flag is only/entirely white, as desired; and given that S1 makes no entailment about whether the flag is partly green, S2 in (32b) is intuited as informative rather than redundant, as desired.

To reiterate, the theories under discussion deal with how *also* interacts with S1 from a *positive* S2. Strictly speaking, Bade (2016) and Aravind & Hackl (2017) do not make predictions for how *also* should interact with S1 from a *negative* S2. Yet, if the long-distance interaction between *also* and S1 is best modelled through domain restriction when S2 is negative, the simplest hypothesis for positive S2s is that the interaction is of the same nature. Thus, the data with negative S2s favour the domain-restriction account generally.

This concludes the theoretical contribution of this section: we have learned that S1s are indeed exhaustified when S2 has an additive particle; the domain of Exh-S1 is restricted, as suggested by Aravind & Hackl (2017). Having made these claims from data with colour terms, we now spend the rest of this section going back to the data with focused individuals with eventive verbal predicates. The judgments with these are not consistently clear; still, I believe my claim stands due to the various examples that are indeed quite clear.

Informativity and redundancy with eventive verbs

The contrast in redundancy and informativity found in (32) can be replicated in sentences with eventive verbs and contrastive topics, like those usually found in the literature on obligatory additives. Consider the following discourses:

(36) a. Carrie married Aisha. She did not marry Ben. (S2 is *redundant*)
b. Carrie married Aisha. She did not also marry Ben. (S2 is *informative*)

There is a strong contrast between these two discourses: (36b) clearly considers the possibility that Carrie might have married both people, while this inference is not observed in (36a). This finding plays off of cultural world knowledge about marriages; in more than a few parts of the world, if Carrie marries Aisha, marrying Ben is not even an option.

The effect is not limited to discourses where there is an expectation of 'uniqueness' as in (36). Consider (37), for example:[8]

(37) A: What shapes did Aisha draw?
B: a. Aisha drew circles. She did not draw triangles. (S2 is *redundant*)
b. Aisha drew circles. She did not also draw triangles. (S2 is *informative*)

Aisha could very well have decided to draw a variety of shapes, but S2 is rather strongly intuited as redundant in (37a); S1 settles that circles are the only shape that Aisha chose to draw. In contrast, in B's answer with *also* (37b), B seems to be open to considering that Aisha might have drawn both circles *and* triangles.

The contrasts are not always empirically as strong as in (36) and (37). For whatever reason, the contrast between (38a) and (38b) is not quite as sharp to my ear, although I find that it is still there.

(38) SCENARIO: A and B are workers at a restaurant where customers can order salad, soup, or both, but nothing else.
A: What did the woman order?
B: a. She ordered salad. She did not order soup. (S2 is *redundant*)
b. She ordered salad. She did not also order soup. (S2 is *informative*)

It is not clear what leads to the slightly weaker judgment in (38). Either way, in all examples seen so far, the analysis given for the colour term data in (32) can be extended here: in (36)–(38), S1 is exhaustified and S2 is not; Exh-S1 is domain-restricted in the (b) examples with *also* in S2, but not in the (a) examples. This explains why S2 in the (a) examples is intuited as redundant whereas S2 in the (b) examples is intuited as informative.

There are other discourses in which the intuited contrast becomes even weaker than (38), and for some speakers non-existent. In (39), I return to the original kind of example I used for much of this chapter.

(39) a. Aisha sang. Ben didn't sing.
b. Aisha sang. Ben didn't also sing.

Speakers vary about how much of a contrast they find between these sentences. To my ear, adding

[8] I thank Douglas Lebo for helpful discussion leading to this example.

an explicit QUD helps bring out a contrast:

(40) Q: Among Aisha, Ben, and Carrie, who sang at the party yesterday?
A: a. Aisha sang. Ben didn't sing. (S2 is *redundant*)
b. Aisha sang. Ben didn't also sing. (S2 is *informative*)

It seems to be that there is a contrast between (40a) and (40b) of the same sort as in (32) and (36)–(38). If examples like (39) require an explicit QUD to bring this out, in contrast to the previous examples, this may simply be because the QUD is not as clear in (39) and a wider variety of interpretations are available when the QUD is not explicit.

In sum, while the data with verbal predicates are in some cases less sharp, what we have observed in this section is that in discourses with negative S2s, S2 is generally still intuited as either redundant or informative depending on whether an additive is present in S2. At first blush, this can be captured either as domain-restriction of Exh-S1 in the presence of *also* in S2, or as post-hoc de-exhaustification of Exh-S2 in the presence of *also*. But the discourses under consideration still have an global exhaustive inference. Given that there is no Exh at all in the negative S2s, the exhaustive inference must come from S1. Thus, it must be that the presence of *also* in S2 corresponds with the domain-restriction of Exh-S1.

3.3.3 Exhaustivity, or a Gricean quantity implicature?

In this last subsection, I step back a bit to verify that it is indeed the case that the exhaustivity inferences just described for bisentential discourses with negated S2s must come from an Exh operator. As noted in chapter 1, the semantic account of strengthening (through Exh) is not necessarily a replacement to the view that pragmatic strengthening effects exist; there is no reason why Gricean maxims could not result in strengthened meanings, too. I could have been overly hasty in assuming that the exhaustive inference of S1s with negative S2s must result from Exh. Could the fact that *Aisha sang* in (41) is interpreted exhaustively be the result of a Gricean implicature, instead?

(41) Aisha sang. Ben didn't sing.

In fact, examples where Exh has to be embedded make such a claim impossible. Consider again (42), repeated from (33a).

(42) The white flag is not green. (*tautologous*)

There is no exhautification of *green* in (42), so any exhaustive inference must come from *white*. Yet, this cannot be captured through pragmatic strengthening.

First of all, on principle, a true Gricean approach would probably not create the tautology observed in (42). For a tautology to be intuited, *white* must exclude *green*—it must be 'white, and no other colour at all.' Truly pragmatic strengthening would presumably not create this: outside of particular contexts like logic classes, why would the listener reason that the speaker wishes to utter a tautology? Doing so would violate Grice's maxim of Relation. In fact, for Grice (1975), even semantically tautologous sentences are acceptable and pragmatically understood as non-tautologous:

(43) a. War is war. (Grice 1975:52)

b. Men are men.[9]

He writes that these "extreme examples of a flouting of the first maxim of Quantity ... are, of course, informative at the level of what is implicated, and the hearer's identification of the informative content at this level is dependent on his ability to explain the speaker's selection of this PARTICULAR patent tautology" (Grice 1975:52, emphasis in original). Still, let's consider for the sake of argument that the avoidance of tautologies is not a good enough reason not to strengthen *white* to exclude *green*.

On a Gricean framework (roughly equivalent to postulating a strictly global Exh, which I will use for familiarity of presentation), attempting to strengthen *white* in (42) is not doable. Let's consider two approaches varying on the nature of the alternatives. One option is to step away from the standard neo-Gricean practice and claim that all non-weaker alternatives are excluded just as they are with Exh. But on this approach, global strengthening would yield inferences about different flags altogether:

(44) a. Exh_{ALT} [$_{v\text{P}}$ not [$_{v\text{P}}$ [$_{\text{DP}}$ the white flag] is green]].

b. $\text{ALT} = \left\{ \begin{array}{c} [_{v\text{P}} \text{ not } [_{v\text{P}} [_{\text{DP}} \text{ the white flag}] \text{ is green}]], \\ [_{v\text{P}} \text{ not } [_{v\text{P}} [_{\text{DP}} \text{ the red flag}] \text{ is green}]], \\ [_{v\text{P}} \text{ not } [_{v\text{P}} [_{\text{DP}} \text{ the blue flag}] \text{ is green}]] \end{array} \right\}$

c. $[\![(44a)]\!] = 1$ iff $\left\{ \begin{array}{l} \text{the unique white}_{\exists} \text{ flag is not green}_{\exists} \wedge \\ \neg[\text{the unique red}_{\exists} \text{ flag is not green}_{\exists}] \wedge \\ \neg[\text{the unique blue}_{\exists} \text{ flag is not green}_{\exists}] \end{array} \right.$

Assuming that the existence presupposition in negated alternatives of the form '$\neg$[the unique red$_{\exists}$ flag is not green$_{\exists}$]' 'projects' so as to entail that there is a unique red flag, (44) entails that for every colour, there is a unique flag that is partially of that colour. They are all partially green, too.[10] This is about as far away from the right truth conditions as we could arrive.

We could instead follow the standard Gricean practice in only excluding stronger alternatives. For any strengthening to take place on such an approach, we have to allow syntactically more complex alternatives: rather than replacing *white* with different colours (which are logically independent rather than being stronger), we must replace *white* with conjunctions like *white and red* (Harnish 1976; Levinson 1983), creating stronger expressions. This approach gets the right results in simple sentences:

(45) Exh_{ALT} [the flag is green] $= 1$ iff $\left\{ \begin{array}{c} \text{the flag is green}_{\exists} \wedge \\ \neg[\text{the flag is green}_{\exists} \text{ and white}_{\exists}] \wedge \\ \neg[\text{the flag is green}_{\exists} \text{ and red}_{\exists}] \wedge \\ \neg[\text{the flag is green}_{\exists}, \text{ white}_{\exists}, \text{ and red}_{\exists}] \end{array} \right.$

But in data like (42), this approach too fails to obtain the desired meaning—or in fact any strengthening at all. (46) represents the LF and proposed alternatives:

(46) a. Exh_{ALT} [$_{v\text{P}}$ not [$_{v\text{P}}$ [$_{\text{DP}}$ the white flag] is green]].

[9] He gives *Women are women.*

[10] This means that green is the only colour for which there is not a unique flag. This is fine: there is no alternative of the form *The green flag is not green* anyway, on the assumption that semantic ill-formedness rules out a sentence from being a candidate alternative.

b. $\text{ALT} = \left\{ \begin{array}{c} [_{vP} \text{ not } [_{vP} [_{DP} \text{ the white flag}] \text{ is green}]], \\ [_{vP} \text{ not } [_{vP} [_{DP} \text{ the white and red flag}] \text{ is green}]], \\ [_{vP} \text{ not } [_{vP} [_{DP} \text{ the white and blue flag}] \text{ is green}]], \\ [_{vP} \text{ not } [_{vP} [_{DP} \text{ the white, red, and blue flag}] \text{ is green}]] \end{array} \right\}$

But the alternatives in (46) are all logically independent: the existence of a unique white$_\exists$ flag does not entail the existence of a unique white$_\exists$ and red$_\exists$ flag, not does the existence of a unique white$_\exists$ and red$_\exists$ flag entail the existence of a unique white$_\exists$ flag. As such, none of them are stronger than the assertion. None are excludable, yielding no strengthening whatsoever:

(47) $[\![(46a)]\!] = 1$ iff the unique white$_\exists$ flag is not green$_\exists$.

In sum, it could in principle be that *bisentential discourses* like (48), which do not require the postulation of embedded strengthening, do not tell us anything about Exh and its alternatives, because S1 could be strengthened pragmatically rather than through Exh:

(48) Aisha sang. Ben didn't sing.

But there is no way to obtain the right exhaustivity inference pragmatically in monoclausal discourses. The inference must be computed locally on the additive's antecedent through an Exh operator. The simplest approach is to claim that the same holds in (48). (48) has an exhaustive inference which must come from S1; my claim in this section is that with the presence of *also* in S2 as in (49), Exh-S1 is domain-restricted so as not to exclude that Ben sang.

(49) Aisha sang. Ben didn't also sing.

3.4 What restricts Exh's domain?

So far, this chapter has brought in new data to argue in favour of the proposal by Aravind & Hackl (2017), according to which S1 is non-contradictory of S2 not due to the post-hoc de-exhaustification of S1, but because Exh-S1 is domain-restricted. In this section, I start by pointing out one change to be made to Aravind & Hackl's theory of how *also* and Exh-S1 interact. For Aravind & Hackl (2017), the reason that Exh-S1 has a restricted domain is because it serves as an answer to a question under discussion (QUD) with a restricted domain. The idea is that Exh-S1 and Exh-S2 both inherit their domain from the QUD the sentence answers; Exh-S1 has the alternatives of QUD-S1, and Exh-S2 has the alternatives of QUD-S2. Aravind & Hackl (2017) posit that the contrastive topic in S2 is not part of the alternatives for QUD-S1. This can be illustrated as follows.

(50) A: Who sang? (= {Aisha sang, Ben sang, Carrie sang})
B: S1: (i) Restricted QUD: Who sang? (= {Aisha sang, Carrie sang})
(ii) Answer: Exh_{ALT} [Aisha$_F$ sang]$_1$.
(iii) ALT = {Aisha sang, Carrie sang} (due to restricted QUD)
(iv) meaning: $[\![\text{S1}]\!] = 1$ iff Aisha sang $\wedge$ Carrie did not sing.
S2: (i) 'Full' QUD: Who sang? (= {Aisha sang, Ben sang, Carrie sang})
(ii) Answer: Exh_{ALT} [also$_1$ [Ben$_F$ sang]].

$$\text{(iii)}\quad \text{ALT} = \left\{ \begin{array}{c} \text{Aisha sang,} \\ \text{Ben sang,} \\ \text{Carrie sang,} \\ \text{Ben also}_1 \text{ sang,} \\ \text{Carrie also}_1 \text{ sang} \end{array} \right\} \text{(due to full QUD)}$$

$$\text{(iv)}\quad \text{meaning: } [\![\text{S2}]\!] = \left\{ \begin{array}{c} 1 \text{ if } \mathsf{sing}(a) \wedge \mathsf{sing}(b) \wedge \neg\mathsf{sing}(c), \\ 0 \text{ if } \mathsf{sing}(a) \wedge \neg[\mathsf{sing}(b) \wedge \neg\mathsf{sing}(c)], \\ \text{\# otherwise} \end{array} \right.$$

The idea is that in (50), for both S1 and S2, the alternatives in the (iii) examples are simply inherited directly from the alternatives in the (i) examples, i.e. those associated with the QUD. Recall from section 3.2.2 that the alternatives for S2 have to be a superset of those for S1; if it was possible for Aisha not to be in the alternatives for S2, there would be no need for the presence of an additive at all, contrary to fact.

In this section, I argue that it is not the QUD that is at cause in restricting the domain of Exh-S1 (section 3.4.1). Or at least, *there are* cases where something other than the QUD is at cause. In section 3.4.2, I suggest that domain-restirction does not arise freely, but is instead caused by the presence of additives in S2; in section 3.4.3, I ask how additives accomplish this, with a rather tentative proposal. Finally, in section 3.4.4, I argue that the interaction between *also* and Exh-X2 is also better thought of in terms of domain-restriction than in terms of syntactic scope.

3.4.1 It is not (always) the QUD that restricts Exh's domain

QUDs are a way of formalizing how sentences fit in a discourse. As summarized by Benz & Jasinskaja (2017:177), the QUD is

> a way to characterize how a sentence fits in its context. The idea is that each sentence in discourse addresses a (often implicit) QUD either by answering it, or by bringing up another question that can help answering that QUD. The linguistic form and the interpretation of a sentence, in turn, may depend on the QUD it addresses.

See e.g. Roberts 2012[1998] and Beaver et al. 2017. Notably, the question–answer form associated with this notion of discourse context means that the QUD is not something that changes within a sentence; a sentence answers one and only one QUD. If this much does not hold for sentences, it at least holds for clauses (one could make the case that conjoined clauses, for instance, can answer different QUDs). In particular, within a clause, it cannot be the case that one constituent answers one QUD and another constituent answers another QUD. To give a somewhat crude example, it would be impossible to analyze (51) by claiming that *Aisha* answers one QUD and *ate an apple* answers another.

(51) Aisha ate an apple.

This much is obvious. But the crucial consequence of this line of thinking is that, if different Exh operators within the same clause have different sets of alternatives, the QUD would not be the right theoretical construct to point to as a cause. There is a single QUD associated with that clause, while the putative Exh operators have different sets of alternatives.

I suspect many readers see where this argument is going. In fact, additive particles can serve to avoid unwanted exhaustivity effects even when their additive presupposition is satisfied sentence-internally, as in the pseudo-repetitions mentioned a few times so far in this chapter:

(52) a. The white flag is #(also) green.
b. The comedy is #(also) a tragedy.

In these examples, we are not dealing with S1 and S2, but A1 and A2 and N1 and N2 (to generalize: X1 and X2). On the view taken so far, the examples in (52) are analyzed as follows. *also* is co-indexed with X1, Exh-X2 scopes above *also* (and therefore does not exclude X1),[11] and Exh-X1 has its domain restricted.

To be sure, *also* is defined in (13) as taking a propositional-type argument. This is not a problem; let's see this for (52a). The propositional-type argument of *also* is $[_{vP}$ the white flag be green$_F]$. This clause has as an alternative the tautologous sentence $[_{vP}$ the white flag be white$]$, which is distinct from *also*'s argument and entailed by the context (the phrase 'the white flag' requires the existence of a flag that is white$_\exists$). Given this, *also* creates the presupposition that the white$_\exists$ flag is white$_\exists$, in addition to the sentence's assertion that it is also green$_\exists$. Either way, it must be that Exh-X1 (the Exh on *white*) does not have *green* in its domain, in which case a contradiction would be created.

Crucially, in (52), the domain restriction on Exh-X1 is in the same clause as Exh-X2. Therefore, if the domain of Exh-X1 is truly a subset of the domain of Exh-X2, it cannot be the QUD that is the cause of this domain-restriction in these examples. If another mechanism has to be found for examples like (52), the simplest hypothesis is that that same mechanism is at play in the bisentential discourses we focused on so far. Thus, it cannot be the QUD that restricts the domain of Exh-X1.

3.4.2 Domain restriction does not arise freely

We have just seen that Aravind & Hackl's (2017) proposal modelled around the QUD does not quite stand up; it can only account for bisentential discourses. We now turn to considering a close equivalent to Aravind & Hackl's proposal. For Aravind & Hackl (2017), the domain of Exh-S1 is restricted through inheritance from the domain of QUD-S1; but the restriction of QUD-S1 comes for free, as opposed to being triggered by the presence of *also* in S2. To remove the notion of the QUD from this theory, then, we could claim that the domain of Exh-X1 is in fact simply restricted for free, rather than inheriting this from the QUD. On this view, there would be a pragmatic principle that allows alternatives to be pruned from X1, as long as they resurface on X2. Call this a pragmatic 'procrastination principle.' Intuitively, alternatives can be put aside without any overt cue at a point in the discourse, as long as they were not mentioned before (and perhaps as long as they resurface afterwards, too). It is necessary to claim that alternatives must resurface later: otherwise, the alternative for X2 corresponding to X1 could be pruned from Exh-X2, and *also* would never be obligatory.

Concretely, for discourses like (53), the intuition is simply that it is possible to put aside X2 (whether it is the individual Ben, or the predicate *green*) at the time of uttering X1.

[11] We will return to the question of whether it is really feasible to claim that in these examples, Exh-X2 is high enough to scope above *also*.

(53) a. Aisha sang. Ben also sang.
b. The white flag is also green.

This allows speakers to procrastinate in discourse: speakers can decide to put aside the alternative expressed by X2 at the time of uttering X1.

This hypothesis is theoretically advantageous because it can explain domain restriction on X1 in clause-internal additivity effects (unlike the QUD account), and it also does not require any new claims about additive particles. On this view, additives do nothing more than add an entailment to X2, as described in section 3.2.1. However, the hypothesis makes two wrong predictions, to which we turn in the next two subsections. First, it cannot capture the intuited meanings of bisentential discourses with a negative S2. In particular, it cannot create different meanings for the discourses that arise according to whether *also* is present in S2. Second, it overgenerates in pseudo-repetitions, in that it does not actually make it possible to generate the observed contradictions using exhaustivity (even with the locality constraint on Exh posited in chapter 2). We will discuss these problems in turn.

Problem #1: no way to capture contrasts in meaning with negative S2s

Data with negated S2s suggest that the domain of Exh-X1 cannot in fact be restricted freely; rather, it is only restricted if *also* is present in X2. Recall that discourses like (54) do not have an Exh operator on X2; only X1 is exhaustified.

(54) a. The flag is white. It is not green. (S2 is *redundant*)
b. The flag is white. It is not also green. (S2 is *informative*)

The S2 in (54) is redundant without *also*; with *also*, the discourse still means that the flag is only white, but S2 makes an informative assertion that the flag is not both white and green. The fact that other colours like red are excluded must come about from the strengthening of *white*, because *green* does not undergo any exhaustification. I therefore analyzed (54) in section 3.3 as involving a domain-restricted Exh on *white*, which excludes all colours except green.

(55) a. Exh_{ALT} [The flag is white]$_1$. It is not also$_1$ green.
b. $\text{ALT} = \{$The flag is white, ~~The flag is green~~, The flag is red$\}$

I relied on this domain-restriction to capture that (54b) has an informative S2, in contrast to (54a). But if this domain restriction on Exh-X1 could come about freely, rather than only occurring in the presence of *also*, there would be nothing stopping it from taking place in (54a) as well. Both discourses would be open to a non-redundant reading for S2. As such, *also* itself must be a factor in distinguishing the intuited meanings of (54a) and (54b)—specifically, by causing the domain of Exh-X1 to be restricted.

Problem #2: pseudo-repetitions would not be contradictory (on some assumptions)

It is not only discourses with negative S2s that pose a problem to the free procrastination principle. Pseudo-repetitions do so too. In particular, I now show that the claims that (*i*) additives necessarily scope below Exh-X2 and (*ii*) additives are not directly responsible for restricting the domain of Exh-X1 are not compatible with the empirical data. As such, one of these two claims must be

incorrect. Either *also* interacts with X2 in a way that is more complicated than simply scoping below Exh-X2 (i.e., by restricting the domain of Exh-X2, or by removing Exh-X2 entirely), or *also* is directly responsible for restricting the domain of Exh-X1.

To see this, consider a pseudo-repetition like (56):

(56) The white flag is #(also) green.

Under the hypothesis that the domain of Exh-X1 can be freely pruned of the material in X2, the alternative *green* can be freely removed from Exh-X1. Thus, *white* does not (necessarily) mean 'not green,' and does not cause a contradiction in (56); the contradiction comes from *green* alone, on this hypothesis.

If additives really do not interact directly with Exh operators beyond scoping below them, Exh-X2's prejacent entails *white* because *also* is co-indexed with *white* and part of Exh-X2's prejacent.[12] Because *also* is adjoined to *v*P (as shown by its linear position in non-finite clauses (57)), this means is that Exh-X2 must be able to scope at least as high as *v*P, too.

(57) a. It would be better for the white flag to also be green.
b. ??It would be better for the white flag to be also green.

Yet, if Exh-X2 can be located as high as *v*P, assuming that subjects can reconstruct, there is a parse of (56) where Exh-X2 scopes above both X1 and X2:

(58) $[_{vP}$ Exh-X2 $[_{vP}$ $[_{DP}$ the $[_{AP}$ Exh-X1 white] flag] $[_{vP}$ be $[_{AP}$ green]]]]

From this position, Exh-X2 entails both *white* and *green without requiring an additive*. Hence, Exh-X2 would not strengthen *green* (X2) to mean 'not white.' And as we just saw, Exh-X1 does not strengthen *white* to mean 'not green' because, by the free application of the procrastination principle, 'green' can be freely removed from the alternatives of Exh-X1. Thus, no contradiction is derived.

This does not conclusively prove that the domain of Exh-X1 is not freely restricted by the procrastination principle; it could be, rather, that Exh-X2 necessarily scopes lower than the additive in these pseudo-repetitions, and therefore it really is X2 that is responsible for the creation of a contradiction. But either way, we have no choice but to claim that additive particles are more powerful than they were depicted in section 3.2.1. They do more than simply scope below Exh-X2. They must either prune the alternatives of Exh-X1, or interfere with Exh-X2 directly (by restricting its domain or removing it entirely)—or both.

Interim summary

Having established in section 3.4.1 that QUDs are not what causes domain-restriction of Exh-X1, in this subsection, I asked whether this domain-restriction could be something that occurs freely, rather than needing to be prompted by an additive or as an inheritance from the QUD. I argued based on data with negative S2s that the pruning of Exh-X1 must in fact be the result of *also* itself, rather than a freely applied procrastination principle. Indeed, the observed contrast in discourses

[12] I repeat that we cannot claim that the domain of Exh-X2 can be *freely* restricted like the domain of Exh-X1, because if both Exh operators could have freely restricted domains, *also* would never be necessary to avoid unwanted exhaustivity effects: speakers could simply restrict the domains of all offending Exh operators.

like (59) could not be explained if the domain of Exh-X1 could be restricted freely, because the alternative 'the flag is green' could be removed from the domain of Exh-X1 either way.

(59) a. The flag is white. It is not green. (S2 is *redundant*)
b. The flag is white. It is not also green. (S2 is *informative*)

I then turned to positive sentences with clause-internal additivity effects ('pseudo-repetitions'), as in (60).

(60) The white flag is #(also) green.

While these do not necessarily show that the domain of Exh-X1 is not freely restricted, they *do* show that *also* does something more than just adding an entailment below Exh-X2, as hypothesized in section 3.2.1. To maintain that *also* does nothing more than adding an entailment, we would have to claim that the domain of Exh-X1 is freely pruned due to the procrastination principle, and that Exh-X2 necessarily scopes higher than *also*, i.e. at least at *v*P. Taken together, these claims would fail to generate a contradiction in (60) even in the absence of *also*: *white* would not exclude 'green' due to the procrastination principle, and *green* would not exclude 'white' because the subject could reconstruct to Spec-*v*P, below Exh-X2. Thus, it must be that *also* is powerful enough to interact directly with Exh-X1 and/or Exh-X2, leaving us with no choice but to abandon the theoretical parsimony of section 3.2.1.

3.4.3 How additives restrict the domain of Exh-X1

The domain-restriction of Exh-X1, we have learned, is accomplished neither through the intermediary of the QUD nor through a freely applied pragmatic 'procrastination' principle. Since the domain-restriction of Exh-X1 corresponds to the presence of *also*, it must be *also* itself is involved in the restriction.

How does this come about? The additive's associate is X2, not X1, so it is not obvious that *also* could change the alternatives for X1. One way to think about this is simply as a condition on the proposition the additive is co-indexed with. Recall the definition of *also* given in (61), repeated from (13); in principle, X1 and X2 ($g(i)$ and p) can be any sort of propositions, as long as X1 is an alternative of X2.

(61) $[\![\text{also}_i]\!]^g = \lambda\,\text{ALT}_{\langle st,t\rangle}.\lambda p.\lambda w : g(i) \in \text{ALT} \wedge g(i)(w) \wedge g(i) \neq p.\ p(w).$

We could try to change this so that *also* requires X2 not to be an alternative to X1, in which case Exh-X1 would not take X2 as an alternative. That is, if *also* is defined as requiring X2 not to be an alternative for X1 (while still allowing X1 to be an alternative for X2), *also* could force the set of alternatives associated with X1 not to include X2. In (62), *also* has two indices; i is co-indexed with the additive's antecedent as previously, and j is co-indexed with the antecedent's set of alternatives. That is, $g(j)$ in (62) is a set of alternatives.

(62) $[\![\text{also}_{i,j}]\!]^g = \lambda\,\text{ALT}_{\langle st,t\rangle}.\lambda p.\lambda w : g(i) \in \text{ALT} \wedge g(i)(w) \wedge g(i) \neq p \wedge p \notin g(j).\ p(w).$

If X2 is not an alternative to X1, then it will not be excluded by Exh-X1. In effect, what this means is that language requires overt signalling (through the additive) in order for the alternatives of X1

to be a proper subset of those for X2. Restricting the alternatives of X1 cannot be accomplished freely without an additive, but language is intelligent enough to prune the alternatives of X1 if one of them would be incompatible with the condition imposed by *also*.

Another similar idea would be to add a condition in *also*'s presupposition that its associate and antecedent must be consistent with each other, something which could in principle occur in a number of ways (but apparently occurs through domain-restriction).

3.4.4 Revisiting *also*'s relationship with Exh-X2

Having focused on additives' relationship with X1 for much of this chapter, I end this section by returning to X2. In section 3.2.1, I claimed that *also*'s relationship with Exh-X2 is simple: it merely scopes underneath it. In doing so, it adds an entailment about X1 in Exh-X2's prejacent, ensuring that Exh-X2 does not exclude X1:

(63) a. Aisha sang.$_1$ Ben #(also) sang.

b. $[\![\text{Exh}_{\text{ALT}}\ [\text{also}_1\ [\text{Ben}_F\ \text{sang}]]]\!] = \begin{cases} 1 \text{ if } \mathsf{sing}(a) \wedge \mathsf{sing}(b) \wedge \neg\mathsf{sing}(c), \\ 0 \text{ if } \mathsf{sing}(a) \wedge \neg[\mathsf{sing}(b) \wedge \neg\mathsf{sing}(c)], \\ \# \text{ otherwise} \end{cases}$

I claimed there are two advantages to this. Empirically, this theory can explain why discourses like (63a) are intuited 'exhaustively': as least as an answer to *Who sang?* (see footnote 2), (63a) bears the inference that other people, like Carrie, did not sing. Second, theoretically, this hypothesis does not require any modifications to the traditional understanding of additives; we only need additives to be anaphoric rather than existential (Kripke 2009[1990]).

Neither of these advantages still clearly hold up. Most clearly, we have demonstrated that (at least with negative S2s, but presumably everywhere) additive particles co-occur with exhaustification on X1:

(64) $\text{Exh}_{\{\text{Aisha sang, Ben sang, Carrie sang}\}}$ [Aisha sang]. Ben did not sing.

Thus, the exhaustive inference intuited in discourses like (63a) is already accounted for by Exh-X1. Stipulating an Exh on X2 makes no empirical change. For all we know, X2 could not be exhaustified at all, or it could have an Exh operator that does not have X1 as an alternative.

On the theoretical side of things, the advantage of the theory in section 3.2.1 has not wholly disappeared, but it has been weakened considerably. Indeed, we now know from our study of additives' interaction with Exh-X1 that in fact, the standard picture is not correct; additives do more than just adding an anaphoric additive presupposition. They also lead to the restriction of the domain of Exh-X1. If long-distance (often cross-sentential) interaction between additives and Exh-X1 is possible, claiming some interaction other than syntactic scope between additives and Exh-X2 no longer seems so theoretically burdensome. Nothing goes wrong if we say that additives restrict the domain of Exh-X2, too.

The tentative proposal sketched out in section 3.4.3 does not necessarily predict additives to do this to Exh-X2. It depends on whether we think that additives and Exh should have the same alternatives. If they must have the same alternatives, we hit a problem: my proposal relied on *also* explicitly rejecting that X2 is an alternative for Exh-X1 (62), but it must be that X1 is an alternative for the *also* in X2, and therefore for Exh-X2 as well. Hence, by this proposal, X1 could not be

pruned from the alternatives of Exh-X2. On the other hand, if we claim that *also* and Exh-X2 do not necessarily have the same alternatives, this concern disappears. X1 would be an alternative for *also* in X2, but not for Exh-X2.

There is in fact an empirical argument in favour of viewing *also* as pruning the alternatives of Exh-X2. Consider (65):

(65) The partially white flag is #(also) green.

If Exh-X2 necessarily scoped above *also*, it would be at least as high as *v*P, because *also* does not scope lower than *v*P; (66) is repeated from (57).

(66) a. It would be better for the white flag to also be green.
b. ??It would be better for the white flag to be also green.

Hence, the view that the interaction between *also* and Exh-X2 involves *also* scoping below Exh-X2 commits us to the view that Exh-X2 in (65) can be as high as *v*P. Without *also*, (65) would have the following LF available (assuming that the subject can reconstruct to Spec-*v*P):

(67) $[_{vP}$ Exh-X2 $[_{vP}$ $[_{DP}$ the $[_{AP}$ partially Exh-X1 white] flag] $[_{vP}$ is $[_{AP}$ green]]]]

The issue with this LF is that it does not in fact predict that *also* would be necessary to avoid a contradiction. Exh-X1 does not result in *white* being true of all parts of the flag here because it is outscoped by *partially* ('partially Exh white' means that there is a part which is only white, not that the entire flag is only white), and Exh-X2's prejacent entails both the partial whiteness and the partial greenness of the flag, so that Exh-X2 would not strengthen either colour term to exclude the other. Hence, if (67) was an available LF, *also* would not be obligatory in this sentence; the sentence would be non-contradictory without needing an additive, contrary to fact. This means that (67) must not be an available LF (see chapters 2 and 6 on the locality constraints on Exh with colour terms and other taxonomic predicates). If Exh-X2 must be lower than Spec-*v*P, it must also be lower than *also*, which cannot scope lower than *v*P (66). Given that *also*'s interaction with Exh-X2 in (65) does not consist of it scoping below Exh-X2, presumably what *also* does is to prune X1 (here *white*) from the alternatives of Exh-X2.

3.5 Conclusion

The premise of this thesis is to use obligatory additive effects as a tool by which to uncover exhaustification in natural language. Indeed, Bade (2016) has shown that additives are not obligatory as a result of *Maximize Presupposition*, but rather because they can circumvent unwanted exhaustivity effects. This begs the question of how exactly additive particles interact with exhaustivity—including Exh operators in additives' prejacent and additives' antecedent.

I started this chapter with what seemed like an empirically well-motivated claim about additives' interaction with S2, which, moreover, came at no theoretical cost whatsoever. The claim was that additives interact with Exh-S2 by taking scope below it, thus ensuring that S1 is entailed by the prejacent of Exh-S2 and is therefore not excluded.

(68) Exh_{ALT} $[\text{also}_1$ $[\text{Ben}_F$ sang]] (where g is such that $[1 \to \lambda w.\ \mathsf{sang}(a)(w)]$)

As for the interaction between additives and Exh-S1, I compared two proposals: either S1 undergoes post-hoc de-exhaustification (Bade 2016), or it remains exhaustified but with an Exh operator that does not have S2 as an alternative (Aravind & Hackl 2017). These claims make the same empirical prediction for positive discourses like (69), if (68) is accepted, because the discourse as a whole will entail that *only* Aisha and Ben sang either way.

(69) Aisha sang. Ben #(also) sang.

While work on the interaction between additive particles and implicatures/exhaustivity focuses on obligatory additive effects (Krifka 1998; Sæbø 2004; Bade 2016; Aravind & Hackl 2017), I suggested that a case of *optional* additivity can help tease apart these theories. Indeed, discourses with a negative S2 optionally take an additive, but its presence comes with a difference in whether S2 is redundant, or contributes something new to the discourse. I showed that S2 is not exhaustified in these examples.

(70) a. Aisha sang. Ben did not sing. (S2 is *redundant*)
b. Aisha sang. Ben did not also sing. (S2 is *informative*)

The discourse inference that *only* Aisha sang can only be captured if S1 is exhaustified, with domain-restriction when *also* is present in S2:

(71) $\text{Exh}_{\{\text{Aisha sang, } \sout{\text{Ben sang}}\text{, Carrie sang}\}}$ [Aisha$_F$ sang]$_1$. not [also$_1$ [Ben sang]].

This disfavours the post-hoc de-exhaustification account of the long-distance interaction between *also* and S1, at least if *also* has the same kind of effect on S1 regardless of whether it acts from a positive or negative S2.

I then asked what causes the domain-restriction of Exh-S1. It cannot be the QUD, because this domain-restriction can occur clause-internally:

(72) The white flag is #(also) green.

Rather, I argued that the domain-restriction of Exh-X1 must be accomplished by the additive itself. The difference in meaning of negative S2s corresponding to whether they are endowed with an additive (70) can only be captured if it is the additive itself that restricts the domain of Exh-X1.

Finally, a consequence of the view that additives have the effect of restricting the domain of Exh-X1 is that the claim exemplified by (68) has lost its empirical motivation. Indeed, the exhaustivity inference in this discourse is fully accounted for by S1. What is more, I took data like (73) to show empirically that additives must be able to prune the alternatives of Exh-X2; they cannot always scope below Exh-X2.

(73) The partially white flag is #(also) green.

Exh-X2 in (73) must be low enough that it would create a contradiction without the additive. This means it must be below Spec-*v*P (the subject's reconstructed position)—and therefore below *also*. As such, the interaction between *also* and Exh-X2 cannot be based in syntactic scope alone. Rather, additives prune the alternatives of *both* Exh-X1 and Exh-X2.

3.A Revisiting a 'specific existential' presupposition for *also*

I spent this chapter assuming an anaphoric definition for additive particles (13). However, some authors have weaker 'specific-existential' presuppositions only requiring the existence of a true proposition in the set of alternatives, but without specifying which proposition that is. Call this proposed *weaker* additive $also_w$.

(74) $[\![\text{also}_w]\!] = \lambda\,\text{ALT}_{\langle st,t\rangle}.\lambda\,p.\lambda w : \exists q[q \in \text{ALT} \wedge p \not\Rightarrow q \wedge q(w) = 1].\ p(w)$. (Göbel 2019:289)

On Aravind & Hackl's (2017) *syntactic* approach to the interaction between additives and Exh, when (74) scopes below Exh, it has radically different results from what I discussed in section 3.2.1. The definition for *also* I assumed throughout this chapter involved anaphoricity, so that if *also* is co-indexed with *Aisha sang* (via the index 1, say), (75a) entails that both Ben and Aisha sang. An Exh operator taking scope above $also_1$ will not exclude that Aisha sang, without there needing to be any pruning of the alternative *Aisha sang*.

(75) a. Ben also_1 sang.

b. $[\![\text{Exh}_{\text{ALT}}\ [\text{also}_1\ [\text{Ben}_F\ \text{sang}]]]\!] = \begin{cases} 1 \text{ if } \mathsf{sing}(a) \wedge \mathsf{sing}(b) \wedge \neg\mathsf{sing}(c), \\ 0 \text{ if } \mathsf{sing}(a) \wedge \neg[\mathsf{sing}(b) \wedge \neg\mathsf{sing}(c)], \\ \#\ \text{otherwise} \end{cases}$

Things are entirely different on the definition of Exh in (74). $also_w$ does not add an entailment about any particular alternative. It only entails that one of them is true. Exh does not know which of them this is. Since Exh cannot arbitrarily decide amongst alternatives, none of them are innocently excludable.

If all Exh does is exclude alternatives, we end up with very weak truth conditions, which are virtually equivalent to there being no Exh at all.

(76) $[\![\text{Exh}_{\text{ALT}}\ [\text{also}_w\ [\text{Ben}_F\ \text{sang}]]]\!] = \begin{cases} 1 \text{ if } \exists q[q \in \text{ALT} \wedge \mathsf{sing}(b) \not\Rightarrow q \wedge q(w) = 1] \wedge \mathsf{sing}(b), \\ 0 \text{ if } \exists q[q \in \text{ALT} \wedge \mathsf{sing}(b) \not\Rightarrow q \wedge q(w) = 1] \wedge \neg\mathsf{sing}(b), \\ \#\ \text{otherwise} \end{cases}$

On the other hand, if Exh also innocently includes alternatives, we end up with quite a problematic meaning. Since no alternative is excludable, they are all included:

(77) $[\![\text{Exh}_{\text{ALT}}\ [\text{also}_w\ [\text{Ben}_F\ \text{sang}]]]\!] = \begin{cases} 1, \text{ if } \begin{cases} \exists q[q \in \text{ALT} \wedge \mathsf{sing}(b) \not\Rightarrow q \wedge q(w) = 1] \wedge \\ \mathsf{sang}(b) \wedge \mathsf{sang}(a) \wedge \mathsf{sang}(c), \end{cases} \\ 0, \text{ if } \begin{cases} \exists q[q \in \text{ALT} \wedge \mathsf{sing}(b) \not\Rightarrow q \wedge q(w) = 1] \wedge \\ \neg[\mathsf{sang}(b) \wedge \mathsf{sang}(a) \wedge \mathsf{sang}(c)], \end{cases} \\ \#\ \text{otherwise} \end{cases}$

As such, the proposed $also_w$ (74) is incompatible with the notion of Innocent Inclusion, which would result in all alternatives being assigned true (unless something independent can rule out the LF in (77), which is doubtful).

If we put aside Innocent Inclusion, (76) is weak but not problematic, and adopting it would have sped up the argumentation in this chapter considerably. In this chapter, I started off assuming a meaning like (75b) for S2s; since S2 in (75b) entails that Carrie did not sing, it was not empirically clear whether this should also be entailed by S1 or not, and much of the chapter was devoted to

proving that in fact, S1 *must* entail this as well (due to negative S2s). I therefore eventually posited an Exh in S1 with domain-restriction to avoid contradicting S2. But with the meaning of $also_w$ giving us the very weak truth conditions in (76), we would have immediately had to reach the same conclusion. (76) does not entail anything about any alternative, so S1 must have a domain-restricted Exh for any exhaustivity inferences in the discourse.

Chapter 4

Alternatives in controlled exhaustivity

4.1 Introduction

In chapter 2, I identified cohyponymic exclusivity as a family of exhaustivity effects in language, and argued that the Exh leading to it is both obligatory with and necessarily syntactically local to the alternative-triggering expression—as in (1), for example.

(1) #Some [Exh_{ALT} comedies] are [Exh_{ALT} tragedies].

In all the examples we saw, the expressions 'controlling' Exh in this way were nouns or adjectives, and their alternatives were their cohyponyms (sisters in a taxonomy). In this chapter, I ask two related questions about these ultra-local Exh operators. First, what is the full range of expressions that control Exh; what are the options for the syntactic size and category of the prejacent of controlled Exh operators? Second, what sorts of alternatives do controlled Exh operators take? Are alternatives for controlled Exh always cohyponyms, as described in chapter 2?

Starting with the first question, we will see that there are in fact syntactically complex phrases that control Exh. In particular, at least some PPs and VPs (or perhaps TPs) are ultra-locally exhaustified, as observed in (2).

(2) a. Some textbooks **about math** are #(also) **about physics**.
b. The train crash that **killed Aisha** #(also) **killed Ben**.

From (2a), I will link the exhaustification of PPs with the phenomenon of *thematic uniqueness* (e.g., Fillmore 1968, Perlmutter & Postal 1977, Chomsky 1981, Bresnan 1982, Carlson 1998, Nie 2020), the generalization that there can only be one instantiation of a given thematic role in an event.

Expanding the discussion of controlled exhaustivity effects to phrases like (2) puts into question the claim that controlled exhaustivity is an effect observed only with cohyponyms. PPs and VPs are not cohyponyms, at least in the typical sense. If we tried to go around this by claiming that these phrases are sisters in a taxonomy, this does nothing but raise the question of what the taxonomy is (and what the advantage of framing the discussion in terms of taxonomies is supposed to be). As such, the observation of controlled Exh with non-cohyponymic alternatives motivates finding a new theory of what sort of relation, if not cohyponymy, must exist between expressions for them to behave as alternatives for controlled Exh.

I suggest that we can capture all the data based on the notion of particular expressions contributing different kinds of information in a given sentence. Cohyponyms are usually alternatives for controlled Exh because they provide the same kind of information (e.g., colour for *white* and *green*, the logic of a story for *comedy* and *tragedy*, and so on; I will not formalize this notion in this chapter, hoping it is intuitive enough that the reader will accept it being left for future work); PPs with the same preposition, and VPs with the same verb, also usually contribute the same kind of information. I will call the kind of information provided by a predicate its 'jurisdiction' (cf. the 'qualia' of Pustejovsky 1995). Of course, descriptively, it is plain as day that different phrases contribute different kinds of information to sentences; the significance of my claim is that this notion has theoretical status for controlled exhaustivity.

I will show that the jurisdictional approach to the alternatives in controlled exhaustivity has broad empirical coverage going beyond solving the initial puzzle of controlled Exh with PPs and VPs. Indeed, I will show that the cohyponymy relation does not manage to capture all the data even with simplex predicates. Notably, on the cohyponymy approach, transitivity is expected to be observed among alternatives for controlled Exh. This is not always the case (3), suggesting something more fine-grained is at play.

(3) a. This mathematician is an accountant.
b. This mathematician is #(also) a surgeon.
c. This surgeon is #(also) an accountant.

Moreover, certain data from cohyponyms fall out for free from the semantics on the jurisdictional approach. In particular, predicates relating to artefacts can (and usually do) contribute information about both the form and the function of an object (e.g., *This object is a fork* provides information about both what the object looks like and what the object can be used for). But artefact predicates are only intuited as mutually exclusive when each predicate is intuited as having the same jurisdiction(s). When one predicate has the 'form' jurisdiction and the other the 'function' jurisdiction, cohyponymic exclusivity disappears:

(4) This shirt is my hat.

Beyond empirical advantages, the jurisdictional approach also has the conceptual advantage of being more principled. We do not need to claim that cohyponyms are alternatives for controlled Exh as a fact based in the lexicon (or something like this); rather, cohyponymic exclusivity has origins more deeply in the elementary fact that predicates contribute different kinds of information.

This chapter is organized as follows. In section 4.2, I bring in data with complex predicates that control Exh, with a particular focus on PPs. I will collapse the controlled exhaustivity effect posited for cohyponymic exclusivity with the well-known phenomenon of thematic uniqueness. Then, sections 4.3 and 4.4 turn to the question of what determines the set of alternatives taken by controlled Exh. In section 4.3, I simply defend the idea that some sort of relation must hold between predicates for them to be alternatives for controlled Exh. In particular, I engage with a debate between Wagner (2005, 2006, 2012) and Katzir (2013, 2014) about the nature of alternatives for predicates in *free* exhaustivity effects (i.e., contrastively focused predicates). Wagner claims that a contrastivity relation must exist, while Katzir claims no particular relation must exists, deriving the observed data through Exh's property of Innocent Exclusion instead. Whatever their status for free exhaustivity effects, I show that neither approach works for controlled exhaustivity.

As such, in section 4.4, I turn to my own approach to alternatives. I will posit 'jurisdictions' as a theoretically substantive part of language. Predicates (whether simplex or syntactically complex) are used in a given sentence to contribute a certain kind of information. If two predicates are used to contribute the same kind of information, they effectively compete for the same jurisdiction and act as alternatives; controlled exhaustivity then results in their becoming mutually incompatible. Finally, section 4.5 finishes the discussion by turning away from the alternatives for controlled Exh and toward the alternatives for *additive particles* in clause-internal additivity. Empirically, I discuss cases where additives are *optional* or *blocked* clause-internally. I argue that the set of alternatives that *also* refers to must be the same as those that form the domain of controlled Exh, because *also* is blocked precisely when the predicates in the subject and predicate of the clause do not share a jurisdiction. Section 4.6 concludes.

4.2 Complex phrases that control Exh

So far, my discussion of controlled exhaustivity has focused exclusively on lexicalized, syntactically simplex predicates: single lexical items like *comedy*, *red*, *fork*, or *federal*. In this section, I ask whether controlled exhaustivity is found anywhere else in language. It is *not* found with various expressions that are standardly discussed as involving exhaustification, like *or* or *some*; as discussed in chapters 1 and 2, these are intuited as weak in DE contexts, for example. Finding instances of controlled exhaustivity is not methodologically easy: the meaning contributed by controlled Exh is not clearly teased apart from the lexical meaning of the controller, because Exh is both obligatory with and necessarily local to the controller. But we have found a way to get around this, namely via clause-internal contradictions that disappear with *also*. The fact that *also* can remove these contradictions shows that exhaustivity was involved in creating them, and the fact that this happens clause-internally means that Exh's syntactic distribution is limited to begin with (it cannot scope globally).

I will show that, beyond simplex nouns and adjectives, there are at least two types of phrases with which we can observe a controlled Exh. The first consists of PPs, and the second will be tentatively described as VPs (although it could be a larger clausal constituent). I describe these in turn in sections 4.2.1 and 4.2.2. In my discussion of PPs, I will extrapolate from the data with PPs in copular sentences and suggest more generally that the control of Exh is a fact of thematic heads—not just prepositions but also other thematic assigners like Voice, the head that assigns roles like Agent to external arguments. In particular, I will suggest that a fairly well-established constraint on language called *thematic uniqueness* is the result of controlled exhaustivity. Finally, I will conclude the section with brief comments on the nature of alternatives for controlled Exh taking a syntactically complex prejacent (section 4.2.3).

One of the upshots of this section will be that the claim from chapter 2 that *cohyponyms* control exhaustivity is insufficiently general. Indeed, there are expressions that control Exh without being cohyponymic, descriptively speaking. This will be the focus of the rest of the chapter.

4.2.1 Focus in PPs; thematic uniqueness as an exhaustivity effect

In this section, I first overview how PPs control Exh, then comment on how this relates to the notion of thematic uniqueness.

An overview of controlled exhaustivity with PPs

To find controlled Exh with complex phrases, what we are looking for is a set of sentences that have the form of (5), but where, rather than having two adjectives triggering a contradiction without *also*, we have two complex phrases.

(5) a. Some **federal** responsibilities are (in fact) #(also) **municipal**.
b. The **white** flag is (in fact) #(also) **green**.

Just to be thorough, I will use a slightly fuller version of the paradigm, as in (6) for *federal* and *provincial*, which shows controlled Exh not only being weakened via *also* but also via *and*.

(6) a. These **federal** responsibilities are #(also) **provincial**.
b. these **provincial** #(and) **federal** responsibilities
c. these **federal** responsibilities which are #(also) **provincial**

In fact, such phrases abound. (7)–(9) give some examples from the pseudo-repetition of PPs.

(7) a. These books **about cats** are #(also) **about bicycles**.
b. these books about cats #(and) about bicycles
c. these books about cats that are #(also) about bicycles

(8) a. This medal **for students** is #(also) **for bus drivers**.
b. this medal for students #(and) for bus drivers
c. this medal for students that's #(also) for bus drivers

(9) a. This book **by Noam Chomsky** is #(also) **by Edward Herman**.
b. this book by Chomsky #(and) by Herman
c. this book by Chomsky that's #(also) by Herman

This paradigm is only observed when the preposition in the pseudo-repeated PPs is the same. Contrast the above examples with (10), which has different prepositions and is fully acceptable.

(10) These books **about cats** are **by Ed Herman**.

Of course, the repetition of a preposition with different meanings behaves as if the prepositions were different (perhaps they are).

(11) The book **by the wall** is **by Ed Herman.**

Unlike the strengthening effect with simplex predicates, here there is a head–complement structure. Crucially, the *complements* of the PPs do not need to be cohyponyms for the effect to arise with PPs. Compare *federal* and *provincial* with *cats* and *bicycles*; the former pair are both types of jurisdiction, but the latter pair has little in common, and are certainly not cohyponyms. To drive home the point that the complements of the PPs in examples like (7)–(9) do not arise from any set of alternatives having roots in the taxonomic organization of concepts, i.e. something 'pre-existing,' we can create examples with prepositions taking syntactically complex complements,

exhibiting the creative and open-ended nature of the generative syntactic module:[1]

(12) a. The tax cut for millionaires is #(also) for Pell Grant recipients who started businesses in disadvantaged communities.
b. this tax cut for millionaires #(and) for Pell Grant recipients who started businesses in disadvantaged communities
c. this tax cut for millionaires that's #(also) for Pell Grant recipients who started businesses in disadvantaged communities

The phrases in (12) are entirely novel; the reader presumably shares the judgment without ever having seen the phrases *millionaires* and *Pell Grant recipients who* … in the same context (or at all, for *Pell Grant recipients who* …), never mind as alternatives to one another. Thus, what we want to blame for the status of these PPs as alternatives is the preposition itself, not the DP it takes as a complement.

Given that this mutual-exclusivity effect with PPs disappears with *and* and *also*, it is presumably an exhaustivity effect. Likewise, because the mutual-exclusivity is found even clause-internally, it must be an obligatory and necessarily ultra-local type of exhaustivity—what I have called controlled exhaustivity. Thus, (7a), without *also*, has the structure and meaning in (13).[2]

(13) ⟦These books [Exh_{ALT} about cats] are [Exh_{ALT} about bicycles]⟧
$= 1$ iff these books which are $\begin{pmatrix}\text{about cats \&}\\ \text{not about bicycles \&}\\ \text{not about …}\end{pmatrix}$ are $\begin{pmatrix}\text{about bicycles \&}\\ \text{not about cats \&}\\ \text{not about …}\end{pmatrix}$
$\Rightarrow$ contradiction

The controlled Exh must take both the preposition and its complement DP in its scope. Exh operators on the DPs, without the prepositions, would at most only yield the meaning that the cats are not bicycles and the bicycles are not cats, and more realistically only yield the meaning that the cats are not other animals, and the bicycles are not other vehicles:

(14) ⟦These books about [Exh_{ALT} cats] are about [Exh_{ALT} bicycles]⟧
$= 1$ iff these books about x $\begin{pmatrix}*\mathsf{cat}(x) \wedge\\ \neg *\mathsf{dog}(x) \wedge\\ \neg *\mathsf{bird}(x)\end{pmatrix}$ are about y $\begin{pmatrix}*\mathsf{bicycle}(y) \wedge\\ \neg *\mathsf{truck}(y) \wedge\\ \neg *\mathsf{skateboard}(y)\end{pmatrix}$
$\not\Rightarrow$ contradiction

The Exh operators on (14) do not strengthen the *about*-phrases but the DPs themselves. They do not entail that the books have only one topic.

In fact, in (13), I am simplifying things in not also putting an Exh on the DPs—these DPs are taxonomic predicates, and as such are exhaustified. The real structure involves at least as many Exh operators as those in (15).

[1](12) is inspired by a tweet by Kamala Harris. `https://twitter.com/kamalaharris/status/1155305122911723526`, accessed March 12, 2021.

[2]With the ellipses in (13), I am currently sweeping an important issue under the rug. When we were discussing controlled exhaustivity with cohyponyms, I took for granted that the alternatives for a given predicate were the set of its cohyponyms. These are closed sets provided by world knowledge or the lexicon. On the other hand, with PPs like *about cats*, it is not entirely clear what the set of alternatives is. I return to this in section 4.4.

(15) These [[Exh$_{ALT}$ books] [Exh$_{ALT}$[about [Exh$_{ALT}$ cats]]]] are [Exh$_{ALT}$ [about [Exh$_{ALT}$ bicycles]]].
≈ these objects which are books (and not newspapers/magazines) are about animals that are cats (and not dogs/birds), and they are about nothing else; and they are about vehicles that are bicycles (and not trucks/planes), and about nothing else
⇒ contradiction (due to Exh operators on *about*-PPs, not those on nouns)

I will keep things simple by only marking down the Exh operators that are relevant to the discussion.

Let me make two more empirical observations about the data with PPs controlling Exh. First, the mutual-exclusivity effect of PPs disappears when one PP entails another. Empirically, we already saw this in chapter 2 with predicates from entailment scales (16a) and with hyponyms and their hypernyms (16b).

(16) a. The warm soups are (in fact) hot.
b. The red flags are (in fact) scarlet.

As such, it is not surprising to find something similar with PPs:

(17) a. These books about wildlife are (in fact) about bears.
b. these books about wildlife about bears

Given that (the relevant) Exh operators are located above the entire PPs rather than only the DPs (see above), I assume that what matters is entailment of the whole PP, not entailment of the complement DP. It happens that the entailment 'percolates' upwards (bears are wildlife, therefore a book about bears is more generally about wildlife).

The last empirical point is that there are cases of PPs not requiring *and* or *also*, even though no logical relation exists between them. These are genuine counterexamples. For whatever reason, these counterexamples seem to correlate with whether a given PP *cannot* be used predicatively in English. Consider (18), which has a non-instrumental (but not prototypically comitative) *with*:

(18) the man with the coat with the boots

There is no controlled-Exh effect in (18) despite the PPs having the same preposition: the two PPs can be stacked without requiring conjunction. Unlike the PPs in (7)–(9), this particular comitative-like *with* cannot be used predicatively:

(19) a. #The man is with the coat.
(# on the intended meaning that he is wearing a coat)
b. #The man with the coat is with the boots.
c. #the man with the coat who is with the boots

On the other hand, a truer comitative reading of *with* does allow a predicative use (20), and here the contradiction (controlled-Exh) effect returns (21).

(20) Aisha is with her friends.

(21) a. the woman with her friends #(and) with her family
b. The woman with her friends is #(also) with her family.

It is not clear why this should be. I put the matter aside for this section, leaving for future research a proper understanding of why a particular instance of comitative *with* (and presumably other PPs) behaves differently from the PPs in (7)–(9).[3]

Thematic uniqueness as a controlled exhaustivity effect

To recap, I suggest that the mutual-exclusivity effect observed in the examples in (22) are two instantiation of the same underlying grammatical phenomenon.

(22) a. The **white** flag is #(also) **green**.
b. The letter **for Aisha** is #(also) **for Ben**.

Specifically, the paradigm in both examples results from phrases, whether APs or PPs, controlling exhaustivity.

Beyond syntactic complexity, (22a) and (22b) differ in that the PPs in (22b) involve the assignment of a thematic role, in this case a recipient or beneficiary role. It is already known in the literature on thematic roles that a given role can only be assigned to one entity per event (on some assumptions about 'entities' and 'events,' at least). This constraint is known as *thematic uniqueness*. In this subsection, I therefore take the obvious step of suggesting that thematic uniqueness is a result of controlled exhaustivity.

Since the work of Davidson (1976), it is largely accepted that events have a theoretical (indeed primitive) status in semantics.[4] Thus, a sentence like (23) has existential import: it asserts that there is an event in which Aisha is eating the cake with the fork.

(23) Aisha is eating the cake with the fork.

One hypothesis on the internal structure of events is that their participants are assigned *thematic roles* (see e.g. Blake 1930; Gruber 1965; Fillmore 1968; Jackendoff 1972; Carlson 1984; Dowty 1991; according to Dowty (1991) and Coppock & Champollion (in progress), the idea essentially goes back to the kārakas introduced by Pāṇini in his *Aṣṭādhyāyī*). It will suffice for our purposes to suppose that there are roles like those listed in (24), and that each participant in an event receives at least one such role. This is shown informally in (25a) and more formally in (25b). (25b) constitutes a 'neo-Davidsonian' semantics in combining Davidson's events with thematic roles.

(24) **Thematic roles:**
agent, theme, recipient, beneficiary, instrument, ...

(25) a. Aisha_{Ag} is eating the cake_{Th} with the $\text{fork}_{\text{Instr}}$.
b. $\exists e[\mathsf{eat}(e) \wedge \mathsf{agent}(e,a) \wedge \mathsf{theme}(e,c) \wedge \mathsf{instrument}(e,f)]$

[3] There are also examples whose status is not clear without a more precise generalization of where we expect to find mutual exclusivity with PPs. For instance, the PPs *in the kitchen* and *on my shoe* in (i) are not mutually exclusive. This would be surprising if we focused on the fact that they are both locative in some sense, but it is unsurprising from the perspective that the PPs have different prepositions.

(i) a. The book in the kitchen is on my shoe.
b. The book is in the kitchen, on my shoe.

[4] In what follows, I use the term 'event' to refer to both events and states indiscriminately.

As mentioned above, one generalization about thematic roles in events is called thematic uniqueness (see e.g. Fillmore 1968, Perlmutter & Postal 1977, Chomsky 1981, Bresnan 1982, Carlson 1998, Nie 2020). The concept of thematic uniqueness has roots in the observation that no verb assigns the same thematic role to more than one argument. Carlson (1984) makes this argument by inventing a hypothetical verb *to skick*; this verb would assign the role of agent to the external argument, and would take two internal arguments, one of which would be a theme and the other of which would be another agent. As such, this made-up verb would allow sentences like (26), which would mean that Aisha and Ben both kicked the ball.

(26) #/*Aisha$_{\text{Ag}}$ skicked the ball$_{\text{Th}}$ Ben$_{\text{Ag}}$.

Naturally, Carlson's point is that no such verb exists, in any language. This is taken to motivate thematic uniqueness generally.

In fact, sentences like (26) motivate thematic uniqueness in a rather obscure way. (26) only shows that it is impossible for a verb to have two *DP arguments* with the same thematic role. It gives us nothing to go on regarding whether the impossibility of a verb like *skick* comes from the syntax or semantics. It is not difficult to imagine what a syntactic theory of the illicitness of (26) would look like. Much modern work in syntax suggests that DP arguments are introduced syntactically by heads that assign thematic roles, such as Kratzer's (1996) Voice[5] head that introduces external arguments (often corresponding to the thematic role of agent). Yet, Voice is not a syntactically recursive head; there can only be one Voice head per clause for syntactic reasons. The syntactic non-recursivity of Voice effectively prevents a verb from being associated with more than one agent.

However, it is easy to show that thematic uniqueness cannot only be a syntactic effect. While granting that syntactic factors are probably at play in ruling out sentences like (26), there are instances of thematic uniqueness that cannot be given a syntactic explanation. We can observe this by handling PP adjuncts instead of DP arguments. In principle, syntax should allow PP adjuncts to be stacked recursively without limit, no matter the thematic role assigned by P to its complement DP. Indeed, Ps are like Voice in assigning a thematic role to a DP, but unlike Voice, nothing syntactic prevents the adjunction of multiple PPs involving the same thematic role being assigned by the P to its complement. Thus, if thematic uniqueness and the impossibility of verbs like *skick* were a purely syntactic phenomenon, it should only be observable with DP arguments, not PP adjuncts.

Yet, the same kind of effect as in (26)—the impossibility of assigning a particular thematic role more than once in a given domain, perhaps an event or a clause—is, in fact, also observed with PP adjuncts:[6]

(27) #Aisha is eating with a fork with a spoon.

Intuitively, (27) is illicit because there are two instruments in a single event (or clause), namely the fork and the spoon. In particular, (27) forms a minimal pair with (28), where one of the two instances of *with* is comitative rather than instrumental, and where the illicitness disappears accordingly.

[5]Voice is also known as *v* for languages like English that have a single head both assigning the agentive thematic role and verbalizing roots (Harley 2017). I use 'Voice' and '*v*' interchangeably in this thesis.

[6](27) is similar to the pseudo-repetition examples we have already seen, but does not have one PP in the subject and the other in the predicate of the sentence.

(28) Aisha is eating with a fork with her girlfriend.

(27) does not only show that there is a constraint in language against the presence of multiple bearers of a particular thematic role in a given domain (presumably an event or clause). More specifically, (27) shows that this constraint is the product of something semantic, rather than syntactic. It is almost certainly the case that examples like (26) also involve a syntactic problem, having to do with the non-recursivity of heads like Voice. But if it is accepted that (26) and (27) are two instances of the same 'thematic uniqueness' phenomenon, then it would be insufficient to understand thematic uniqueness in (26) in purely syntactic terms, because the effect persists even when the syntax is properly cared for, as in (27).[7]

Thematic uniqueness is perhaps best known from the second clause (italicised here) of Chomsky's (1981) theta-criterion.

(29) **The theta-criterion** (Chomsky 1981:36): Each argument bears one and only one θ-role, and *each θ-role is assigned to one and only one argument*.

Since, as we have just seen, it is undesirable to limit the observation of this phenomenon to *arguments*, I prefer Carlson's (1998:40) formulation (which, besides, is more clearly semantic in nature).

(30) An event has at most one entity playing a given thematic role.[8]

While the first clause of Chomsky's θ-criterion is controversial (see in particular Jackendoff 1972 and Hornstein 1999), the same kind of critical attention has not been paid to the concept of thematic uniqueness. Critical discussion of thematic uniqueness has focused on apparently 'symmetric' examples like (31). In (31a), Aisha and Ben are both agents; this is fine by (30) if we take conjoined individuals to form single entities. On the other hand, if we assume a particularly close parallel between (31a) and (31b), such that Aisha and Ben have the same thematic role in both examples, then in (31b) we have two *non-conjoined* individuals who are both agents.

(31) a. Aisha and Ben embraced (each other).

[7]Languages vary in the way that we can ensure that particular thematic uniqueness effects are only due to the semantics, and not the syntax. In English, we do this by relying on PP adjuncts instead of DP arguments. Other languages, like Kinyarwanda, have recursive applicative heads which can introduce DP arguments where English would rely on PP adjuncts. Therefore, in Kinyarwanda, we can use certain DP arguments rather than PP adjuncts to observe thematic uniqueness effects that are necessarily semantic in nature. Indeed, DP arguments introduced even by recursive applicative heads must be the unique bearers of a particular thematic role (i)—that is, thematic uniqueness is observable (Nie 2020). The numbers in (i) refer to Bantu noun classes or to persons.

(i) Umugoré a-ra-na-ha-ki-zi-ba-ku-n-som-**eesh-eesh-er-er-ez**-a. (Nie 2020:86)
woman 1-PRES-also-16-7-10-2-2SG-1SG-read-**CAUS-INSTR-APPL-APPL-APPL**-FV
'The woman is also making them read it (cl. 7, book) with them (cl. 10, glasses) to you for me there (cl. 16, in the house).'

The two different heads in bold (*-eesh-* and *-er-*/*-ez-*) can both merge onto verbal expressions recursively, and they allow (require) the merger of a DP argument, to which they assign a thematic role. But crucially, each DP argument must receive a different thematic role—even the different DP arguments of different instantiations of the same recursive applicative head.

[8]I assume that the conjunction of two individuals, e.g. *Aisha and Ben*, counts as a single 'entity' in Carlson's formulation.

b. Aisha embraced Ben.

Dowty (1991) suggests that this symmetry is only superficial. He refers to (32) (which he attributes to Noam Chomsky, via Quang 1971).

(32) a. The drunk embraced the lamppost. (Dowty 1991:583)
b. #The drunk and the lamppost embraced.

Dowty (1991:583) writes, "The oddness in [(32b)] is of course that it implies that the lamppost somehow took part in the act of embracing. Once we see this, it suddenly becomes quite apparent that [(31a)] is not really synonymous with [(31b)] either: the same asymmetry in who is responsible for the action appears there too." In fact, the intuition for (31) is quite straightforwardly that while (31a) entails that Ben embraced Aisha, (31b) does not (Carlson 1998). He could have remained immobile and only passively received the embrace. I conclude that the conjoined DP *Aisha and Ben* is the sole agent in (31a), while in (31b), the subject is the agent and the object is the theme, rather than both being agents.[9]

This leaves as perhaps the trickiest case stative predicates like the symmetric *similar* and *identical* (Carlson 1998) or *to the left/right of* (Dowty 1991). In (33), there is not much identifiable difference between *the rock* and *the tree* in relation to the state described. If the arguments of states are themes, for instance, one would think that the rock and the tree are both themes in (33).

(33) a. The rock is to the left of the tree. (Dowty 1991:563)
b. The tree is to the right of the rock.

To be sure, the sentences contrast in meaning in regards to the 'perspective' taken; indeed, the contrast in (33) can be understood using Talmy's (1978) notion of figure and ground (where the subject in either example in (33) is the figure, the DP in the predicate is the ground). But this does not make (33) less of a counterexample to thematic uniqueness; figure and ground are not thematic roles (Dowty 1991). Thus, it may be that sentences like (33) constitute a genuine exception to thematic uniqueness; perhaps thematic uniqueness is not quite the right generalization. I will return to this in section 4.4, suggesting that the distinction between figure and ground may indeed be what makes examples like (33) acceptable, and this may be understood as part of the broader theory developed in that section. For now, I put aside apparent counterexamples like (33).

Naively, the easiest way to understand thematic uniqueness would be as a hypothesis on the

[9] On the other hand, Carlson (1998) discusses some trickier examples like the following, where neither argument is identifiable as an agent.

(i) a. The car and the bus collided.
b. 1) The car collided with the bus.
2) The bus collided with the car.

Even here, however, the examples are not as symmetrical as they might first appear; (ia) implies movement by both vehicles into the other, while the sentences in (ib) do not. Carlson (1984:38) points out that the apparent symmetry in (i) disappears if one argument is necessarily stationary:

(ii) a. Last night, a bus collided with a bridge abutment, injuring eleven.
b. #Last night, a bridge abutment collided with a bus, injuring eleven.

lexical meaning of the heads that assign thematic roles. Taking thematic uniqueness seriously, one could claim that these heads do not introduce their associated DP as one instance of a bearer of that thematic role in the event, but rather as *the only* bearer of that thematic role in the event. Let me use meta-language constants like Agent as inherently existential; that is, $\mathsf{Agent}(e,x)$ means that x is *a member of* the set of participants in the event e assigned the role of agent. The thematic uniqueness phenomenon could be understood as showing that the meaning of heads that assign thematic roles is not as in (34a), which is only existential, but as in the lexically strong (34b) (shown here for Voice, putting aside the fact that it can introduce external arguments with a broader range of thematic roles than just agents, such as experiencers).[10]

(34) $[\![\text{Voice}]\!] =$

a. $\lambda x.\lambda e.\ \mathsf{Agent}(e,x)$.

b. $\lambda x.\lambda e.\ \mathsf{Agent}(e,x) \wedge \forall y[\mathsf{Agent}(e,y) \rightarrow y \sqsubseteq x]$.

In other words, the observation of thematic uniqueness could be taken to teach us that thematic heads are lexically strong rather than weak. We now turn to seeing that this is not so; (34a) is in fact the right lexical entry, and thematic uniqueness arises from controlled exhaustivity.

The first step in appreciating the weakness of thematic uniqueness is to observe that thematic roles are not only a relation between verbs and their arguments/adjuncts, but also something found internally to DPs (e.g., Abney 1987, Parsons 1995):

(35) Aisha_{Ag}'s destruction of the $\text{spaceship}_{\text{Th}}$ (Abney 1987)

The existence of thematic roles within DPs as in (35) raises the question of which sort of entity they modify. Adjectival modification as in (36) suggests the presence of an event within the DP, so I will continue speaking of 'events.' Nothing hinges on this.

(36) Aisha's rapid destruction of the spaceship

It would also be possible that DP-internal thematic roles involve direct modification of the individuals themselves; this would require rethinking the nature of thematic roles, which would lie beyond the scope of this section.

Given the existence of thematic roles within DPs, a question that arises naturally at this point is whether thematic uniqueness is also found within DPs, or whether the effect is restricted to clausal thematic roles as in (37), repeated from (27).

(37) #Aisha is eating with a fork with a spoon.

In fact, we do observe thematic uniqueness within DPs (38). To observe this, we must adhere to thematic roles assigned by prepositions, in order to focus on the semantics and avoid purely syntactic problems arising.

(38) a. #Aisha's handling of the food with a fork with a spoon

b. #a letter for Aisha for Ben

c. #the book by Chomsky by Herman

[10] See for instance Legate (2014:39ff) for an explicit walk-through of how a head like Voice composes with both the event and the DP to which it assigns a thematic role.

The examples in (38) are clearly deviant on the intended readings. Some of them can be acceptable, so let me comment briefly on this. Most saliently, (38b) is only deviant if *for Aisha* and *for Ben* are understood in the same way (most saliently, as meaning that the letter is directed at Aisha and directed at Ben). However, the *for*-PPs can also be understood as meaning that the letter was written to please someone. The phrase can be interpreted as meaning that the letter is directed at Aisha in order to please Ben, and the phrase's deviancy disappears under this reading. But this is not a counterexample to thematic uniqueness; in fact, it emphasizes the point. If Aisha and Ben are both recipients, the thematic uniqueness effect holds and the phrase is deviant; but if Aisha is a recipient and Ben is a beneficiary, the thematic uniqueness effect is not predicted to be observed, and indeed the phrase becomes acceptable.[11]

We have now observed both clause-level and DP-internal instances of thematic uniqueness. Is it possible to also observe thematic uniqueness across these two domains? To see this, let us check whether thematic uniqueness still holds if one of the modifiers containing a thematic role (i.e., one of the PPs) is within the DP subject proper, while another is predicated of the entire DP through the verb *to be*. I will refer to this configuration as a *distributed* double-predication. And indeed, keeping in mind the caveats just discussed for these examples, the thematic uniqueness effect is still observed:[12]

(39) a. #The letter for Aisha is for Ben.
b. #The book by Chomsky is by Herman.

Thus, we have observed thematic uniqueness with PPs in three different types of environments: within clauses, within DPs, and distributed between the clause and a DP. (40) summarises this.

(40) a. #Carrie wrote the letter for Aisha for Ben.
b. #the letter for Aisha for Ben
c. #The letter for Aisha is for Ben.

We now have the full empirical paradigm that we need to observe that thematic uniqueness behaves like a controlled exhaustivity effect. As we have seen, controlled exhaustivity yields mutual-exclusivity inferences between phrases within clauses and DPs (41), but this mutual exclusivity disappears with conjunction or additive particles (42).

(41) a. #The white flag is green.
b. #the white green flag

(42) a. (i) The flag is white and green.
(ii) the white and green flag
b. The white flag is also green.

[11] A similar point holds in a much less subtle way for (38c), which is acceptable if one of the *by* prepositions is understood as introducing a location—that is, if the DP means 'the book that Chomsky wrote that is sitting next to Herman' (or vice-versa). Obviously, this reading should be put aside.

[12] I am putting aside (38a) because it is independently deviant even with only one of the PPs in the VP position:

(i) *Aisha's handling of the food is with a fork.

Putting a second *with*-PP inside the subject DP is not expected to improve the problem with (i), so the example does not provide good testing grounds for distributed thematic uniqueness.

As we have already seen at the start of this section, PPs—now discussed in terms of their thematic properties—behave in the same way:

(43) a. #The letter for Aisha is for Ben.
b. #the letter for Aisha for Ben

(44) a. (i) The letter is for Aisha and for Ben.
(ii) the letter for Aisha and for Ben
b. The letter for Aisha is also for Ben.

As such, I suggest that thematic uniqueness is simply a case of controlled exhaustivity. This claim involves some extrapolation from the data in (43)–(44), which makes use of PPs distributed between the subject and PP, to the data with argument DPs or with PPs that are stacked in the VP. I have already discussed the argument DPs: since there is something going wrong syntactically, there is no simple way to test whether something is going wrong semantically too. The relevant phrases are simply not generated by the syntax. As for the PPs that are stacked in the VP rather than distributed, nothing is wrong syntactically, so one might think it could be shown that the illicitness derives from controlled exhaustivity. On the one hand, we cannot test this through *also* due to its syntax: it requires a particular syntactic configuration to be able to pick up a constituent as its clause-internal antecedent. Distributed thematic double-predication of *for*-PPs can be fixed by *also* (45a), but this is no longer the case when the *for*-PPs are stacked within the VP (45b).[13]

(45) a. The letter [for Aisha]$_i$ is #(also$_i$) [for Ben]$_F$.
b. (i) #Carrie wrote a letter for Aisha [for Ben]$_F$.
(ii) #Carrie wrote a letter [for Aisha]$_i$ **also**$_i$ [for Ben]$_F$.
(iii) #Carrie wrote a letter [for Aisha]$_i$ [for Ben]$_F$ **too**$_i$.
(iv) #Carrie **also**$_i$ wrote a letter [for Aisha]$_i$ [for Ben]$_F$.
(v) #Carrie wrote a letter **also**$_i$ [for Aisha]$_i$ [for Ben]$_F$.

At the same time, data with *and* do suggest that this effect disappears in the same conditions as other controlled exhaustivity effects:

(46) Carrie wrote a letter for Aisha **and** for Ben.

To state the obvious, this is not a conjunction of Aisha and Ben, but of two *for*-PPs which have Aisha and Ben as their respective complements. Thus, on the assumption that *and* is Boolean here (see chapter 5), (46) must be interpreted as showing that there is nothing lexically inconsistent with the predication of multiple *for*-PPs.

Let us therefore claim that thematic uniqueness is always a result of controlled exhaustivity (as well as the syntax of clausal thematic heads in the case of argument DPs). I already discussed controlled exhaustivity with PPs at the beginning of this section 4.2.1 (see the discussion of examples (13)–(15)); in the remainder of this section, I merely extend the discussion to thematic effects with DPs. I remind the reader that I took the mutual-exclusivity effect on pseudo-repeated PPs to

[13]This paradigm might teach us something about *also*—perhaps that material in the subject can behave as an antecedent for *also*, while material in the VP cannot. But it would be strange to conclude from (45b) that the mutual exclusivity of the *for* PPs cannot be fixed by *also* as such, given that we see in (45a) that, given the right syntactic setup, it can.

involve an Exh operator immediately above P, and taking as its alternatives a set of expressions where the P's DP complement is replaced by other DPs (see section 4.2.3 for commentary). (47) is repeated from (13).

(47) ⟦These books [Exh_{ALT} about cats] are [Exh_{ALT} about bicycles]⟧
= 1 iff these books which are $\begin{pmatrix}\text{about cats \&}\\ \text{not about bicycles \&}\\ \text{not about …}\end{pmatrix}$ are $\begin{pmatrix}\text{about bicycles \&}\\ \text{not about cats \&}\\ \text{not about …}\end{pmatrix}$

To extend this analysis to DP arguments to clausal thematic heads like Voice, we need thematic heads to have a weak lexical meaning (as already discussed), as in (48) for Voice (repeated from (34a)).

(48) $[\![\text{Voice}]\!] = \lambda x.\lambda e.\ \mathsf{Agent}(e,x)$.

Heads like Voice must be in the scope of an Exh operator that takes as its prejacent both Voice itself and its DP argument, and nothing else. This should look something like (49).

(49) [[Exh_{ALT} [DP_{Ag} Voice]] VP]

This is not a standard constituent structure, however. On the standard view (e.g., Harley 2017), Voice takes the VP as its first argument and the DP as its second; the DP is in its specifier:

(50) [$_{\text{VoiceP}}$ [$_{\text{DP}}$ (external argument)] [$_{\text{VoiceP}}$ Voice [$_{\text{VP}}$ …]]]

One possibility would be to claim that thematic roles like Agent are actually introduced internally to the DP, rather than by a clausal head (as done in, e.g., Coppock & Champollion in progress:§11.2). The clausal head Voice might simply require a DP that is already marked as an agent. If this is so, the puzzle about constituency disappears.[14]

Interim conclusion

The logic of the argumentation in this section can be summed up with the following paradigm:

(51) a. #The green flag is white. *distributed AP adjuncts*
b. #The letter for my mother is for my sister. *distributed PP adjuncts*
c. #I'm eating with a fork with a spoon. *clause-level PP adjuncts*
d. #*Aisha_{Ag} skicked the ball Ben_{Ag}. *clause-level DP arguments*

In chapter 2, I discussed sentences like (51a), where simplex predicates (nouns or adjectives) are intuited as mutually incompatible. In this section, we observed that the same incompatibility is observed with PPs (51b). In both cases, the incompatibility disappears with conjunction or additive particles:

[14] A slightly more complex possibility would be to claim that Voice forms a complex head with a restrictor, a variable or set of φ-features which ends up being identified with the DP agent higher up in the syntax. This is what Legate (2014) pursues for at least some languages based on data entirely independent from our concerns around thematic uniqueness (the reader is once again directed to Legate 2014:39ff). In that case, the complex Voice–restrictor head could be the constituent that is exhaustified.

(52) a. The green flag is also white.
b. The letter for my mother is also for my sister.

PPs, as such, control exhaustivity just like cohyponyms.

From here, there is a fairly direct analytic line to claiming that thematic uniqueness is generally the result of controlled exhaustivity. Empirically, the copular sentence in (51b), with PPs distributed between the subject and predicate, is not substantially different from the non-copular (51c), where both PPs modify the event and are syntactically positioned in the VoiceP. In both cases, the problem with the sentence intuitively arises because there are two PPs involving the assignment of the same thematic role to different constituents. While (51c) cannot be tested with *also* for syntactic reasons (*also* cannot take a PP clausally adjoined inside its own VoiceP as its antecedent), we still observe that the infelicity disappears with conjunction, as seen in (53c) ((53a–b) emphasize that the clausal PPs in (53c) behave in the same way as other controlled Exh effects).

(53) a. The flag is green and white.
b. The letter is for my mother and for my sister.
c. I'm eating with a fork and with a spoon.

If we accept (51c) as an effect arising from controlled exhaustivity, at this point, we are already analyzing what is descriptively a violation of thematic uniqueness as the result of controlled exhaustivity. Presumably, then, controlled exhaustivity is also involved in the thematic uniqueness violation effect observed with DP arguments (51d), although the badness of (51d)-type sentences is also syntactic, and therefore harder to discuss with any certainty.[15]

Thus, we have collapsed the semantic illicitness of cohyponymic pseudo-repetitions like (51a) and the semantic badness of thematic uniqueness violations like (51c). This suggests that my discussion of controlled exhaustivity in chapter 2 as a fact of cohyponymy, and the discussion of thematic uniqueness effects in the literature on thematic roles, are both insufficiently general. We will return to this question in later sections of this chapter, effectively asking what cohyponyms and heads that assign thematic roles have in common. Before this, however, let us observe controlled exhaustivity with one more set of complex expressions.

4.2.2 Focus in VPs

We just saw that, with some extrapolation, one might reasonably take all thematic uniqueness violation effects to result at least in part from controlled exhaustivity. This would mean that when Voice introduces an external argument and gives it the thematic role of agent, there is an Exh operator whose effect is to strengthen the nature of the thematic role assigned to the DP. Rather than being *an* agent of the event, it becomes the *only* agent. In this section, we turn to another area where Voice—or perhaps a slightly lower or higher clausal head like V or T; I will write V—does seem to require Exh, and where this is motivated directly by empirical facts rather than extrapolation from the behaviour of PPs. The data are an instance of controlled exhaustivity observable from focus effects within the verbal domain.

[15] Aron Hirsch (p.c.) points out that it may be that the observed syntactic non-recursivity of heads like Voice is in fact due to controlled exhaustivity. On this view, Voice is not in fact syntactically constrained (there could in principle be many Voice heads in the clausal spine), but syntactic derivations with more than one such head end up being semantically inconsistent (the argument of each Voice head is the *only* agent of the event).

As the reader will recall, the literature on obligatory additive particles (in particular Krifka 1998, Sæbø 2004, Bade 2016, and Aravind & Hackl 2017) focuses on data like the following.

(54) a. A: Who went to the party?
B: Aisha went to the party. Ben #(also) went to the party.
b. A: What did Aisha and Ben do last night?
B: Aisha went to the party. Ben #(also) went to the party.

Exhaustivity comes into play because *Aisha* and *Ben* are focused in B's answers above. In B's answer in (54b), while *Aisha* and *Ben* are old material, they are focused due to being contrastive topics. By all appearances, this is a *free* exhaustivity effect, rather than a controlled one. Indeed, Bade (2016) finds experimentally that the additive in (54a) becomes fully optional if the two sentences are conjoined:

(55) A: Who went to the party?
B: Aisha went to the party and Ben (also) went to the party.

This is easily understood if Exh is free, and can therefore scope anywhere within a given sentence. In B's answer in (55), Exh can take global scope (56) (Bade 2016). As such, it does not strengthen either clausal conjunct to be contradictory of the other, although it does still exclude other alternatives.[16]

(56) $[\![\text{Exh } [\text{Aisha}_F \text{ went to the party and Ben}_F \text{ went to the party}]]\!]$
$= 1$ iff Aisha went to the party $\wedge$ Ben went to the party $\wedge$ Carrie did not go to the party.

It is not immediately clear what distinguishes these types of exhaustivity effects from those that involve controlled exhaustivity; one might point to the fact that these involve contrastive topics/focused *e*-type arguments, whereas the controlled exhaustivity effects seen so far all involved various $\langle et \rangle$ predicates (nouns, adjectives, and PPs). However, there are examples that bridge this divide—examples where Exh is controlled, but which are more akin to Bade's in having focused arguments rather than non–contrastively-focused predicates. Again, to find instances of controlled exhaustivity, we are looking for sentences with some material from the subject pseudo-repeated in the predicate, and for the sentence to be illicit without *also*, as shown in (57) for nouns,

[16] These conjunction data may remind the reader that I too discussed conjunctions like (i) in chapter 2, while claiming that these predicates involve a *controlled* rather than *free* exhaustivity effect.

(i) a. The flag is white and green.
b. This is a comedy and a tragedy.

As we will return to in chapter 6, a syntactic constraint on Exh in examples like (i) is well motivated from the lack of biclausal paraphrases:

(ii) #The flag is white and it is green.

I will claim in chapter 6 that Exh in (i) can take scope above the entire conjunction phrase (iii), but no higher.

(iii) The flag is $[_{\text{AP}}\ \text{Exh}_{\text{ALT}}\ [_{\text{AP}}\ \text{white and green}]]$.

That is, Exh here is subject to a locality constraint, but the constraint must be loose enough to allow it to take two conjoined cohyponyms in its scope.

adjectives, and PPs.

(57) a. The **train** is #(also) a **plane**.
b. The **white** flag is #(also) **green**.
c. The letter **for Aisha** is #(also) **for Ben**.

As it turns out, controlled exhaustivity is found with VP or VP-like material (the exact size of the constituent is not immediately clear). Indeed, if a VP that exists as part of a relative clause in the subject is pseudo-repeated in the predicate, *also* is required.

(58) The train crash **that killed Aisha** #(also) **killed Ben**.

In (58), the bolded phrases are mutually exclusive without the additive. Reapplying the analysis proposed elsewhere in chapter 2 and previously in this chapter, then, these phrases control exhaustivity:

(59) $[\![$The [train crash that [Exh_{ALT} [killed Aisha_F]]] [Exh_{ALT} [killed Ben_F]]$]\!]$
$= 1$ iff the train crash that $\begin{pmatrix}\text{killed Aisha \&}\\ \text{didn't kill Ben \&}\\ \text{didn't kill Carrie}\end{pmatrix}$ $\begin{pmatrix}\text{killed Ben \&}\\ \text{didn't kill Aisha \&}\\ \text{didn't kill Carrie}\end{pmatrix}$ $\Rightarrow$ contradiction

It is not immediately clear where the Exh above *killed Aisha* should be: at VP, as written out, or somewhere higher, like above the CP *that killed Aisha*? I will have little to say about this, unfortunately.

The fact that (58) is an example about a train crash is not incidental; the nature of the example makes it pragmatically likely that the two dying events happened at the same time. When events happen one after another, *also* is not obligatory anymore:

(60) The person who hugged Aisha (also) hugged Ben.

On its more natural reading, (60) means that the person first hugged Aisha, then hugged Ben, and *also* is fully optional. On the other hand, if the speaker intends that the person hugged them both at the same time, *also* is obligatory. In fact, going back to (58), Justin Royer (p.c.) points out that there is a way to accept (58) without *also*—specifically, if the dying events happened at a large enough interval, for instance if Aisha died on the spot, whereas Ben died significantly later, and especially if this occurred for reasons only indirectly related to the train crash. This observation aligns (58) with (60), and we can think of it in the following way. The VPs, whether in (58) or (60), are exhaustified so as to exclude the other. But tense kicks in as follows. If the phrases are intuited as part of events that occur at the same time, a true contradiction is intuited (p and $\neg$p cannot be true at the same time). If the events happened at different times, however, no contradiction is intuited (p and $\neg$p can both be true if they are true at different times; for instance, it can be raining at 5pm and not raining at 6pm). Thus, (60) still involves controlled Exh. At some point, the person hugged Aisha and not Ben. Afterwards, the person hugged Ben and not Aisha. No contradiction results.[17]

I leave for future research a theory of how these types of examples, which apparently must

[17]This issue of tense is only showing up now because all the data in chapter 2, and much of the initial discussion of controlled Exh with PPs in section 4.2.1, involved *states* rather than dynamic events.

involve local exhaustification, relate to and differ from the non-controlled exhaustification of examples with contrastive topics discussed elsewhere in the literature (e.g., Bade 2016).

4.2.3 Some comments on complex prejacents and alternatives

Above, I assumed that the alternatives for controlled Exh with PPs and VPs are obtained by replacing the DP complements of P and V with different DPs. (61) is repeated from (13)/(47).

(61) $[\![$These books $[\text{Exh}_{\text{ALT}}$ about cats$]$ are $[\text{Exh}_{\text{ALT}}$ about bicycles$]]\!]$
$= 1$ iff these books which are $\begin{pmatrix}\text{about cats} \wedge \\ \text{not about bicycles} \wedge \\ \text{not about} \ldots\end{pmatrix}$ are $\begin{pmatrix}\text{about bicycles} \wedge \\ \text{not about cats} \wedge \\ \text{not about} \ldots\end{pmatrix}$

Things could have been different. Sticking to the data with PPs, it could have been that the preposition itself was replaced by other prepositions, which, for (61), would have yielded meanings like 'the only relation between the book and cats is aboutness' (the books are about cats but are not by cats, for cats, with cats, near cats, etc.).

In fact, which form the alternatives take may be a matter of choice; (62) is best understood as taking alternatives where the DP stays constant and where the preposition is replaced by alternative prepositions.[18]

(62) Some books by Aisha are ??(also) ABOUT Aisha.

(62) seems to require *about* to be given focus intonation, but this could simply be because its complement *Aisha* is old information. However, even here the effect is apparently sensitive to the DP complement of the preposition. The judgment flips if we replace *Aisha* by *herself*.

(63) Some books by Aisha are (#also) about herself.

Moreover, it's also not clear how consistently different-preposition, same-complement PPs require *also*. For instance, (64) is not contradictory and does not require *also*.

(64) The Liberals' political program is BY the business class, FOR the business class.

However, the obligatory comma-intonation suggests there may be an elided *and* in (64).

For what it's worth, it does not seem possible for the alternatives to involve changes in both the DP complement and the preposition itself. Indeed, the judgment in (62) changes if we not only change *by* with *about*, but also *Aisha* with some other DP.

(65) Some books by Aisha are (#also) about wildlife.

This is not a trivial finding; it is often taken (e.g., Sauerland 2004) that if there are two (or more) alternative-triggering expressions in a given domain, the set of alternatives involves all the phrases

[18] It is not clear how to integrate examples like (62) into the theory I will develop in section 4.4, where I claim that expressions are alternatives for controlled Exh if they contribute the same kind of information in a given sentence. For (62), it is true that both contribute information about Aisha, but they contribute different kinds of information about her. I leave this for future work.

possibly obtained by replacing one or more alternative-triggering expression with one of its alternatives. On the other hand, if this 'controlled exhaustivity' effect with prepositions necessarily involves the assignment of semantic focus in determining which PPs are alternatives, it could be that it is simply difficult to assign focus to both a preposition and its DP complement.

As a bit of housekeeping, note that complex phrases that control exhaustivity need not be of the same size:

(66) The letter [for the student] is #(also) [for the old woman living down the road].

As discussed in chapter 1, I assume with Katzir (2007) that alternatives are syntactic objects and that they generally cannot be more syntactically complex than the prejacent of Exh. But Katzir (2007) makes an exception for this: alternatives can be more complex if they are provided contextually. As such, there is no problem with claiming that in (66), *for the student* and *for the old woman living down the road* are alternative PPs.

4.2.4 Interim summary

In this section, we have seen that some expressions that are more syntactically complex than cohyponyms control exhaustivity. In particular, we find the signature of controlled exhaustivity with both PPs and a clausal constituent that might be VP, CP, or anything in between:

(67) a. The letter **for Aisha** is #(also) **for Ben**.
b. The train crash that **killed Aisha** #(also) **killed Ben**.

In my discussion of prepositions as controllers, I suggested that thematic uniqueness more generally is a case of controlled exhaustivity, using as a springboard the observation that prepositions assign thematic roles to their complements. My discussion of VP data like (67b) was less conclusive, but provides a significant empirical bridge between the sort of obligatory-additive effects discussed in this thesis and the effects discussed elsewhere in the literature, for instance by Bade (2016). While more work is needed to fully appreciate the significance of this corner of the paradigm, it is clearly an important part of the empirical picture.

In chapter 2, I gave an initial definition of controlled exhaustivity as in (68).

(68) **Controlled exhaustivity**: Some alternative-triggering expressions ε 'control' Exh, meaning that
a. there must be an Exh operator scoping above ε; and
b. the Exh operator must scope locally above ε.

At the time, there was no reason to think that ε did not simply refer to all and only cohyponyms. We now have reasons to think otherwise. ε includes PPs and VPs (or something like them).

In some of these cases, like (67b), there is clearly a focused element within the phrase controlling exhaustivity. In others, however, this is not so clear. Consider (69):

(69) A: What did Carrie do today?
B: #She wrote a letter for Aisha for Ben.

Given A's question, B's answer in (69) involves focus on the entire VP:

(70) #She $[_{\text{VP}}$ wrote a letter for Aisha for Ben$]_F$.

The focus is not squarely on Aisha or Ben. A theory that tried to explain the obligatory and necessarily local exhaustivity effects found in (67) as arising exclusively from the focus on *Aisha* and *Ben* would have to account for why the effect persists even when focus is manipulated as in (69). Of course, it may simply be that the focus structure is more complex than showed in (70).

We have just improved our definition of controlled exhaustivity by enlarging the set of expressions that control Exh. Let's continue working on our definition of controlled exhaustivity, in order to better understand alternatives.

4.3 A relation between alternatives

We have just observed controlled exhaustivity with syntactically complex $\langle et \rangle$ predicates, namely PPs and VPs. This raises an important question: in chapter 2, I claimed that the alternatives for controlled exhaustivity were cohyponyms, as in (71).

(71) a. This is a $[\text{Exh}_{\text{ALT}}$ comedy$]$.
b. $\text{ALT} = \{\text{comedy, tragedy, epic, } \dots\}$
c. $[\![(71a)]\!] = 1$ iff this is a (comedy & not tragedy & not epic & $\dots$).

This was quite natural given the goal of strengthening cohyponyms vis-à-vis one another. This is motivated empirically by the fact that non-cohyponymic predicates are, in general, perfectly compatible with one another, e.g. *comedy* and *short*:

(72) This comedy is short.

But PPs and VPs do not participate in the cohyponymy relation, at least in the typical sense. If we tried to go around this by claiming that these phrases are sisters in a taxonomy, this does nothing but raise the question of what the taxonomy is (and what is supposed to be advantageous about framing the discussion in terms of taxonomies). As such, the observation of controlled Exh with non-cohyponymic alternatives motivates a new theory of which relation, if not cohyponymy, must exist between expressions for them to behave as alternatives for controlled Exh.

In this section, I will begin investigating alternatives more seriously first by asking whether a particular relation must hold at all for predicates (whether lexical or syntactically complex) to be alternatives for controlled Exh. I will focus on the alternatives for lexical predicates (descriptively, cohyponyms) rather than phrases, because there is already a debate in the literature touching on predicates undergoing *free* (non-controlled) exhaustification via contrastive focus. Indeed, in some of his work on alternatives with free exhaustivity effects, Katzir (2013, 2014) claims that any predicates can be alternatives to one another. He claims that, to the extent that this seems to counter intuitions, it is because the Innocent Exclusion property of Exh ensures that some predicates are not excluded. I will show that, whatever the status of this hypothesis for free exhaustivity effects, it makes incorrect predictions for *controlled* exhaustivity effects. There must therefore be a kind of relation between predicates leading them to be alternatives for controlled Exh, which I will return to in section 4.4.

4.3.1 When predicates are alternatives in free exhaustivity effects

Before asking what sort of relationship must exist between predicates in order for them to be alternatives for controlled Exh, it is worth asking if such a theoretically important kind of relation exists at all. Might it be that any predicates can behave as alternatives to one another for controlled Exh? This exact question is already discussed in the literature on free exhaustivity effects, specifically through a debate between Wagner (2005, 2006, 2012) and Katzir (2013, 2014). Wagner takes a stance of some limited comparability to mine, in that he claims that a constraint exists on which predicates can serve as alternatives to which. Katzir claims no such constraint exists, and suggests to derive the data discussed by Wagner through the 'Innocent Exclusion' property of exhaustivity.

For our purposes, it will suffice to sum up the data they discuss through the following two examples.

(73) Aisha only owns RED convertibles.

a. $\rightsquigarrow$ Aisha does not own blue convertibles
b. $\not\rightsquigarrow$ Aisha does not own expensive convertibles
c. $\not\rightsquigarrow$ Aisha does not own cheap convertibles

(74) SCENARIO: *Aisha's uncle, who is very rich and makes expensive convertibles, came to Aisha's wedding. I wonder what he brought as a present.*

a. He brought a CHEAP convertible
b. #He brought a RED convertible.

For Wagner, these are taken to show that *red* and *blue* can act as alternatives, but *red* cannot be an alternative to *expensive* or *cheap* (or vice-versa).

The first thing to note about these examples is that what Wagner and Katzir are discussing is a free exhaustivity effect, not a controlled exhaustivity effect. In this sense, they are not in direct alignment with my discussion. I am bringing their debate into this chapter because, as we will see, Katzir's argument could in principle be carried over to the data I am analyzing as controlled exhaustivity. To see that (73)–(74) are a free exhaustivity effect, recall from chapter 2 (section 2.3.2) my discussion of the colour terms in these types of examples. In (73), the meaning of the assertion is not 'Aisha owns convertibles that are only red' (which would be mum on whether she also owns convertibles of other colours), but rather 'Aisha *only* owns convertibles that are only red'—with two *only* particles in the paraphrase. Wagner and Katzir's discussion is about the topmost *only* in my paraphrase. They are not concerned with the strengthening of *red* from $red_{\exists}$ to $red_{\forall}$ (that is, *red* is lexically existential but intuited as universal in (73)–(74)). Thus, the debate between Wagner and Katzir pertains to an empirically different domain from my examples.

In this section, I first show that Wagner's account of a constraint on alternatives straightforwardly does not carry over to the controlled-exhaustivity empirical domain, whatever its status for the domain of free exhaustivity. Then I turn to Katzir's alternative theory, which is premised on the idea that there *is* no constraint on which predicates can be alternatives to which. I show how one could try to carry this idea over to the domain of controlled exhaustivity, and offer some arguments against doing this. The conclusion of the section is that we really do need to assume (as I did in chapter 2) that controlled exhaustivity involves a particular relationship (such as cohyponymy) between sets of predicates.

4.3.2 Wagner's account: alternatives must be contrastive

Wagner takes the data in (73)–(74) as showing that there exists a constraint on the types of expressions that can act as alternatives—a constraint going beyond requiring the expressions to be of the same semantic type. In particular, he proposes (e.g., Wagner 2005:253, Wagner 2012:§6.2.5) that predicates can only be alternatives if they come from a natural partition. While convertibles can be divided into a natural partition according to their colour(s), or according to whether they are cheap or expensive, they cannot be partitioned between 'red' and 'cheap' convertibles; some convertibles are both red and cheap and others are neither red nor cheap. Katzir (2013) sums up Wagner's proposal as in (75), where $\Rightarrow$ denotes cross-categorial entailment.

(75) **Contrasting Alternatives:**
A node α' is a true alternative to a node α in the context of a sister node β only if it contrasts with α in the context of β; that is, only if $[\![[\alpha'\ \beta]]\!] \Rightarrow \neg[\![[\alpha\ \beta]]\!]$.

Regardless of its status for the alternatives in the domain of free exhaustivity, (75) does not work as a hypothesis for controlled exhaustivity: the lexical item *red* (meaning 'partially red') does not contrast (in the sense of (75)) with *blue* (meaning 'partially blue'), nor does *comedy* (lexically/conceptually meaning 'either a true comedy or a tragicomedy') contrast with *tragedy* (lexically/conceptually meaning 'either a true tragedy or a tragicomedy'), and so on. That is, cohyponyms are lexically or lexically–conceptually weak, as argued for at length in chapter 2, and as such cannot form partitions. Prior to the effect of controlled exhaustivity, the assertion of one cohyponym does not exclude the assertion of another.

As I wrote above, Wagner is concerned with semantic processes occurring *above* the outcome of controlled exhaustivity; once *red* or *comedy* have been strengthened through their controlled Exh operator, they do in fact contrast with *blue* or *tragedy* in the sense of (75). Thus, Wagner's theory may well hold for free exhaustivity effects, but the proposal does not carry over to the domain of controlled exhaustivity.

4.3.3 Katzir's account: Innocent Exclusion

Katzir (2013, 2014) takes a different position, namely that predicates like *red* and *expensive* may very well be alternatives. He points out, for example, that while *red* and *cheap* may not *appear* to act as alternatives in cases like (73), they do appear to behave as alternatives under universals like *require*:

(76) Aisha is only required to own RED convertibles. (Katzir 2013:340)
a. $\rightsquigarrow$ Aisha is not required to own blue convertibles.
b. $\rightsquigarrow$ Aisha is not required to own cheap convertibles.
c. $\rightsquigarrow$ Aisha is not required to own expensive convertibles.

Katzir takes (76) to be the 'basic case' showing us that these predicates really can be alternatives; the puzzle is therefore why this is not intuited in examples like (73)–(74).

For Katzir, the effect in (73)–(74) is due to Innocent Exclusion (see chapter 1). The assertion in (77a) may well have the alternatives in (77b-i), but the ones with *cheap* or *expensive* are not innocently excludable (77b-ii). The idea is that something cannot both be 'not cheap' and 'not

expensive,' so neither alternative can be excluded.

(77) a. Exh_{ALT} [the car is red].
b. (i) ALT = {The car is red, The car is blue, The car is cheap, The car is expensive}
(ii) innocently excludable alternatives: {The car is blue}

The reverse presumably holds for sentences like (78). On Katzir's proposal, these do not carry inferences about colour terms because the alternatives with colour terms are not innocently excludable: a car cannot have no colour.

(78) a. Exh_{ALT} [the car is cheap].
b. (i) ALT = {The car is cheap, The car is expensive, The car is red, The car is blue}
(ii) innocently excludable alternatives: {The car is expensive}

Can Katzir's proposal be used to understand alternatives for controlled Exh, too? Recall we are trying to understand why (for example) *comedy* is incompatible with *tragedy*, but not with *short*:

(79) a. #This comedy is a tragedy.
b. This comedy is short.

I have been claiming in this thesis that this is because *tragedy* is an alternative to *comedy*, while *short* is not (for the alternatives associated with the controlled Exh). Taking inspiration from Katzir (and his discussion of *free* exhaustivity), one could rebuke that what stops non-cohyponymic predicates like *comedy* and *short* from becoming mutually exclusive is in fact Innocent Exclusion. Predicates which are non-cohyponymic to a given asserted predicate are not innocently excludable. Thus, (71) should be amended to (80) (keeping to the predicates *comedy*, *tragedy*, *short*, and *long*, for ease of presentation).

(80) a. This is a [Exh_{ALT} comedy].
b. (i) ALT = {comedy, tragedy, short, long}
(ii) innocently excludable alternatives: {tragedy}
c. $[\![(80a)]\!] = 1$ iff this is a (comedy & not tragedy).

The claim would be that a comedy cannot be neither long nor short, therefore these predicates are not innocently excludable. In this way, one could try to claim that the controlled exhaustivity of predicates can be understood without reference to a special, theoretically important relation between predicates.

4.3.4 Difficulties for the IE approach to alternatives for controlled Exh

There are several points to make about this alternative proposal. The first point touches on Katzir's argument for both the domains of free exhaustivity and of controlled exhaustivity. The issue is the following: it is not actually clear that Innocent Exclusion would really act as in (77) and (80). In the literature on gradable predicates like *cheap* and *expensive* or *long* and *short*, it has been noted (e.g., Alxatib & Pelletier 2011, Križ 2015) that sentences where both predicates are negated are acceptable:

(81) The car is neither cheap nor expensive.

(81) means that the car is in an extension gap between the two predicates: too expensive to be cheap, but too cheap to be expensive. If this is taken seriously, then in fact, (77a), with the alternatives in (77b-i), would mean that the car is neither cheap nor expensive. The sentence would come out to meaning that the car is in the extension gap in between these gradable predicates. This is clearly not the meaning we should be capturing. The point holds in exactly the same way with *long* and *short* (80). As such, 'this is a comedy' would mean that the comedy is not a tragicomedy, and what is more, it is of medium length (!).

As for (78), where colour terms are posited not to be innocently excludable upon the utterance of *cheap*, here too, I am not convinced that the colour terms are not innocently excludable, but for a different reason. My reason is different for free Exh and controlled Exh; let's start with free Exh. Recall that Wagner and Katzir are taking for granted that colour terms are universal: in my terms, they have already been locally exhaustified, so that *red* means 'entirely red,' and so on. In light of this, it is in fact perfectly possible to exclude all colour terms. They are universal, and it is possible for a car to not be *entirely* of a single colour. If the car is half white and half red, then it is neither entirely white nor entirely red (nor entirely of any other colour). As such, on their universal meaning (in my terms, once they have been ultra-locally exhaustified through their controlled Exh), colour terms actually *are* innocently excludable. If so, then (78a), with the alternatives in (78b-i), would end up meaning that the car is not entirely of a single colour: not all parts are red, not all parts are blue, and so on. 'The car is cheap' would mean that the car is not expensive, and moreover, it has more than one colour.[19]

The Innocent Exclusion approach to colour terms is also problematic for controlled Exh (i.e., Katzir's theory as applied to the empirical domain I am analyzing as involving controlled exhaustivity). Consider in particular (82). On the Innocent Exclusion approach, the idea is that *car* is not strengthened to exclude other predicates like colour terms.

(82) a. This is a $[\text{Exh}_{\text{ALT}}$ car$]$.
b. (i) $\text{ALT} = \{\text{car, train, red}_{\exists}\text{, blue}_{\exists}\text{, green}_{\exists}\text{, white}_{\exists}, \dots\}$
(ii) innocently excludable alternatives: $\{\text{train}\}$
c. $[\![(80a)]\!] = 1$ iff this is a (car & not train).

Since colour terms are existential prior to their controlled exhaustification, negating all colours would indeed mean that the car is colourless. As such, my previous criticism of Katzir's assumption that colour terms are not innocently excludable only applies in the domain of free exhaustivity, where colour terms have already been strengthened and are thus universal. In (82), excluding all colour terms would end up meaning that the car has no colour, in violation of our conceptualization of the world; we think of surfaces as having colours.[20]

Colour terms on this view are only not innocently excludable due to world knowledge or conceptualization. No logical contradiction arises if we exclude all colour terms. There is no logical issue with claiming that a surface has no colour; it is not like statements of the form 'p $\wedge\neg$p,' which are logically invalid regardless of the content of p. (83), corresponding to the strengthened meaning of *car* if colour terms are excludable alternatives, is only 'contradictory' in the sense that

[19] One way around this could be for Innocent Exclusion to refer to falsity conditions rather than truth conditions. Given that 'not red' means 'not read at all' (and so on), colours would not be innocently excludable because the result would be that the car has no colour (but see below).

[20] I am putting aside the tricky issue of transparency/invisibility.

it is at odds with the idea that cars necessarily have a colour.

(83) $\lambda x.\ \mathsf{car}(x) \wedge \neg\mathsf{train}(x) \wedge \nexists y[y \sqsubseteq x \wedge \mathsf{green}(y)] \wedge \nexists y[y \sqsubseteq x \wedge \mathsf{red}(y)] \wedge \nexists \ldots$

As such, if we are to avoid colour terms all being excluded in (82), we need Exh to make reference to world knowledge about colours. But this is incompatible with arguments made elsewhere in the literature. Indeed, Magri (2009) discusses the obligatory nature of exhaustivity effects precisely in terms of the interaction between these effects and world knowledge. In particular, he shows that Exh does *not* take world knowledge into account. Magri (2009) focuses on examples like (84), which are judged as deviant because *some* is strengthened to mean *not all*, even though all Italians come from the same country.

(84) #Some Italians come from a beautiful country.
$\rightsquigarrow$ not all Italians come from a beautiful country

It is not clear on what grounds Exh would fail to take world knowledge into account in (84), but behave intelligently with colour terms (and presumably at least some other cohyponyms), treating them as not innocently excludable due to world knowledge about surfaces having a colour. Thus, I think it is unclear at best whether Exh would treat colour terms as non-excludable, given that their collective exclusion does not lead to the kind of logical contradiction we expect Exh to actually take into account.[21]

The 'world-knowledge' domains of nationalities and colour terms could be argued to be two different types of knowledge. Intuitively, there could be something 'deeper' at play in our cognition about colours and surfaces, than in the knowledge we have about nationality. A sceptical reader could counter that I should not put Magri's data around nationality in the same category as anything to do with colour terms. I have two reasons not to believe that this is a proper counterargument. The first is weak and conceptual: distinguishing between world knowledge about colour terms and world knowledge about nationalities strikes me as making a rather thin cut. It may be doable, but it is a challenge I leave for those who wish to defend the Innocent-Exclusion theory of controlled exhaustivity. The second is more empirical: while I have focused on data around colour terms, my objection to taking Katzir's theory and applying it to controlled exhaustivity holds for predicates from other domains too. Consider genres. We need a predicate like *novel* to be strengthened in a way that does not exclude all genres in a sentence like (85a). The predicate *novel* is subject to controlled exhaustification (85b), so something must prevent the local Exh from excluding predicates like *comedy*, *tragedy*, and so on, in which case (85a) would mean something quite stronger than it does—stronger and indeed absurd.

(85) a. This is a novel.
b. This novel is #(also) a film.

[21] It is not clear how this argument fits with the definition of Innocent Exclusion given in chapter 1, where alternatives were innocently excludable if there is a possible world in which they are all false and the prejacent of Exh is true. Since this refers to possible worlds, it is possible that world knowledge could 'sneak in' in this way (perhaps there is no world in which cars have no colour at all). However, the same issue would arise for (84)—if 'Italians' is understood as the nationality and not the ethnic group (which is necessary for the judgment to go through; on the ethnicity reading, there is nothing wrong with claiming that some but not all Italian come from a beautiful country), then presumably in all worlds the members of a nationality come from the same nation.

On Katzir's approach, (85a) would not deny that the novel has a genre because genres, taken as a set, are not innocently excludable. But the fact that stories have a genre is intuitively quite a similar kind of world knowledge as the fact that all members of a nationality come from the same country. It is not a 'deep' cognitive process like our conceptualization of surfaces as necessarily having colours.

If what I have just argued is incorrect and Exh does in fact take world knowledge into account, Roni Katzir (p.c.) points out that a different problem would arise for controlled Exh, on the view that there is no relation determining sets of alternatives. At least, this is the case if the notion of Innocent Inclusion is correct (see chapter 1). If Exh includes alternatives that are not excluded and that can be consistently included, controlled Exh operators would strengthen predicates in unattested ways. For instance, colour terms, being lexically weak, can all be included consistently with world knowledge; thus, *fork* in (86a), being subject to controlled exhaustivity (86b), would end up being strengthened to *include* all colour terms; the utensil would be a fork that is partly of all colours.

(86) a. This is a fork.
b. This fork is #(also) a spoon.

This is what would arise if Innocent Inclusion is real and if Exh *does* take world knowledge into account (contrary to what I have just argued). On this view, because Exh takes world knowledge into account, it cannot innocently exclude all colour terms, but it can include them.

In sum, while Katzir provides data from *free* exhaustivity effects which suggest that there is no particular relation needing to hold between predicates for them to be alternatives to one another (e.g., (76)), this should not be extended to *controlled* exhaustivity effects. Katzir's Innocent-Exclusion approach would not yield the right results for controlled Exh. Assuming Exh does not take world knowledge into account, it would exclude far too many alternatives. What is more, I have also identified some problems with his account for free exhaustification. Of course, this raises the question of what to do with the datapoint with *require* that Katzir uses to motivate his view that all predicates can be alternatives; (87) is repeated from (76).

(87) Aisha is only required to own RED convertibles. (Katzir 2013:340)
a. ⇝ Aisha is not required to own blue convertibles.
b. ⇝ Aisha is not required to own cheap convertibles.
c. ⇝ Aisha is not required to own expensive convertibles.

It strikes me that the inferences in (87b–c) might not truly be entailed by the sentence. The speaker is simply not making a claim about such requirements, and we infer that no such requirements exist. It seems the statement in (87) can be answered by asking about whether there are other requirements not pertaining to colour:

(88) A: Aisha is only required to own RED convertibles.
B: Does it have to be {cheap, Italian, electric, big enough for four, second-hand, . . . }?

4.3.5 Interim conclusion

In this section, I have considered whether previous arguments about contrastively focused predicates can be adopted to describe alternatives for controlled Exh. For Wagner (2005, 2006, 2012), predicates can be alternatives if they represent cells in a natural partition. This cannot work to strengthen the weak meaning of cohyponyms (i.e., for controlled Exh) precisely because cohyponyms are underlyingly weak. While it is true that 'all green' excludes 'all white,' 'partially green' does not exclude 'partially white.' For Katzir (2013, 2014), there is no relation determining whether a predicate can be an alternative to a contrastively focused predicate. To the extent that we intuit otherwise, this is because certain sets of alternatives are not innocently excludable. For controlled Exh, the problem with this approach is that it requires Exh to take world knowledge into account in ways that it has independently been argued by Magri (2009) not to. It would have to take into account that all surfaces have a colour, that all stories have a genre, and so on, in order for sets of predicates like {*green*, *red*, *white*, ...} or {*comedy*, *tragedy*, *epic*, ...} not to be innocently excludable upon the utterance of other predicates. Without taking world knowledge into account, (89) would mean that the novel has no genre, for example.

(89) This is a novel.

Another issue arises if gradable predicates like *short* and *long* are alternatives for controlled Exh; they can in fact both be consistently negated, so (89) would come to mean that the novel, in addition to not having a genre, is of medium length.

In sum, some relation must hold between predicates for them to be alternatives for controlled Exh.

4.4 Jurisdiction in predication

In chapter 2, I argued that the alternatives for Exh operators controlled by nouns and adjectives are the controller's cohyponyms. I'll refer to this as the 'cohyponyms-as-alternatives hypothesis' (CAH), on the understanding that this is meant as a claim specifically about controlled exhaustivity, and not free exhaustivity effects.

(90) **Cohyponyms-as-alternatives hypothesis (CAH):**
In controlled exhaustivity effects, predicates are alternatives iff they are cohyponyms. Hence, *also* or *and* is required to co-predicate expressions when the cohyponymy relation holds.

We have just seen that the CAH is at least partly right in positing that *some* relation must exist between expressions for them to be alternatives for controlled exhaustivity. In this section, I ask whether cohyponymy is really the right relation. That is, are taxonomies (whether lexical—mentally 'fixed' due to conceptual structure—or context-dependent) the best way to characterize the alternatives for controlled Exh?

I answer in the negative. Clearly, cohyponymy is at least *insufficient* if we do not want to consider PPs and VPs to be 'cohyponyms.' But the problems go deeper. Even in the domain of simplex predicates, we will see that certain problems emerge when trying to account for the data in terms of cohyponymy. For example, some predicates referring to artefacts will cause a problem.

On the one hand, claiming that all artefacts count as cohyponyms for the CAH will undergenerate (some artefacts do not behave as alternatives, i.e. are not intuited as mutually exclusive), but attempts at having finer-grained taxonomies to separate artefacts overgenerate (many artefact that have nothing in common beyond being artefacts still display cohyponymic exclusivity, removable by *and* or *also*).

I therefore suggest an alternative relation to characterize the alternatives for controlled exhaustivity. The idea is that predicates are grammatically made mutually exclusive not when they are cohyponyms, but when they contribute *the same kind of information* in a sentence (a notion I will not formalise in this thesis). It happens that cohyponyms usually do contribute the same kind of information. I will call the kind of information provided by a predicate its 'jurisdiction.' For example, *red* has the jurisdiction of COLOUR, while *fork* has the jurisdictions FORM and FUNCTION (cf. the 'qualia' of Pustejovsky 1995).

For many predicates, the particular way they are used in a given sentence affects the sort of information they contribute. Jumping on this, I will return to some of the non-literalist readings discussed in chapter 2, and integrate them properly in the theory of predication being developed in this thesis, rather than dismissing them as pragmatic complications. Indeed, the jurisdiction of a predicate can in principle change from sentence to sentence; I will show that predicates are only interpreted as exclusive when they share a jurisdiction:

(91) a. #This shirt is a hat.
b. This shirt is {my hat, a good hat}.

Likewise, there are examples where PPs that apparently involve the same thematic role can be predicated of the same event; this is surprising if thematic uniqueness is a primitive notion, but less so if what matters if whether two PPs share a jurisdiction (which could be finer-grained than thematic roles, for instance due to Talmy's (1978) notion of figure and ground).

Beyond empirical advantages, the jurisdictional approach has the conceptual advantage of being more principled. The claim that cohyponyms are alternatives for controlled exhaustivity is no longer a primitive notion, but has roots in the general observation that different predicates contribute different kinds of information—a simple and obvious observation.

Of course, at the level of description, it is obvious that different phrases contribute different kinds of information, but the significance of my claim is that this is a theoretically important notion for controlled exhaustivity. A proper formalization of the notion of jurisdiction is left for future work, however, although I take limited steps in this section: taking inspiration from work in degree semantics, I tentatively formalize jurisdictions as predicates' first argument.

In this section, I first overview the appeal of the CAH and various apparent counterexamples that I argued in chapter 2 should not be taken to be serious problems for the CAH (section 4.4.1); then, I lay out two difficulties for the CAH beyond the existence of controlled exhaustivity with PP and VP alternatives (section 4.4.2); finally, I outline a notion of jurisdiction and show that it has several kinds of advantages (section 4.4.3).

4.4.1 The cohyponymic approach and its complications

The claim advanced in chapter 2 is that the alternatives for controlled Exh are *cohyponyms*. For example, *tragedy* is an alternative to *comedy* in (92) (repeated from (71)), but *green* is not.

(92) a. This is a [Exh_{ALT} comedy].
b. $\text{ALT} = \{\text{comedy, tragedy, epic, } \dots\}$
c. $[\![(92a)]\!] = 1$ iff this is a (comedy & not tragedy & not epic & …).

Let us review once again the arguments in favour of viewing cohyponymy as the relevant relationship for alternatives, and overview again the kinds of data that appear to counter the generalization, but that I cautioned in chapter 2 not to take too seriously as a core part of the semantics of predicates.

The appeal of cohyponymy as the relevant relation behind alternatives

As shown many times at this point, mutual predicational exclusivity is present with cohyponyms generally. These require conjunction or an additive particle in order to be mutually compatible. (93) shows this many times over, repeating some examples we have already seen in chapter 2, and adding some new ones as well. The bolded terms on the right-hand side suggest the sort of conceptual domain or taxonomy that the predicates come from.

(93) a. A tragicomedy is a tragedy that is #(also) a comedy. **(genres)**
b. The white flag is #(also) green. **(colours)**
c. A spork is a fork that is #(also) a spoon. **(utensils)**
d. Some live-action movies are #(also) animated. **(film type)**
e. (i) This car is #(also) a boat. **(vehicles)**
(ii) This train is #(also) a plane.
f. (i) Some snowshoes are #(also) skis. **(gear)**
(ii) There's a new kind of bicycle that is #(also) a skateboard.
g. (i) Some federal responsibilities are #(also) provincial. **(jurisdictions)**
(ii) City-states are cities that are #(also) countries.
h. Some residential neighbourhoods are #(also) industrial. **(zoning)**
i. (i) SCENARIO: *Apple starts selling computers with two operating systems.*
Now, some Macs are #(also) PCs. **(brands)**
(ii) SCENARIO: *McDonald's and A&W make a two-in-one fast-food joint:*
This McDonald's is #(also) an A&W.
j. He made a sling that is #(also) a bandaid.[22] **(medical equipment)**
k. Futons are couches that are #(also) beds. **(furniture)**
l. Cyborgs are humans that are #(also) robots. **(humanoids)**
m. Ben is an otter who's #(also) a twink. **(us gays)**
n. Are any derivational morphemes #(also) inflectional? **(morphology)**
o. Some left-wing ideas are #(also) right-wing. **(politics)**

As we have seen, the effect can hold in some cases even when a lexical item like *like* is required in order to deal with our world knowledge about certain taxonomies. At present, I've only found this

[22]This sentence was retrieved from social media.

where we know that certain things (species, languages) involve descent from a common ancestor.[23]

(94) a. (i) A platypus is like a beaver that {was, is} #(also) a duck. **(species/breeds)**
(ii) A labradoodle is like a poodle that {was, is} #(also) a labrador.
b. (i) Michif is like a Plains Cree dialect that {was, is} #(also) French. **(languages)**
(ii) English is like a Germanic language that {was, is} #(also) Italic.

The effect in (93) holds steady regardless of the particular predicate chosen from a taxonomy. For instance, whatever predicate is chosen from the domain of vehicles (93e), the exclusivity effect remains the same:

(95) a. This bike is #(also) a car.
b. This bike is #(also) a bus.
c. This bike is #(also) a train.

Even with *also*, some of the sentences in (95) correspond to truth conditions that are difficult to imagine, but with some imagination (e.g., for (95c), picture a very long communal tandem bike running on train tracks), they are all good with *also*.

Thus, there are good positive reasons to view the cohyponymic relationship as the culprit behind the behaviour of simplex predicates and *also*. The other side of the coin, of course, is that the effect disappears if we mix-and-match predicates from the different conceptual domains identified in (93). For example:

(96) a. Some live-action movies are comedies. **(film type + genre)**
b. The train is a provincial responsibility. **(vehicles + jurisdiction)**
c. Some industrial areas are a federal responsibility. **(zoning + jurisdiction)**
d. This robot is a car. **(humanoid + vehicle)**

Moving beyond the particular predicates in (93), I also attempted in chapter 2 to highlight the mutual compatibility of different-domain predicates by putting together predicates that are chosen precisely because they refer to things we do not except to overlap; (97) are some new such examples.

(97) a. Aisha's skin is a plant.
b. This mermaid is a figure skater.

These are odd, in the sense of requiring imagination and not being well suited to the real world, but they are far from being contradictions like the sentences in (93) as intuited without *also*.

[23] Michif is a so-called 'mixed language.' Roughly speaking, its verbs come from Plains Cree and its nouns come from French. Perhaps (i) is a more natural way to express (94b-i).

(i) Michif is **as if** Plains Cree was #(also) French.

Either way, (i) still requires *also* in the expected way.

Some things to be careful about in claiming that predicates are (not) compatible

As discussed in chapter 2, there are cases that seem to go against this trend, specifically sentences involving pseudo-repeated cohyponyms without a contradiction being intuited. But these examples can be thought of as simply being more complicated in various pragmatic or semantic ways, rather than being true counterexamples.

The first such apparent counterexample is the 'function-as' reading of predicates, were 'P is Q' means that something which is literally P (and not Q) serves the function of Q (and not P). The examples in (98a–b) are clear cases of this, and (98c) is probably one as well (among other things, this depends on whether *comic book* and *PhD dissertation* are really cohyponyms).[24]

(98) a. The fork is the spoon. *(about a camping set without a spoon)*[25]
b. This shirt is my hat.
c. This comic book is a PhD dissertation.

The second such case pertains to pseudo-repetitions where one cohyponym is true at one point in time or in one world, and the other is true at another point in time or in another world. For our purposes, we can simply analyze these examples as involving tense and world pronouns respectively. (99) is an example based around cohyponyms being true at different, non-overlapping times; t_1 is some time in the past, while t_0 includes the present.

(99) a. *A formerly entirely white shirt has emerged from the wash fully green.*[26]
The white shirt is green.
b. The [white t_1] shirt is [green t_0].

(100) is based around cohyponyms being true in different worlds, where w_0 is the real world of utterance and w_1 is the imaginary world of a theatre play.

(100) a. *We are setting up a play and decide to represent the character of a fox with a cat.*
The cat is the fox.[27]
b. The [cat w_0] is the [fox w_1].

Such examples are simply more complicated than the basic examples we are looking at; they do not counter the paradigm. Besides, notice that *white* still means 'entirely white,' and *green* 'entirely green,' in (99). The same goes with the world-pronoun cases:

(101) a. *We are setting up a play and decide to represent a red couch with a blue one.*
The blue couch is red.
b. The [blue w_0] couch is [red w_1].

While I will revisit my claim that the 'function-as' reading is merely a complication that should not influence our basic understanding of the paradigm, I think this claim stands firmly for the examples involving tense and world pronouns, and I will not be revisiting these examples.

[24] Looking ahead, I will be returning to examples like (98) later on, to claim that they might better not be dismissed in this way.

[25] Michael Wagner, p.c.

[26] I thank the audience at WCCFL 38 at UBC for this datapoint.

[27] Inspired by an example by Michael Wagner, p.c.

4.4.2 Two difficulties for the cohyponymy approach

In this section, I point to two general problems for the CAH. First, in a sort of reductio ad absurdum, I will show that many of the taxonomies proposed so far are in fact too narrow in scope. For the CAH to be sustained, we would in fact need a single grand taxonomy of all concrete form-denoting nouns (like *dog*, *plane*, *shirt*, etc.). It is not clear what the benefit of the notion of 'taxonomy' is at this point; it would be better to simply view such nouns as being grouped due to all providing the same kind of information about an individual (viz. its form). The second, stronger argument against the CAH is that it makes predictions about transitivity that are not borne out. If A and B are alternatives for controlled Exh, and so are B and C, then the CAH predicts A and C to be alternatives as well, because cohyponymy is a transitive relation. I will show that this is not always the case, however, requiring a more fine-grained notion than cohyponymy to explain predicates' behaviour.

These two problems (enormous, hard-to-identify taxonomies and the lack of transitivity for the alternatives of controlled Exh) come in addition to one of the main points to arise in this chapter, which is that alternatives for controlled Exh can be complex phrases of the form *about cats* or *killed Aisha*. These are not cohyponymic.[28]

A grand taxonomy of form-denoting concrete nouns

In this section, I show that the CAH makes it unavoidable to postulate a general taxonomy of form-denoting concrete nouns, something which is not well understood from the point of view of 'hyponymy' but better understood from the point of view of 'the sort of information a predicate contributes.' To be sure, the cohyponymy approach is hard to conclusively disprove: the approach does not make hard predictions, absent a clear theory of what sorts of taxonomies exist mentally, and whether these are fixed or can be created and/or changed in conversation depending on the situation. But I think the data in this section suggest that the CAH is at least unappealing. This section is essentially a reduction to absurdity.

I start the argumentation with data suggesting that some of the taxonomies identified in (93) are too narrow in scope. Indeed, (102a–c) mix-and-match predicates from various taxonomies identified in (93), and the sentences still require a clause-internal *also* to avoid a contradictory meaning. This is in contrast to the mix-and-matched predicates in (96) which did *not* require *also*.

(102) a. This couch is #(also) a car. **(furniture + vehicles)**
b. He made a sling that is #(also) a ski. **(medical + gear)**
c. Now, some Macs are ??(also) skateboards. **(brands + gear)**

[28]There is possibly a fourth problem for the CAH, also not discussed in this subsection (but see section 4.4.3), in the form of certain 'non-literalist' examples that I just dismissed (section 4.4.1 and chapter 2), like the following:

(i) a. This shirt is my hat.
b. CARTOON SCENARIO: *A horse has been turned into a sheep.*
That sheep is a horse.

Both examples in (i) involve two predicates from the same taxonomy (one would think), but they can be co-predicated consistently; this is not predicted by the CAH, although it could be compatible with it if we dismiss such examples as I did in section 4.4.1. However, in section 4.4.3, I will claim that these examples can be understood from the alternative to the CAH that I will propose, motivating that they should be viewed as a core part of the paradigm.

Given that *also* is required in (102), the predicates must act as alternatives. This is unexpected only if (*i*) predicates are alternatives if they come from the same taxonomy, and (*ii*) the taxonomies are those identified in bold on the right-hand side.

On the other hand, the nouns in these sentences can all be understood as being *artefacts*. While couches and cars or slings and skis are rather different kinds of artefacts corresponding to different spheres of life (the household, medical aid, motorised transportation, leisurely athletic transportation), they are all human-made objects made to serve particular functions. As such, (102) suggests that VEHICLES (93e), GEAR (93f), MEDICAL EQUIPMENT (93j), FURNITURE (93k), and even BRANDS (93i) should really be collapsed into a single grand taxonomy of ARTEFACTS; UTENSILS (93c) would be part of this taxonomy as well. Note that while brands are not artefacts per se, I am using the term 'brand' in (93) and (102) to refer to artefacts created by the corporation that has that brand; that is, 'Macs' in (102c) refers to computers, which are indeed artefacts.

Our grand taxonomy of artefacts quickly runs into problems, however. The first is that certain predicates which do not refer to artefacts would also have to be included, by virtue of requiring *also* to avoid contradictions. In (93), I had a 'HUMANOID' taxonomy to group predicates like *human* and *robot*; (103) is repeated from (93l).

(103) Cyborgs are humans that are #(also) robots. **(humanoid)**

This proposed taxonomy is not the most intuitively appealing: robots are in fact also artefacts, and many of them are not humanoid at all. These problems become clearer when we see that *human* is in fact also subject to a controlled exhaustivity effect with words like *plane*, which are quite clearly non-humanoid artefacts:

(104) Flying cyborgs are humans that are #(also) planes. **(humanoid + vehicles)**

Our grand ARTEFACT taxonomy, then, is too narrow for at least some examples. But it's gotten really unclear what conceptual domain is supposed to group not only *human* and *plane*, but also (presumably) *sling*, *ski*, *Macs*, and *skateboard*. The non-identifiability of certain taxonomies is a major problem for the CAH. The CAH centres on the claim that these predicates are cohyponyms, but what eactly is their common hypernym?

In my view, the take-away is that any concrete noun that contributes information about the form (and sometimes function) of an individual is an alternative to all other such concrete nouns. It's not just *human* and artefacts. The sentences in (105) are obviously cartoonish but make the point:

(105) The nation's mad scientists have just created ...
  a. a (very long) dog that is #(also) a ski.
  b. an apple that is #(also) a plane.

In contrast, concrete form-denoting nouns are *not* alternatives for (concrete) nouns contributing *other types of information*—as in the examples in (106), which involve a concrete form-denoting noun and another noun contributing information about what an individual does with their time. (106a) is repeated from (97b)).

(106) a. This mermaid is a figure skater. / This figure-skater is a mermaid.
  b. This man is a politician. / This politician is a man.

Mermaid provides information about the physical form of the individual, something which *figure skater* gives no information about; likewise with the predicates in (106b). As such, form-denoting concrete nouns are alternatives to one another, but not alternatives to other types of predicates like *figure-skater* or *politician*. The notion of taxonomy has no role in this description; what matters is the kind of information contributed by a predicate.

Transitivity in taxonomies

There is another serious drawback to using taxonomies as a way to model which expressions are alternatives for controlled Exh. Sisterhood in a taxonomy is necessarily a transitive relation: if A and B are sisters and so are B and C, then A and C must be sisters as well. However, we do not always find transitivity in alternatives for controlled Exh, specifically in cases involving different expressions contributing different kinds of information despite all being kinds of the same thing.

Let's start along the same empirical lines as above, by observing another difficulty in the identifiability of a taxonomy. Consider careers. While *doctor* and *bus driver* behave as the CAH predicts cohyponyms to behave, suggesting that there is a taxonomy of careers, *doctor* and *spy* do not:

(107) a. This doctor is #(also) a bus driver.
b. This doctor is (also) a spy.

Yet, being a spy is a job, too. It is not entirely clear what to do with this example, because the problem may only be apparent. Perhaps (107b) is acceptable without *also* because the fact that the doctor is a spy means that they are not a 'real' doctor (the example is really paraphrasable as 'the doctor is not a doctor at all, they are a spy'—so there *is* exclusivity).

However, the following minimal pair suggests that the issue is more general than the rather particular case of being a spy.

(108) a. This mathematician is an accountant.
b. This mathematician is #(also) a surgeon.

These are all cohyponyms referring to careers. But in (108a), the two predicates contribute different kinds of information: *mathematician* refers to training or education, while *accountant* refers to actual employment. The predicates do not act as alternatives for controlled Exh. On the other hand, in (108b), both predicates must at least contribute information about training (perhaps also current employment; this is less clear), and as such they do act as alternatives for controlled Exh. The contrast in (108) is not well captured on the CAH.

If *mathematician* and *surgeon* are alternatives for controlled Exh, and *mathematician* and *accountant* are not, the CAH predicts that *surgeon* and *accountant* would not be alternatives either, by transitivity. But in fact, *accountant* and *surgeon are* alternatives for controlled Exh:

(109) This surgeon is #(also) an accountant.

Again, due to the different training involved in both careers, this can be understood in terms of there being too much overlap in the kind of information contributed by the predicates. But on the CAH, given that cohyponymy is a transitive relation, (109) does not fit well with the data in (108). According to the CAH, we would have to claim that *mathematician* and *surgeon* (108b) are cohyponyms, as are *surgeon* and *accountant* (109); this would make the hard prediction that

mathematician and *accountant* should be cohyponyms, and therefore alternatives for controlled Exh. This is not borne out (108a).

The way to make sense of these data, I think, is that (108a) does not require *also* because *mathematician* can be interpreted as providing information about training (and not career), and *accountant* can be interpreted as providing information about career (and not training) *in the context that* being trained as a mathematician is viewed as good enough for a career as an accountant. This is not the case in (109), where the irrelevance of being trained as a surgeon to becoming an accountant makes it harder to interpret *accountant* as only providing information about the present career of the individual and not their training. In sum, it is possible to understand why these examples pattern the way they do, but doing so requires a finer-grained notion than merely cohyponymy.

4.4.3 Advantages and implementation of jurisdictions

Recall that the hypothesis we are investigating is the one in (110), repeated from (90).

(110) **Cohyponyms-as-alternatives hypothesis (CAH):**
In controlled exhaustivity effects, predicates are alternatives iff they are cohyponyms. Hence, *also* or *and* is required to co-predicate expressions when the cohyponymy relation holds.

We have just gone through two sets of data casting some doubt on the idea that cohyponymy is the relation at play in grammatically making predicates mutually exclusive. First, nouns that provide information about the physical form of an individual all act as alternatives for controlled Exh, while they do not act as alternatives with nouns that provide other sorts of information. Second, even among nouns providing types of information other than physical form (e.g., career nouns), whether or not they act as alternatives for controlled Exh depends on what exactly those nouns are intuited as entailing about an individual.

This suggests a generalization along the lines of (111) rather than (110), introducing the concept of 'predicational jurisdiction.'

(111) **A new hypothesis based on predicational jurisdiction:**
For controlled exhaustivity, predicates act as alternatives (and hence require *and* or *also*) if they contribute the same kind of information about their argument.

I will not be formalizing the notion of *kinds of information*; I hope the reader will accept it as intuitive, and leave for future work a better-defined approach making stronger predictions. The idea behind (111) is that, when we apply predicates to an individual, they do not (necessarily) apply in a 'totalitarian' way, but rather provide certain limited kinds of information about it. Predicates are assigned particular *jurisdictions* within a given sentence. I refer to two instances of predication that contribute the same kind of information as *same-jurisdiction predication*. Of course, it's descriptively not very interesting to point out that different predicates contribute different kinds of information (who would say otherwise?); what is new is that I claim that there is theoretical status to this notion. Finally, note that one important way that (111) is different from the CAH is that we are not dealing with the classification of predicates as such, but their uses in a given sentence; in principle, the same predicate can contribute different kinds of information in different sentences.

In this section, I make various sorts of comments about jurisdictions. It is organized as fol-

lows. First, expanding on the immediately preceding discussion, I make some caveats about 'form-denoting concrete nouns' (in particular what it means to be 'form-denoting'). I then point out three advantages to the jurisdictional approach. First, the 'function-as' use of artefactual nouns can be viewed as a core part of the semantics rather than a pragmatic complication; second, the open-ended nature of the alternatives for PPs and VPs can be collapsed with the open-ended nature of alternatives for nouns and adjectives; and third, apparent counterexamples to thematic uniqueness can be understood based on some PPs having different jurisdictions despite assigning the same thematic role to their complement. With this empirical discussion under our belt, I then give some limited formalization of how jurisdictions compose in the semantics, and in particular how they play out in creating contradictions with controlled Exh. Finally, I end by briefly connecting my notion of jurisdictions to Pustejovsky's (1995) notion of qualia in the lexicon, without thoroughly comparing and contrasting the notions.

Physical form vs. particular aspects of physical make-up

I have just claimed that form-denoting concrete nouns are all alternatives for controlled Exh. But of course, it could not be that any predicates that refer to any component of physical make-up are alternatives. For example, colour terms tell us something particular about the physical make-up of an individual, but they are not incompatible with the form-denoting nouns like *car*:

(112) This car is green.

The same goes for plenty of other types of predicates, like material terms (e.g., *metal*), terms relating to dimensions (e.g., *long*), and so on:

(113) a. This car is metal.
b. This car is long.

(112) and (113) are examples involving a noun and an adjective, but category is not the right construct to blame for the non-exclusivity of the predicates. There are sets of both adjectives (114a) and nouns (114b) that are mutually non-exclusive:

(114) a. This **long** car is **green**.
b. The **man** is a **politician**.

The issue is perhaps terminological: by 'form,' I am largely just referring to shape. True, certain concepts involving a form jurisdiction can also come with information about other prototypical aspects of physical make-up; for example, *cow* suggests an animal that is black, white, and/or brown, beyond just the shape of the animal. For simplicity, and to avoid turning this chapter into a research project on concepts and prototypes, I'll assume that to the extent that predicates like *cow* suggest things going beyond shape, these are incidental prototypical properties. This is probably wrong, and a better understanding of the 'holistic form' denoted by a noun like *cow*, but not an adjective like *green*, will ultimately be necessary to properly understand jurisdictions.

Another caveat: 'form' predicates like *car* or *cow* are compatible with adjectives relating quite literally to shape:

(115) a. This cow is circular.

b. This car is triangular.

But these examples mean that the cow is roundish as far as normal cows are concerned, and likewise with the car being triangular. They somehow only *weakly* modify the shape-related information contributed by the form-denoting noun. Unfortunately, I must leave for future work how to properly tease apart the notion of form contributed by nouns like *cow* and the different notion of shape contributed by adjectives like *circular*.

I now turn to various advantages of using predicational jurisdictions to describe the alternatives for controlled Exh.

Reintegrating some non-literalist examples

If the hypothesis in (111) is correct, this suggests it might have been better not to put aside the 'function-as' examples from section 4.4.1 as mere pragmatic complications rather than a central part of the paradigm. (116) is repeated from (98b).

(116) This shirt is my hat.

It is possible to view the acceptability of (116) as stemming from the fact that the two predicates, despite being cohyponyms (presumably), have different jurisdictions. Thus, (116) does not involve same-jurisdiction predication. Indeed, in the FORM jurisdiction, $\mathsf{shirt}(x)$; in the FUNCTION jurisdiction, $\mathsf{hat}(x)$. Moreover, the predicates are exhaustified in a way that makes them mutually incompatible once we control for jurisdictions. That is, the form in (116) is a shirt that is not a hat; the function is a hat that is not a shirt.

The same point goes for examples like (117), repeated from chapter 2.

(117) That sheep is a horse.

In chapter 2, I discussed how this example is only felicitous in a magical setting where a horse has been turned into a sheep (and perhaps vice-versa as well; I am not too sure). The animal has the form of a sheep but the inner essence of a horse. I pointed out in chapter 2 (footnote 8) that, unlike the other non-literal readings of sentences that I was putting aside entirely, this particular non-literal reading is not clearly reducible to independent components of language like tense pronouns, world pronouns, or metalinguistic uses of phrases. Rather, it seems to be an important observation about predicates that they can be true of an individual in one way and false of that same individual in another way. This is exactly what predicational jurisdiction are able to describe; *sheep* and *horse* are non-contradictory in (117) if they hold of a different jurisdiction (the form vs. the inner identity) but contradictory if they are both meant to be predicated on the same jurisdiction(s).

While a predicate like *shirt* has both the FORM and FUNCTION jurisdictions available to it, it seems that not all imaginable possibilities exist. The possibilities for the jurisdictions taken by predicates must somehow be limited (in ways I do not have a theory of). For example, (118) cannot be taken to mean that the car has the form of a car but the *size* of a tree:

(118) #This car is a tree.

There is no SIZE jurisdiction available to the predicate *tree*, apparently. I will briefly (but inconclusively) return to this point below when I compare my jurisdictions to the qualia of Pustejovsky

(1995).

Open-ended sets of alternatives: collapsing simplex and complex alternatives

There is another empirical advantage to analyzing the controlled exhaustification of simplex predicates like nouns and adjectives as involving sets of alternatives organized around predicational jurisdictions rather than cohyponymy. In particular, this helps to collapse the theory built for simplex predicates with the facts for complex predicates.

On the CAH, simplex predicates' alternatives were provided by the lexicon or the conceptual module, being peers in particular taxonomies, such as:

(119) a. COLOURS: {green, white, red, ...}
b. GENRES: {comedy, tragedy, epic, ...}
c. UTENSILS: {fork, spoon, knife, ...}
d. DOG BREEDS: {poodle, labrador, bulldog, ...}
e. FURNITURE: {couch, bed, table, ...}
f. LANGUAGES: {French, English, Plains Cree, ...}

These are more or less finite, well-defined sets. For the new hypothesis based in predicational jurisdictions, on the other hand, the alternatives are much more open-ended. Consider the general FORM jurisdiction, needed to capture examples like (120), repeated from (102a).

(120) This couch is #(also) a car.

For the form-jurisdictional predication in (120), is the set of alternative composed of literally all predicates that could refer to the physical make-up of anything? I leave this open, but simply show that this is in fact reminiscent of the alternatives for complex predicates.

Indeed, I noted in section 4.2 that the sets of alternatives for phrases like PPs are not cohyponymic, at least if cohyponyms are thought of as involving cognitively more or less fixed taxonomies as organised in accordance with world knowledge about kinds. With PPs, the alternatives are simply other PPs, obtained by replacing the preposition's complement with other DPs, which do not need to come from a particular taxonomy.

(121) a. The letter $[\text{Exh}_{\text{ALT}}$ for $[\text{Aisha}]_F]$ is #(also) $[\text{Exh}_{\text{ALT}}$ for $[\text{Ben}]_F]$.
b. ALT = {for Aisha, for Ben, for ...} ($\Rightarrow$ all possible recipients)

The jurisdictional theory makes the thematic-role alternatives look much more similar to the alternatives for predicational jurisdictions: in both cases, we are in fact dealing with effectively open-ended sets, and in both cases the alternatives are properly characterized as involving the same kind of information, moreso than as being cohyponymic.

Something similar goes for examples like (122), where the two alternative predicates do not come from hard-wired conceptual taxonomies: they involve novel phrases created by the open-ended power of generative syntax. On the normal view, cohyponymy is a lexical relation between (lexicalized) predicates, and as such cannot relate *floor wax* to *dessert topping*.[29]

[29](122) is a modified sentence from a *Saturday Night Live* sketch: `https://snltranscripts.jt.org/75/75ishimmer.phtml` (accessed November 21, 2019). I thank Jacob Hoover for suggesting this sentence.

(122) This floor wax is #(also) a dessert topping.

At the same time, the CAH could be specified to be flexible enough to allow for ad-hoc taxonomies being created in particular discourse conditions, in which case (122) would not be a problem for it as such. Either way, the jurisdictional approach can deal with (122) effortlessly, since *floor wax* and *desert topping* both contribute information about the function and make-up of a substance.

Jurisdictions and counterexamples to thematic uniqueness

The jurisdictional approach has another advantage, this time inherent to the data around thematic uniqueness. Recall that sentences like (123), repeated from (33a), constitute a problem for thematic uniqueness as an empirical generalization, on the hypothesis that *the rock* and *the tree* are both themes (given the stative nature of the predicate).

(123) The rock is to the left of the tree. (Dowty 1991:563)

While they have the same thematic role, *the rock* is the figure whereas *the tree* is the ground (Talmy 1978); in this sense, they contribute different kinds of information. Figure and ground are not thematic roles (Dowty 1991), but they can be reasonably conceived of as different jurisdictions.[30] Here, of course, we are not talking about the kind of information given by a predicate to an individual, but to an event or state.

Also consider apparent counterexamples to thematic uniqueness like (124):[31]

(124) a. #Aisha is eating with a fork with a spoon.
b. Aisha is eating with a fork with her left hand.

In both examples, both *with*-PPs are instrumental, but the illicitness associated with thematic uniqueness violations is not intuited in (124b). With the right set of jurisdictions, this could be understood in terms of *with a fork* and *with her left hand* not sharing a jurisdiction; it's not clear what those jurisdictions would be, however. But this is at least possible in principle, unlike on a purely thematic approach.

Something similar goes with other PPs, such as the *about*-PPs in (125):

(125) a. This book **about cats** is #(also) **about bicycles**.
b. This book **about cats** is (??also) **about friendship**.

The concept of thematic uniqueness is too coarse to make the cut between these examples. But on the jurisdictional approach, it is intuitively appealing that the *about*-PPs in (125a) are both 'the same kind' of topic (the literal thing that the book discusses) whereas the topics in (125b) are not quite the same kind: *cats* is the literal subject of discussion whereas *friendship* is the moral theme of the story. With the right approach to jurisdictions, it could be possible to claim that the PPs share a jurisdiction in (125a) but not (125b).

[30] Unfortunately, the word 'reasonably' is pulling a lot of weight in this sentence, given that I have not given a formal or even particularly specific definition of what can and cannot count as a predicational jurisdiction.

[31] I thank the audience at McGill's syntax/semantics reading group for (124b).

Some formalization: jurisdictional arguments

We have started putting together a notion of jurisdiction and seeings its advantages for controlled exhaustivity. In this subsection, I offer some limited formalization and spell out exactly how jurisdiction interacts with Exh.

In the standard approach to predication, a noun like *car* simply means (126), where car denotes the set of cars (or to put the matter in clearer terms: the set of exemplars of the concept referred to as *car*; see chapter 1).

(126) $[\![\text{car}]\!] = \lambda x.\text{car}(x)$.

I argued in chapter 2 that taxonomic predicates like *car* in fact undergo exhaustification, making (126) an oversimplification. But there is now a second way in which (126) looks like an oversimplification. As a matter of fact, we want *car* to be able to contribute different things in different sentences, so it must have meanings more akin to those in (127).[32]

(127) $[\![\text{car}]\!] =$

a. $\lambda x.\ x \in \{y : y \text{ has all the properties of a car}\}$.
b. $\lambda x.\ x \in \{y : y \text{ has the form of a car}\}$.
c. $\lambda x.\ x \in \{y : y \text{ has the function of a car}\}$.

That is, *car* can denote the set of things that have both the form and function of a car, the set of things that have the form (but not necessarily the function) of a car, or the set of things that have the function (but not necessarily the form) of a car. In addition to the examples given so far in this section, we can observe this specifically for *car* in examples like those in (128):

(128) a. This is a car.
⇒ *form and function of a car*
b. This (broken-down) car is a good trampoline.
⇒ *form of a car*
c. This couch with wheels is Aisha's car.
⇒ *function of a car*

These are the different 'jurisdictions' available to *car*. Jurisdictions, then, mean that the domain of individuals can be partitioned in different ways; a given individual could fall in the cell of cars by one partition of the domain, but outside of this cell by another partition.[33]

Generalizing from the set of possible meanings for *car* in (127), what we want is for predicates

[32] I will be using rather informal paraphrases in the coming discussion; early attempts at formalizing things further led to undue complexity for the purposes of this chapter.

[33] An anonymous reviewer for *Sinn und Bedeutung* 27 suggests using category theory to model this. They write that category theory makes it possible that

> you have the same set of objects but various operations that you can apply to them, which gives you the option to 'analyze' them differently, thus leading to a) a different vocabulary and b) different interpretations of the same vocabulary. Moreover, category theory gives you perfect tools to model transition from one jurisdiction to another.

I leave this as a promising avenue for future research.

like *car* to have output conditions like in (129):

(129) $x \in \{y : y \text{ has the (jurisdiction) of a (predicate)}\}$.

As such, predicates should have lexical entries like (130), where their particular jurisdiction in a particular sentence is set by an abstract jurisdictional argument j. I am taking inspiration from the literature on degree semantics, where degree predicates are often suggested to take a degree as their first argument.

(130) $[\![\text{car}]\!] = \lambda j.\lambda x.\ x \in \{y : y \text{ has the } j \text{ of a car}\}$.

Let's see how this works, and in particular how this relates to controlled exhaustivity. Let's take two predicates that are most saliently understood as making a claim about the form and function of an individual:

(131) a. $[\![\text{couch}]\!] = \lambda j.\lambda x.\ x \in \{y : y \text{ has the } j \text{ of a couch}\}$.
b. $[\![\text{car}]\!] = \lambda j.\lambda x.\ x \in \{y : y \text{ has the } j \text{ of a car}\}$.

In a contradictory sentence like (132a), both these predicates are given both FORM and FUNCTION for j and are then locally exhaustified, creating a contradiction:[34]

(132) a. #This couch is a car.
b. (i) $[\![\text{FORM}\oplus\text{FUNCTION couch}]\!]$
$= \lambda x.\ x \in \{y : y \text{ has the form and function of a couch}\}$.
(ii) $[\![\text{Exh}_{\text{ALT}}\ [\text{FORM}\oplus\text{FUNCTION couch}]]\!] =$
$$\lambda x.\ \begin{cases} x \in \{y : y \text{ has the form and function of a couch}\} \wedge \\ \quad x \notin \{y : y \text{ has the form and function of a car}\} \wedge \\ x \notin \{y : y \text{ has the form and function of a pillow}\} \wedge \\ \qquad\qquad x \notin \ldots \end{cases}$$
c. (i) $[\![\text{FORM}\oplus\text{FUNCTION car}]\!] = \lambda x.\ x \in \{y : y \text{ has the form and function of a car}\}$.
(ii) $[\![\text{Exh}_{\text{ALT}}\ [\text{FORM}\oplus\text{FUNCTION car}]]\!] =$
$$\lambda x.\ \begin{cases} \quad x \in \{y : y \text{ has the form and function of a car}\} \wedge \\ x \notin \{y : y \text{ has the form and function of a couch}\} \wedge \\ x \notin \{y : y \text{ has the form and function of a pillow}\} \wedge \\ \qquad\qquad x \notin \ldots \end{cases}$$

Once exhaustified, these predicates are contradictory due to overlapping in jurisdiction. In contrast, non-contradictory sentences like (133a) involve a different jurisdictional arguments for each predicate:

(133) a. This couch is my car.
b. (i) $[\![\text{FORM couch}]\!] = \lambda x.\ x \in \{y : y \text{ has the form of a couch}\}$.
(ii) $[\![\text{Exh}_{\text{ALT}}\ [\text{FORM couch}]]\!] =$

[34] I am assuming that the alternatives do not vary in the j argument. That is, the alternatives for the phrase 'FORM⊕FUNCTION car' are all of the form 'FORM⊕FUNCTION P.' It might be possible to derive this by having Exh below the jurisdictional argument, rather than above it as in the main text.

$$\lambda x.\ \begin{cases} x \in \{y : y \text{ has the form of a couch}\} \wedge \\ x \notin \{y : y \text{ has the form of a car}\} \wedge \\ x \notin \{y : y \text{ has the form of a pillow}\} \wedge \\ x \notin \ldots \end{cases}$$

c. (i) $[\![\text{FUNCTION car}]\!] = \lambda x.\ x \in \{y : y \text{ has the function of a car}\}$.

(ii) $[\![\text{Exh}_{\text{ALT}}\ [\text{FUNCTION car}]]\!] =$

$$\lambda x.\ \begin{cases} x \in \{y : y \text{ has the function of a car}\} \wedge \\ x \notin \{y : y \text{ has the function of a couch}\} \wedge \\ x \notin \{y : y \text{ has the function of a pillow}\} \wedge \\ x \notin \ldots \end{cases}$$

The result for (133) is that the artefact is asserted to have the form but not the function of a couch, and to have the form but not the function of a car. Hence, there is no contradiction.

This provides some formalization of jurisdictions. Of course, many questions remain, which I must leave for future work. First, there is the question of what exactly is the set of jurisdictions. Presumably it is quite large. Recall from section 4.4.2 that I argued that the career nouns in (134) (repeated from (108)–(109)) are (in)compatible due to whether they contribute information about training, employment, or both; these notions would therefore be jurisdictions, on the current view. This means that the set of jurisdictions is not a small set of cognitive primitives but rather a large set whose membership is richly informed by world knowledge.

(134) a. This mathematician is an accountant.
b. This mathematician is #(also) a surgeon.
c. This surgeon is #(also) an accountant.

Likewise, in my discussion of how subtle differences in the meaning of PPs can affect whether they are interpreted as thematic-uniqueness violation effects, I pointed to contrasts such as (135), repeated from (124).

(135) a. #Aisha is eating with a fork with a spoon.
b. Aisha is eating with a fork with her left hand.

I have not given any formalization of how jurisdictions work with PPs as predicates of events (rather than individuals), but these would somehow have to make finer-grained distinctions than only thematic roles.

A second important unanswered question is what constrains the kinds of jurisdictions that a predicate can have. Recall for example that the noun *tree* cannot have LENGTH as a jurisdiction; (136), repeated from (118), cannot mean that the car has the length of a tree.

(136) #This car is a tree.

I return to this issue shortly below. Still, it is important to note that the semantics I have just proposed overgenerates in not constraining the kinds of jurisdictions that a given predicate can have.

Jurisdiction and Pustejovksy's qualia structure

I end this section by pointing out a link to the existing literature, namely the 'qualia structure' in lexical items that Pustejovsky (1995) proposes. Qualia structure refers to "the representation of the defining attributes of an object, such as its constituent parts, purpose and function, mode of creation, etc." (Pustejovsky 1995:3). Pustejovsky uses this notion to capture the 'sense in context' of predicates as they combine together. Some but not all of the data he uses this notion to discuss is reminiscent of my discussion.

At the less clearly related side of things, he uses qualia structure to predict the different meanings of *use* in the following examples (Pustejovsky 1995:87), and how they are conditioned by the direct object:

(137) a. Aisha *used* **the new knife** on the turkey. (*use* ≈ *cuts with*)
b. Aisha has *used* **soft contact lenses** since college. (*use* ≈ *wears*)
c. This car *uses* **unleaded gasoline**. (*use* ≈ *burns*)
d. My wife *uses* **the subway** every day. (*use* ≈ *travels on*)

He argues that *use* is semantically underspecified and 'reacts' to the qualia structure of its complement. For instance, if *knife* is specified as having the telos of cutting, the notion of cutting is present in *knife*, which *use* picks up on.

In examples like (137), Pustejovsky uses qualia to talk about the interplay between predicates, rather than the referents themselves (which is what I stipulate jurisdictions for, at least for nouns and adjectives). But Pustejovsky comes closer to my claim with examples like (138) (Pustejovsky 1995:91):

(138) a. John crawled through **the window**.
b. Mary broke **the window**.

The alternation of *window* as referring to the aperture vs. the physical object is, in fact, a distinction about the referent itself.

I leave for future research a stronger theoretical connection between Pustejovsky's argumentation (and data) and what I have proposed in this section. The only point I will make here is that as Pustejovsky (1995) implements it, qualia structure is integrated in the meaning of lexical items. The predicate *shirt*, for instance, would have both information about the form of shirts and the function of shirts as part of its lexical meaning. This differs significantly from my suggestion above that jurisdictions are arguments of predicates. I made this claim in order to capture that *shirt* can refer to anything with the function (but not necessarily the form) of shirts in certain sentences; it is not clear how this could arise if both the form and function of shirts are part of the lexical entry for *shirt*. In this sense, Pustejovsky's (1995) qualia structure is the mirror image of jurisdictions.

On the other hand, one advantage to Pustejovsky's approach is that it provides a natural way to delimit the possible jurisdictions for a given predicate. I noted above that examples like (139), repeated from (118)/(136), show that not all imaginable jurisdictions are actually observed for a given predicate. Specifically, (139) cannot mean that the car has the *length* of a tree, which would be expected if *tree* took LENGTH as its jurisdictional argument (in which case the denotation of *tree* would be the set of all things, tree or otherwise, that are the length of a tree).

(139) #This car is a tree.

A response to this inspired by Pustejovsky's qualia is that this is because there is nothing in the lexical representation of *tree* touching on length. In contrast, *shirt* corresponds to an artefact and therefore comes lexically with information about form and function. If jurisdictions/qualia are lexically represented, it is possible to model limits on possible jurisdictions lexically. To capture both that predicates lexically put limits on their possible jurisdictions, and that jurisdictions are not fully lexical (since they can vary sentence by sentence), it may be possible to require 'agreement' or overlap between lexical items and their jurisdictional arguments. I leave this for future work.

4.4.4 Interim conclusion

In section 4.3, I defended the view that there is a particular sort of relation that must hold between two phrases for them to be interpreted as mutually exclusive in most cases, unless conjoined by *and* or *also*. What sort of relation could this be? In this section, I revised the hypothesis from chapter 2, which was that the relation in question is cohyponymy. I argued that restricting the grammatically produced mutual-exclusivity effect to cohyponyms does not produce as many contradictions as we actual intuit. Examples like (140) do not involve cohyponyms, but still require expressions like *also* to be acceptable.

(140) This sling is #(also) a ski.

Using the intuition that all form-denoting concrete nouns appear to be alternatives for controlled Exh, I suggested that the notion of the sort of information an expression contributes is theoretically important. To state the obvious, predicates are not applied to individuals in 'totalitarian' ways, but rather, they contribute limited types of information, such as what the form of that individual is, what its function is, and so on. In (140), both predicates contribute information about the form of the individual, and therefore act as alternatives to one another.

Much remains for future work, in particular a formalization of what exactly it means for predicates to contribute 'different kinds of information.' There is also an empirically interesting domain that remains to be investigated. Indeed, different readings of predicates seem to come and go depending on fairly subtle factors. Consider the sentences in (141):[35]

(141) a. #This shirt is a hat.
b. This shirt is my hat.
c. This shirt is a good hat.

In (141a), only the form jurisdiction is available for *hat*, but this disappears by using a possessive or modifying it with *good* (141b–c). In addition to the sorts of questions I have already identified, an interesting path forward for work on predicational jurisdictions could start by identifying the semantic environments that allow certain readings, and those that do not.

4.5 Optional and unacceptable clause-internal additivity

As a last piece of commentary on the nature of alternatives in controlled exhaustivity effects, I move in this short section from discussing alternatives for Exh to discussing alternatives for one of

[35] I thank Michael Wagner for emphasizing these kinds of data to me.

the expressions removing problems created by Exh, namely *also*.

This thesis focuses largely on the obligatory additive effects found with predicates. These are 'clause-internal' additivity effects, in that the additive's presupposition is entailed by the subject. That is, 'α is also β' entails that α is α. But there are also clause-internal additivity effects with predicates which are just about optional, and others that are impossible. In this section, I claim that this can be understood if the relation governing what can be an alternative for controlled Exh is the same as the one governing what can be an alternative for clause-internal *also*.

Let's consider optional and illicit clause-internal additivity in turn. We start with the optional *also* in (142).

(142) Bears are (?also) scavengers.

In (142), *also* here is perfectly possible—but a little odd. It makes it sound as though the speaker considers that being a bear and being a scavenger are somehow comparable, 'on a par,' or two options in the same paradigm. Of course, being a bear and being a scavenger aren't really any of those things, at least in the sort of conversational contexts one imagines to make sense of (142) out of the blue. The distinction in meaning corresponding with the presence of *also* is subtle, but I think real.

This has an interesting consequence. Indeed, the presence of *also* must add more than just *also*'s additive presupposition, which is entailed by the sentence and therefore not expected to correspond to any detectable changes in meaning. The most obvious path forward is to claim that the presence of *also* corresponds to whether *bear* and *scavenger* are alternatives to one another. If they are not alternatives, *also* is not needed, and in fact not licensed, presumably due to its presupposition not being met. If they are alternatives, we intuit them as being 'on a par,' and *also* is not optional but *required* to prune the alternatives of the controlled Exh operators on each predicate. On this view, *also* in (142) is not really optional: it is either illicit or obligatory, in correspondence with whether *bear* and *scavenger* are alternatives.

On the theory developed in chapter 2, it's not clear how to make sense of this. On the simplest approach (where taxonomies are provided by the conceptual space, and are not mutable according to particular conversations), *bear* and *scavenger* are either cohyponyms (and therefore alternatives), or they are not. A more complex approach to cohyponymy would allow discourse to spontaneously create ad hoc taxonomies. In this case, *bear* and *scavenger* could in principle be cohyponyms. But what would be the nature of the taxonomy in which they are sisters? The point should be that, for *bear* and *scavenger* to be alternatives, they must be understood as contributing the same kind of information in a conversation, i.e. share a jurisdiction. Modelling this as a taxonomy is beside the point.

Beyond apparently optional additivity effects like (142), there are also predicates which (outside of special contexts) refuse the presence of *also* altogether:

(143) Some cooks are (#also) Italians.

Of course, on the theory developed here, (143) is not surprising: *cook* and *Italian* do not contribute the same kind of information in most contexts. The fact that *also* is disallowed indicates that these predicates are not alternatives for *also* any more than they are alternatives for controlled Exh. Crucially, (143) improves in a context where they do contribute the same kind of information. For example, imagine a job posting looking for an applicant who has one of the following three

properties: being a cook, being Italian, or being LGBT+. (144) is fine in this scenario:

(144) We should hire Aisha. She's a cook who's also Italian!

Note that the additive in (144) is optional, however; perhaps there is optionality about whether *cook* and *Italian* are alternatives for the purposes of Exh and *also* even in this context.

In sum, *also* is blocked from copular sentences unless the subject and predicative predicates share a jurisdiction. This means that the notion of jurisdiction is not only at play in determining what is an alternative for Exh, but also for *also*. We need to claim that Exh's alternatives are same-jurisdiction predicates because predicates are understood as contradictory when they share a jurisdiction. But we also need to claim that *also*'s alternatives are same-jurisdiction predicates (for clause-internal additivity effects) in order to capture that *also* is illicit if it lacks a same-jurisdiction antecedent (143).

4.6 Conclusion

This chapter has made progress on the profile of controlled exhaustivity in two ways.

First, I showed that controlled Exh is observed with phrases (PPs and constituents that might be identifiable as VPs). The fact that controlled Exh exists with PPs lead to the conclusion that thematic uniqueness is generally due to controlled exhaustivity. I focused less on the observation of controlled Exh with VPs, but I think the data are clearly promising as a way to bridge the empirical discussion in the present thesis with obligatory-additive effects with contrastive topics as discussed by Krifka (1998), Sæbø (2004), Bade (2016), and Aravind & Hackl (2017); I leave this for future work.

Second, I showed that the notion of cohyponymy is not quite right in characterizing the alternatives for controlled Exh. First of all, PPs and VPs are not cohyponyms in the normal sense. Second, even in the domain of nouns and adjectives, what seems to be at play is whether two predicates contribute the same kind of information in a given sentence, rather than whether they belong to a particular identifiable taxonomy. Form-denoting concrete nouns like *dog* and *ski* are all alternatives for controlled Exh, for instance; these are not cohyponyms (what would be their common hypernym?), but they all contribute information about the physical form of an individual. I therefore suggested to give theoretical status to jurisdictions; much like degrees for degree predicates, I suggested that predicates take a jurisdiction as their first argument. This can capture the ambiguity of certain predicates. For example, artefactual predicates can refer to an object's function, form, or both, as determined by the jurisdictional argument they take in a given sentence.

In the two remaining chapters of this thesis, I turn first to a particular empirical subset of controlled exhaustivity effects, namely such effects with colour terms (chapter 5). Then, in chapter 6, I suggest a theory of why Exh is ever 'controlled,' i.e. obligatory with and necessarily local to certain alternative-triggering expressions.

Chapter 5

Controlled Exh and subatomic homogeneity

5.1 Introduction

So far in this thesis, I have discussed a variety of predicates to show that their meaning is systematically strengthened in basic sentences through the exclusion of related predicates. In doing so, I have grouped together all taxonomic predicates. Descriptively, some differences between taxonomic predicates exist, but I have not treated these differences as having theoretical significance for controlled exhaustivity. In particular, consider the pair in (1):

(1) a. The white flag is #(also) green.
b. This comedy is #(also) a tragedy.

The effect in (1a) touches on quantification over the parts: without the additive, the sentence means that all parts of the flag are white, and all parts are green, clashing with world knowledge about surfaces and colours. This is not the case in (1b), which has to do with the set of individuals that count as comedies or tragedies: the predicates have larger extensions in the presence of *also* than without it. With the additive, (1b) is therefore intuited as referring to an exemplar which is about halfway between a prototypical comedy and tragedy. This is quite different from (1a), where the additive does not allow the sentence to mean that the entire flag is halfway between greenness and whiteness (i.e., pale green); rather, it means that the flag has a part which is a true green, and another part which is a true white.

Despite this descriptive difference between the sentences in (1), I have collapsed them in previous chapters as two instances of the same phenomenon of controlled exhaustivity. After all, all else being equal, it is better to give a single unified explanation to the paradigm in (1a) and (1b) than to give each example a different explanation. In both cases, I have suggested weak lexical meaning paired with the exclusion of related predicates via ultra-local exhaustification, as in (2) (where colour terms are subscripted with $\exists$ to emphasize their existential lexical meaning).

(2) a. $[\![\text{The } [\text{Exh}_{\text{ALT}} \text{ white}] \text{ flag is } [\text{Exh}_{\text{ALT}} \text{ green}]]\!]$
$= 1$ iff the $\begin{pmatrix} \text{green}_{\exists} \,\& \\ \text{not white}_{\exists} \,\& \\ \text{not red}_{\exists} \end{pmatrix}$ flag is $\begin{pmatrix} \text{white}_{\exists} \,\& \\ \text{not green}_{\exists} \,\& \\ \text{not red}_{\exists} \end{pmatrix} \Rightarrow$ contradiction
b. $[\![\text{This } [\text{Exh}_{\text{ALT}} \text{ comedy}] \text{ is a } [\text{Exh}_{\text{ALT}} \text{ tragedy}]]\!]$

$$= 1 \text{ iff this } \begin{pmatrix} \text{comedy } \& \\ \text{not tragedy } \& \\ \text{not epic} \end{pmatrix} \text{ is a } \begin{pmatrix} \text{tragedy } \& \\ \text{not comedy } \& \\ \text{not epic} \end{pmatrix} \Rightarrow \text{contradiction}$$

For 'non-quantificational' predicates like *comedy*/*tragedy*, postulating 'weak lexical meaning' is a claim about the concepts referred to by these predicates: they overlap in the conceptual space. That is, lexically–conceptually, *tragedy* and *comedy* have a non-empty intersection. For quantificational predicates like colour terms, 'weak lexical meaning' is a claim about quantification in the lexicon: colour terms are lexically existential (Harnish 1976; Levinson 1983):

(3) $[\![\text{green}]\!] = \lambda x.\exists y[y \sqsubseteq x \wedge \mathsf{green}(y)]$.

Notably, on this view, simple sentences like (4), while intuited as involving universal quantification, are not 'really' universal.

(4) The flag is green.

(4) means that the flag has a green part; following the exclusion of other colour terms through Exh, it means that the flag has a green part and does not have a part of any other colour:

(5) $[\![\text{The flag is } [\text{Exh}_{\text{ALT}} \text{ green}]]\!] = 1 \text{ iff the flag is } \begin{pmatrix} \text{green}_{\exists} \ \& \\ \text{not white}_{\exists} \ \& \\ \text{not red}_{\exists} \end{pmatrix}$

Given world knowledge that all areas of a surface must have a colour, (5) is pragmatically strengthened to mean that the flag is entirely green, in the same way that (6) with an overt *only* means that the flag is entirely green (presumably due to *only* excluding other colour predicates; (6) is virtually synonymous with (4), so it is not clear what else *only* would be doing).

(6) The flag is only green.

This claim essentially follows the spirit of Harnish (1976) and Levinson (1983) in strengthening colour terms through the exclusion of other colour terms, although as discussed in chapters 2 and 6, a substantial difference arises due to the necessity of claiming that the exhaustification of colour terms is *obligatorily local*—just as is the case with other taxonomic predicates like *comedy*/*tragedy*.

Would it be better to analyze (4) as involving true, *semantic* universal quantification over parts? In fact, there is a substantial literature analyzing the intuited meaning of colour terms in sentences like (4) as semantically universal. This is the literature on homogeneity (all-or-nothing) effects (see e.g. Löbner 2000, Križ 2015, 2019, Bar-Lev 2021, Križ & Spector 2021, and citations therein for homogeneity effects in general, although not all these sources discuss colour terms). The homogeneity literature usually focuses on examples like (7), involving a predicate composing with a plurality. The observation about such sentences is that the predicate is implied to hold of all atomic parts of the plurality in positive sentences (7a), and to hold of none of them in negative ones (7b).[1]

[1] Pragmatic factors can weaken the effect to allow exceptions. We will return to this in section 5.4.4.

(7) a. Aisha saw the children.
≈ she saw all of the children
≉ she saw at least some of the children
b. Aisha didn't see the children.
≉ she didn't see all of the children
≈ she saw none of the children

However, even though it is most often discussed for pluralities, the homogeneity paradigm is also found with atoms; the go-to example involves colour terms (e.g., Löbner 2000, Križ 2015).

(8) a. The flag is green.
≈ all of the flag is green
≉ at least some of the flag is green
b. The flag isn't green.
≉ not all of the flag is green
≈ none of the flag is green

I will refer to the effect in (7) and (8) as 'plural' and 'subatomic' homogeneity respectively, and refer to both the universal quantification in positive sentences and the negated existential quantification of negative sentences as a 'homogeneity' effect.

As I will show, the appeal to world knowledge present in my own proposal for colour terms does not carry over to the plural homogeneity paradigm. It is therefore sensible for theories of plural homogeneity to take positive sentences to involve universal quantification. Yet, in this chapter, I will defend that subatomic homogeneity is best analyzed as I have proposed, without real universal quantification. In particular, I will discuss two theories of homogeneity in addition to my own proposal; while I will not make any claims about the theories' validity for plural homogeneity, I will show that they come with substantial problems for subatomic homogeneity. The two theories I will focus on are the presuppositional account of homogeneity (e.g., Löbner 2000), which claims that predication involves a presupposition that the predicate is true of all or none of the (atomic or subatomic) parts of its subject, and Bar-Lev's (2018; 2021) exhaustivity-based account of homogeneity, which creates the homogeneity effect through the Innocent Inclusion of subdomain alternatives. The upshot of this discussion is that a clearer picture will emerge of whether it is right to collapse plural and subatomic homogeneity (as first done by Löbner 2000, and taken up by Spector 2013 and Križ 2015, 2019), and if so, to what degree. I will suggest that the two phenomena are distinct, but related in both involving weak basic meanings that are strengthened in positive sentences through Exh.

This chapter is organized as follows. In section 5.2, I discuss at length the judgment for colour terms in positive sentences, then turn to describing plural and subatomic homogeneity effects and subatomic part structure. Section 5.3 then outlines the two theories of homogeneity we will be evaluating for their ability to deal with subatomic homogeneity: the hypothesis that homogeneity is the result of an excluded-middle presupposition, and the exhaustivity-based account of homogeneity proposed by Bar-Lev (2018, 2021). Then, in section 5.4, I discuss the matter of conjoined colour terms; colour conjunctions will play a pivotal role in making the cut between competing theories. In particular, I will show that predicate conjunctions with atomic subjects are always interpreted intersectively, something which will cause problems for some theories of homogeneity. I will also show that we should not understand the intersective conjunction of colour terms as involving any

pragmatic *weakening* of the normal semantic meaning of colour terms; hence, colour terms must be lexically weak. With this much background in our pocket, we turn in section 5.5 to evaluating the two theories of homogeneity. We will see that, regardless of their status for pluralities, they do not succeed in capturing the meaning of colour term predicates (and by extension subatomic homogeneity), in particular colour term conjunctions. Section 5.6 concludes.

5.2 An empirical look at summative predicates

This section provides an empirical starting point for comparing different theories of subatomic homogeneity, including the one I proposed in previous chapters. I did so without referring to the notion of homogeneity explicitly, and I collapsed predicates like colour terms that participate in the homogeneity paradigm with other predicates like *comedy*/*tragedy* that do not.

We begin in section 5.2.1 by looking narrowly at colour terms in positive sentences. I will defend in detail the claim from previous chapters that colours are interpreted as modifying all subatomic parts of their argument in positive sentences; that is, they are interpreted as universal/'total.' While I claimed in previous chapters that this is so despite their weak lexical entries, the goal for now is simply to focus on the intuited meaning in simple positive sentences—I will be describing them as strong (universal) without this contradicting my claims in previous chapters. Then, in section 5.2.2, we will link colour terms to a broader pattern of summative predication and homogeneity effects. Finally, section 5.2.3 offers some comments on part-structure within atoms.

5.2.1 Colour predicates receive a total interpretation

Let's begin with a famous observation about colour terms (see Kennedy & McNally 2010 and citations therein), apparently at odds with the rhetoric so far in this chapter. Colour terms are often used to refer only to some salient component of their subject. Hence, the two sentences in (9) could both be true of the same grapefruit, if its skin is yellow (9a) and its flesh is pink (9b).

(9) a. The grapefruit is yellow.
b. The grapefruit is pink.

In these examples, the colour predicates are only true of some part of the grapefruit, not all of it: a part of the grapefruit is yellow, a part of it is pink. As a result, it seems obvious that colour terms are interpreted as applying to only a part of their subject. To borrow terminology from Yoon (1996), colours appear to receive a 'partial' (existential) interpretation rather than a 'total' (universal) one. For example, (9a) appears to mean (10a)—not (10b), which would be false in our scenario.

(10) a. $\exists y[y \sqsubseteq \iota x[\mathsf{grapefruit}(x)] \wedge \mathsf{yellow}(y)]$.
b. $\forall y[y \sqsubseteq \iota x[\mathsf{grapefruit}(x)] \rightarrow \mathsf{yellow}(y)]$.

The apparently obvious conclusion is that colour terms are lexically existential:

(11) $[\![\text{yellow}]\!] = \lambda x.\exists y[y \sqsubseteq x \wedge \mathsf{yellow}(y)]$.

But in fact, (11) is not the right way to characterize the meaning of colour terms, even when they are apparently partial as in (9). Indeed, we will see that the truth conditions in (10a) fail

to capture that the intuited meaning is stronger than mere existential quantification by the colour term over parts of the subject. In what follows, I will argue that examples like (9) are more of a 'special case' where colours' meaning is pragmatically made weak due to the complex structure of the subject. That is, world knowledge informs us that grapefruits have different components (their skin, flesh, and so on), and this complication obfuscates colours' real meaning. We will work toward controlling for this pragmatic complication in colour terms' meaning, and as we do so, a different picture will emerge—one in which colour terms do not receive a *partial* interpretation, but actually a *total* one (as already noted by Harnish 1976, Levinson 1983, Krifka 1990, Lasersohn 1995, Winter 1998, and indeed previous chapters of this thesis).

The first reason to question the weak truth conditions in (10a) is the following: while not all parts of the grapefruit in (9) are of a certain colour, all parts of the component that the speaker has in mind are in fact of that colour. (9a) means the skin is entirely yellow, and (9b) means the flesh is entirely pink. It would be false to describe a grapefruit whose flesh was half pink and half orange with (9b). Of course, such multicoloured grapefruits are not common in the real world, so it may help to consider an example where there is no prior expectation that a particular part of the subject will be of only one colour. Take (12), for example. Many cars are painted with more than one colour, but in (12), *red* is interpreted in a universal way vis-à-vis the relevant component of the car. The relevant component of the car, of course, is the painted exterior—and not the tires, the steering wheel, or other components.

(12) The car is red.

That is, *red* in (12) means 'all red,' not 'partially red,' but it only quantifies over the relevant component of the car. Hence, truth conditions like in (10a), where colours only quantify existentially over their subject, are too weak. Once we control for the component of the subject that the speaker has in mind, colour terms are in fact interpreted as quantifying over it universally.

It is easy to think of ways that only the skin or flesh of the grapefruit in (9) could be relevant (Kennedy & McNally 2010): for example, the skin is what one actually sees of an uncut grapefruit, and the flesh is what distinguishes the different varieties (and it is the part that one actually eats). The way that the subject is reduced to a pragmatically relevant component can be appreciated by contrasting (9) with an equivalent sentence referring to a fruit whose varieties are *not* told apart by the colour of the flesh. Apples are one such fruit: their flesh is always the same kind of off-white. As a result, an apple with a red skin can be described by (13a), but (13b) is sharply degraded (outside of a special context, such as comparing a healthy apple with a rotten one).

(13) a. The apple is red.
b. #The apple is white.

This pragmatic 'reduction' of the subject to a particular salient subcomponent (i.e., the domain-restriction of the subject) is not unique to colour terms. There is nothing special about colour terms. Consider the meaning of *small* in (14). *Small* is clearly not a partial adjective: if it was, everything would be small by virtue of having some mereological part that is small. Now imagine that Aisha's house is normal-sized on the outside, but the walls are all several metres thick, so that there is only room to crawl around on the inside. (14) is a true description of such a house.

(14) Aisha's house is very small.

That is, (14) can be interpreted to mean that *the relevant part* of the house (its interior) is small for such a part of a house. Hence, this type of pragmatic reduction of the subject is not unique to colour terms.

If it is true that colour terms are total once we put aside the pragmatic reduction of subjects, we should be able to observe this in a more empirically straightforward way by controlling for the kind of individual we use as a subject. If the subject has no pragmatically relevant components to speak of, then my description leads to the expectation that colours should necessarily modify *all* subatomic parts of that subject. Colour terms' total meaning would be more easily appreciated, since we would have controlled for the pragmatic complications of multi-component subjects like grapefruits, cars, and houses.

The prediction is borne out. As we will see in a moment, colours predicated of things like walls, lines, or flags are all given a true universal interpretation (vis-à-vis the entire subject). These are subjects that, at least in the basic case, lack components. Flags, for example, are made up of only the colours that constitute them; they have no components other than the shapes created by the colours themselves.

Of course, particular colour designs on things like flags or walls may lead us to conceptualize of them as having a background and foreground, in which case flags are more like grapefruit in exhibiting multiple components. The flag of Japan is conceptualized as having a red circle on a white background, so that one may say (15), where *white* only means that the relevant component of the flag (its background) is entirely white.

(15) The flag of Japan is white, with a red circle on it.

The background–foreground distinction is yet another non-linguistic complication. (15) means that the flag's background is *entirely* white; it does not mean that the flag is only partially white.[2]

We need to use subjects lacking any components, including the background–foreground distinction. Tricolours, like the flags of France or Italy, have more than one colour, and yet these colours are conceptualized as on a par; it is not the case that one of the colours is the background on which the others are superimposed. Such flags are therefore truly free of subcomponents. Thus, if the claim that colours are given a total interpretation vis-à-vis the pragmatically relevant component of the subject is correct, it should be infelicitous to describe these flags by only one of their colours. This is indeed the case:

(16) #The flag of Italy is green.

(16) is clearly false. More generally, simple sentences like (17) mean that the subject is entirely of

[2]This is in line with the following description by Kennedy & McNally (2010):

> A t-shirt which is entirely white except for a few small red flowers around the neck or three thin pale-grey stripes will pass as white, but a shirt which is white on the entire left-hand side, front and back, and red on the other, will not. The less the color in question predominates on an object, the less likely we are to describe the object as being of that color and the more likely we are to use complex descriptions such as "white with red flowers." (Kennedy & McNally 2010:92)

The shirt that is "white with red flowers" is entirely white (that is, the background is entirely white), and there are red flowers imposed onto it. A half-white shirt does not count as being 'white' because it is not conceptualized as having an entirely white background.

the colour denoted by the colour term.

(17) The flag is green.

What (17) means is that the flag is just about entirely green:

(18) $\forall y[y \sqsubseteq \iota x[\mathsf{flag}(x)] \rightarrow \mathsf{green}(y)]$.

The same observation holds for other complication-free subjects like lines or walls (19); I will adhere to flags in my examples.

(19) a. The line is green. ($\approx$ the line is all green)
b. The wall is green. ($\approx$ the wall('s surface) is all green)

I described (18) as meaning that the flag is 'just about' all green because there is in fact another qualification to make about colours' totality: the concept of 'totality' should not be interpreted too rigidly. It has already been noted that total adjectives are not necessarily interpreted in a literally total way. To see this, consider other total adjectives, for example *clean* (Yoon 1996).

(20) The table is clean.

(20) means the entire table is clean, not only that the table has a clean part; it would be false if the table was half dirty, for example. But it has been observed that context can significantly affect the extent to which the totality is interpreted in a fine-grained manner. As Sassoon & Zevakhina (2012:227) write, "by default, statements like *The garage is clean* may be accepted as 'true enough' even when they are actually false, e.g., there are stains on the garage floor." If such 'coarse' readings of total adjectives like *clean* are possible, this should be the case with colours, too. The ability to tolerate some degree of exception has been called 'non-maximality' in the homogeneity literature, something we return to in section 5.4.4.

In addition to simple sentences like (17), we can also observe colours' totality by attempting to predicate more than one colour term of a single individual. If colour term predication was partial rather than total, we would expect to be able to do this. As an experimental set-up, (21) does this with colour terms which are ensured to be weak due to being modified by *partially*:

(21) The partially white flag is partially green.

If colour terms were lexically partial, it should be possible to remove both instances of *partially* from (21), as in (22).

(22) #The white flag is green.

As a matter of fact, as we have discussed throughout this thesis, (22) is contradictory. Intuitively, the contradiction comes from the fact that (22) means both that the flag is entirely white, and that it is entirely green. If colour terms were given a partial interpretation, (22) would simply mean that the flag has a white part and a green part. Thus, we have evidence from both consistent (17) and

inconsistent (22) sentences that colour terms are interpreted as total.[3]

In summary, colour terms do not receive a partial interpretation in the kinds of sentences looked at so far, even though this can be obscured by the pragmatic domain-restriction of the subject to one of its components (including the distinction between background and foreground). This is why all the examples with colour terms in the previous chapters of this thesis had flags as subjects, a practice that I will maintain in this chapter too. Hence, the contrasts reported in this thesis must point to the meaning of colour terms themselves, properly isolated and controlled for, rather than complications stemming from world knowledge about the salient components of subjects. What is more, as we have seen in previous chapters and will continue to observe, colour terms show interesting behaviour in terms of their totality or partiality even when we control for their subjects being flags.

5.2.2 Subatomic homogeneity

Given the universal interpretation accorded to colour terms in positive sentences, we would expect negative colour predication to negate the universal meaning. Negative sentences with colour predication would mean that the subject is not entirely of that colour. Yet, this is not the case (Löbner 2000; Spector 2013; Križ 2015, 2019). Rather, negative sentences with colour terms mean that *no* part of the subject is of that colour:

(23) The flag is not green.
≈ the flag is not green at all

Thus, we have an effect where a predicate is true of all subatomic parts of its subject in positive sentences, but of no subatomic parts in the negative.

Löbner (2000) points out that this all-or-nothing effect is found generally with what he calls 'summative' predicates. Summative predicates are predicates that are true of an individual by virtue of being true of that individual's material parts (24a). This is opposed to 'integrative' predicates (24b) which do not refer to their argument's material parts.

(24) a. SUMMATIVE PREDICATES ((i) examples from Löbner 2000)
(i) The cow is **black**.
(ii) The desk is **metal**.
b. INTEGRATIVE PREDICATES
(i) The cow is **mad**.
(ii) Aisha is a **professor**.

That is, a desk is metal by virtue of its material parts being metal, but Aisha is not a professor by virtue of her material parts being professors. While I will continue focusing on colour terms, summative predicates show the same effect in general, as with the material term in (24a-ii):

(25) a. The desk is metal.

[3] In fact, as noted in chapters 2 and 4, (22) also has a non-contradictory reading. Imagine you put a white flag in the laundry and it comes out green; the flag that was formally entirely white is now entirely green. In such a context, you could utter (22). But even on this non-contradictory reading of (22), both colour terms are given a total interpretation: the flag was previously all white, now it is all green.

≈ all parts of the desk are metal

b. The desk is not metal.
≈ no part of the desk is metal

Call the paradigm exemplified by (23) and (25) 'subatomic homogeneity.'

Something very similar holds for pluralities. Predicates applied to pluralities hold of all of the plurality's atomic parts in positive sentences, and of none of its atomic parts in negative sentences. (26) is repeated from (7).

(26) a. Aisha saw the children.
≈ she saw all of the children
$\not\approx$ she saw at least some of the children

b. Aisha didn't see the children.
$\not\approx$ she didn't see all of the children
≈ she saw none of the children

Call this paradigm 'plural homogeneity.'[4] I will refer to as 'homogeneity' the meaning of both positive and negative sentences (i.e., the universal quantification in positive sentences and the meaning of negative sentences as negated existentials) either as a joint paradigm or taken in isolation.

As Löbner (2000) points out, the plural and subatomic homogeneity effects can be observed together when a summative predicate has a plural subject:

(27) a. The flags are green.
≈ all the flags are entirely green
$\not\approx$ at least some of the flags are entirely green
$\not\approx$ all of the flags are at least partly green
$\not\approx$ at least some of the flags are at least partly green

b. The flags are not green.
≈ none of the flags are green at all
$\not\approx$ not all of the flags are not green at all
$\not\approx$ none of the flags are entirely green
$\not\approx$ not all of the flags are entirely green

The effect in (27) could in fact also be described without reference to atoms at all: (27a) means that all subatomic parts of the plurality denoted by *the flags* are green, and (27b) means that no subatomic parts of the plurality are green. From this example alone, it is not immediately clear that we need to worry about 'stacking' plural and subatomic homogeneity effects. Marcin Wągiel (p.c.)

[4] Križ (2015) points out that plural homogeneity is also found outside of distributive predication, as with the collective predicate *perform Hamlet*:

(i) a. The children are performing *Hamlet*. (Križ 2015:7)
≈ all the children are performing *Hamlet* together
$\not\approx$ some of the children are performing *Hamlet* together

b. The children are not performing *Hamlet*.
$\not\approx$ not all of the children are engaged in performing *Hamlet*
≈ none of the children are engaged in performing *Hamlet*

I will not discuss this in this chapter.

points out that in examples involving lexically existential ('partial') predicates like *dirty*, however, we do observe plural homogeneity:

(28) a. The table is dirty.
≈ at least part of the table is dirty
b. The tables are dirty.
≈ all of the tables are at least partly dirty

(28) suggests that reference to atoms is necessary even with predicates referring to subatomic parts. I put this issue aside for this chapter.

5.2.3 Part structure within atoms

So far, we have observed a subatomic homogeneity paradigm, a phenomenon which has to do with the parts of atoms. But what exactly is the nature of these parts? In this section, I first point out that we are talking about the 'arbitrary pieces' rather than the 'salient parts' of atoms, then make explicit some assumptions about quantification in the meta-language I have been using so far.

Subatomic homogeneity refers to arbitrary pieces

In one respect, discussing plural homogeneity is a little easier than subatomic homogeneity: for the plural paradigm, we have some baseline for what the parts being quantified over are. Plurals are made up at least of atoms, so that (29) means that for each atomic child, Aisha saw/didn't see that child (again, some exceptions are possible; see section 5.4.4).

(29) Aisha {saw, didn't see} the children.

Work on plural homogeneity does not usually dwell at length on part structure, precisely because the parts are given: they are the atoms. At most, some work (e.g., Bar-Lev 2018, 2021 and Križ & Spector 2021) asks whether 'subpluralities' are also manipulated as parts. That is, the parts for plurals are at least atoms, and maybe also subpluralities made up of the atoms.

But what, for the subatomic homogeneity paradigm, is the nature of the parts of atoms? Work on part structure (see Wągiel 2021 and references therein) sometimes distinguishes between the '(salient) parts' and the '(arbitrary) pieces' of atoms. Wągiel (2021:5) gives as an example the parts of a table: the leg of a table is a salient part, which we conceive of as having some individuality, while some splinter or a randomly selected cubic centimeter of the table is just an arbitrary piece. In this chapter, I will keep using the word 'part' neutrally, and refer specifically to 'salient parts' or '(arbitrary) pieces' when I need to make the distinction between these notions. In this section, I ask whether subatomic homogeneity makes reference to salient parts or to arbitrary pieces.

In section 5.2.1, I noted that the sorts of objects we want to use as the arguments of colour terms are things like flags, lines, or walls, which do not have any salient parts other than the parts made up by the colours themselves. This is to observe that colour terms really are total in positive sentences. But we wouldn't want to build a theory around only these types of examples, precisely because they are cases where salient parts are by definition made up of only one colour. To see this, consider (30).

(30) The flag is green.

We cannot know from (30) alone whether the universal meaning of *green* is best understood as quantification over salient parts or arbitrary pieces. Truth-conditionally, in (30), these are equivalent. Paraphrasing (30) as 'all arbitrary pieces of the flag are green' straightforwardly captures that there are no other colours on the flag. But paraphrasing (30) as 'all salient parts of the flag are green' captures the same truth conditions: salient parts of flags are by definition made up of only one colour, so if all salient parts of the flag are (at least partly) green, then all salient parts are entirely green throughout, and all arbitrary pieces of the flag are themselves entirely green.

Thus, for objects like flags, the three attempts in (31) of giving truth conditions for (30) are equivalent. In (31), $\sqsubseteq_{\text{SP}}$ refers to the salient-part parthood relation and $\sqsubseteq_{\text{AP}}$ refers to the arbitrary-piece parthood relation.

(31) a. $\forall x[x \sqsubseteq_{\text{SP}} \iota y[\mathsf{flag}(y)] \rightarrow \forall z[z \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{green}(z)]]$.
b. $\forall x[x \sqsubseteq_{\text{SP}} \iota y[\mathsf{flag}(y)] \rightarrow \exists z[z \sqsubseteq_{\text{AP}} x \wedge \mathsf{green}(z)]]$.
c. $\forall x[x \sqsubseteq_{\text{AP}} \iota y[\mathsf{flag}(y)] \rightarrow \mathsf{green}(x)]$.

All of these mean in one way or another that all arbitrary pieces of the flag are green.

As such, for the purposes of this section, we must turn away momentarily from objects like flags and consider objects that can have salient parts of more than one colour. For example, imagine a table that has been painted with stripes of every basic colour. It seems that in such a case, the conceptually salient parts of the table (outside of special contexts made specifically to bring out colours) are is still the tabletop and the legs. Thus, a table, unlike a flag, can have salient parts that have more than one colour.

Focusing on objects like tables instead of flags, then, we observe that it is just about unnecessary to bring in the notion of salient parts to describe the subatomic homogeneity paradigm, although getting to this observation will not be without complications. To start things off, consider (32):

(32) The table is green.

There are two things to point out about (32). On the one hand, it does *not* mean that literally all arbitrary pieces of the table are entirely green. (32) would be true of a table that is made of wood and painted green on the exterior. In this sense, the table has arbitrary pieces which are not green, viz. its non-visible wooden interior. On the other hand, (32) *does* mean that all *visible* pieces of the table are green. For a wooden table that has been painted, (32) must mean that the table has been painted green all over.[5]

The picture from (32) is complex, but best understood as involving quantification over arbitrary pieces rather than salient parts. The most obvious point to make is that the universal quantification of *green* cannot be *only* over salient parts. If (32) was paraphrasable as 'all salient parts of the table are at least partly green,' (32) would be true as long as each salient part of the table had *some* green piece. (32) would be true if the whole table was painted in green and white stripes, for instance. In fact, putting aside the possibility of small exceptions, the sentence means that the table is just about entirely green. As such, if we wanted to talk about salient parts, we would have to describe the truth conditions of (32) rather redundantly as 'all the salient parts of the table are made up

[5] A small caveat is that a part of the table like the *underside* of its tabletop wouldn't need to have been painted green for (32) to be true. But this can be described as the underside not being 'visible' for all pragmatic purposes. Thus, this is aligned with the fact that the non-visible wooden interior of the table need not be green.

entirely of green pieces,' as in (33a).

(33) ⟦The table is green⟧
  a. $= 1$ iff $\forall x[x \sqsubseteq_{\text{SP}} \iota y[\mathsf{table}(y)] \rightarrow \forall z[z \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{green}(z)]]$.
  b. $\neq 1$ iff $\forall x[x \sqsubseteq_{\text{SP}} \iota y[\mathsf{table}(y)] \rightarrow \exists z[z \sqsubseteq_{\text{AP}} x \wedge \mathsf{green}(z)]]$.

Prima facie, the double universal quantification in (33a) is hardly necessary. We could have written out the truth conditions without reference to salient parts at all:

(34) $\forall x[x \sqsubseteq_{\text{AP}} \iota y[\mathsf{table}(y)] \rightarrow \mathsf{green}(x)]$.

(34) is obviously simpler and, if equivalent to (33a), preferable.

In fact, (33a) and (34) are different in one way, depending on whether we claim that objects must be partitioned into salient parts, or that objects can have pieces which are not part of a salient part at all (e.g., the interior pieces of a table are not salient). For the sake of argument, assume the latter. (33a) immediately captures the fact that not literally all arbitrary pieces of the table must be green for the sentence to be true: it is only the case that all arbitrary pieces *that are part of a salient part* of the table must be green. Pieces which are not part of a salient part do not need to be green. Proponents of the truth conditions in (33a) could reasonably argue that non-visible parts of the table are not salient parts at all: the salient parts are the visible exterior of the tabletop and the visible exterior of the legs of the table, and non-visible parts (the wooden interior of the table and the underside of the tabletop) do not qualify as salient parts at all. Thus, (33a) only requires all arbitrary pieces of the visible exterior of the tabletop and legs of the table to be green, without requiring anything of the interior of the table. (34), on the other hand, does not immediately capture this, because it requires all arbitrary pieces of the table to be green.

This is not an insurmountable problem for (34): one could rely on an independent domain-restriction mechanism to reconcile (34) with the fact that non-visible pieces of the table do not need to be green. If we find that such domain-restriction is independently required to discuss the subatomic homogeneity paradigm, (34) would in fact become weakly preferable to (33a) merely by virtue of being simpler.

In fact, this seems to be the case. Consider either of the following discourses:

(35) a. The car is blue, with black seats and a white roof.
  b. The car is blue. It has black seats and a white roof.

The discourses in (35) mean that all of the exterior of the body of the car, except the roof, is painted blue. Crucially, we need a domain-restriction mechanism regardless of whether we take *blue* to mean that all salient parts are blue (36a), or all arbitrary pieces are blue (36b).

(36) ⟦The car is blue⟧ $= 1$ iff
  a. $\forall x[x \sqsubseteq_{\text{SP}} \iota y[\mathsf{car}(y)] \rightarrow \forall z[z \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{blue}(z)]]$.
  b. $\forall x[x \sqsubseteq_{\text{AP}} \iota y[\mathsf{car}(y)] \rightarrow \mathsf{blue}(x)]$.

Indeed, the speaker must find the seats and the roof to be salient parts of the car, if they are worth mentioning as separate from the rest. Thus, these parts must be salient but excluded from the domain in *the car is blue*. Regardless of whether *blue* means that all salient parts of the car are made up of blue arbitrary pieces, or that all arbitrary pieces are blue, we need a domain-restriction

mechanism in (36) to ensure that the roof and the seats are not intuited as blue at all.

If we need a domain-restriction mechanism either way, then it is simpler to take the subatomic homogeneity paradigm to only make reference to arbitrary pieces (modulo domain-restriction), rather than both arbitrary pieces and salient parts. In conclusion, a sentence like (37) has the truth conditions in (37a), not (37b-i) (or (37b-ii), which we already ruled out in (33)).

(37) ⟦The table is green⟧

a. $= 1$ iff $\forall x[x \sqsubseteq_{\text{AP}} \iota y[\mathsf{table}(y)] \rightarrow \mathsf{green}(x)]$.

b. (i) $\neq 1$ iff $\forall x[x \sqsubseteq_{\text{SP}} \iota y[\mathsf{table}(y)] \rightarrow \forall z[z \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{green}(z)]]$.

(ii) $\neq 1$ iff $\forall x[x \sqsubseteq_{\text{SP}} \iota y[\mathsf{table}(y)] \rightarrow \exists z[z \sqsubseteq_{\text{AP}} x \wedge \mathsf{green}(z)]]$.

In other words, given the redundancy of adding in the notion of salient parts, I will simply assume that, for the purposes of describing the homogeneity paradigm, subatomic parthood refers to arbitrary pieces.[6] This is a weak conclusion, because the preference for (37a) over (37b-i) is only due to a notion of parsimony. As such, I will occasionally bring up salient parts when this notion becomes relevant. When I need to do so, I will move on in my examples from discussing flags to discussing tables, because tables have salient parts that do not necessarily correspond to their colour.

The quantificational force of meta-language constants

There is one more point to make about quantification. I have been using meta-language constants like green without being explicit about their meaning. To be in the set of green things, does an individual have to be entirely green, or only partly green? Consider for concreteness the lexical entry I gave for colour terms in chapter 2:

(38) $\lambda x.\exists y[y \sqsubseteq x \wedge \mathsf{green}(y)]$.

Is this best paraphrasable as 'some pieces are *all* green' or as 'some pieces have *some* green on them'? That is: what is the 'covert' quantificational force of green, below the 'overt' $\exists$ quantifier in the lexical entry?

In fact, once the subatomic parthood relation $\sqsubseteq$ is taken to refer to arbitrary pieces, it does not matter whether the constant green is covertly universal or existential. If green in (38) means 'all green,' (38) means that there is an arbitrary piece that is entirely green. If green in (38) means 'partly green,' (38) means that there is an arbitrary piece that has some green on it. Either way, the meaning of (38) is existential. What is more, crucially, the same goes for a putative universal lexical meaning for colour terms:

(39) $\lambda x.\forall y[y \sqsubseteq x \rightarrow \mathsf{green}(y)]$.

If green in (39) means 'all green,' the intended universal meaning of (39) is captured straightforwardly. But, counter-intuitively, (39) is still universal even if green means 'partly green.' Indeed, in the world of arbitrary pieces (unlike salient parts), any piece can be further subdivided into smaller arbitrary pieces. If (39) means that all pieces can be further subdivided into other pieces such

[6] Of course, this is not meant as a general claim on subatomic part-structure (see Wągiel 2021 for more general discussion), but just as a claim about subatomic homogeneity.

that at least one of these subpieces is itself (entirely or partly) green, then (39) does not permit its subject to have any non-green piece—even if green is existential. If x had a piece with a non-green 'subpiece' y, y would *not* have a green 'subsubpiece' z. Yet, (39) claims that *all* arbitrary pieces have a green subpiece, and z is an arbitrary subpiece. Thus, if x is not entirely green, it would not count as green according to (39), even if the constant green is only existential.

For simplicity, I will describe green as meaning 'all green.' The intuition is simply that, if green denotes the set of exemplars of the concept GREEN, it seems more palatable for something to be an exemplar of this concept if it is entirely green (cf. e.g. Gärdenfors (2000)). Nothing hinges on this.

5.2.4 Interim summary

In this section, I identified the intuited meaning of colour terms as universal in positive sentences. I then showed that they are negated existentials in negative sentences. This means that, like other 'summative' predicates, colour terms participate in a subatomic homogeneity paradigm. This paradigm co-exists with a plural homogeneity effect. Plural homogeneity is defined vis-à-vis the atomic parts of pluralities, while subatomic homogeneity is defined vis-à-vis the arbitrary pieces (rather than the salient parts) of atoms.

5.3 Some theories of homogeneity

In this chapter, I will defend my controlled-exhaustivity theory of subatomic homogeneity against two alternatives (of course, there are more alternative theories of homogeneity). The first is the classic suggestion that summative predicates come with an 'all-or-nothing'/'excluded-middle' presupposition. They are only defined for individuals of which the predicate is true of all parts or no parts at all. The second is much closer to mine, in that it also relies on an Exh operator appearing in positive sentences to strengthen weak lexical meaning (Bar-Lev 2018, 2021). In this section, I will take each theory in turn, to both describe how it handles plural homogeneity, and how it could be made to work for subatomic homogeneity.[7] In this section, the theories will come out as working perfectly well for subatomic homogeneity, but in later sections, I will bring in new data with conjoined summative predicates which will create problems for them.[8]

[7] Subatomic homogeneity is less often discussed in the homogeneity literature than plural homogeneity. To my knowledge, Löbner 2000 is the only article to focus on it at least as much as plural homogeneity; but my discussion of Löbner's proposal will in fact focus on a formalization by Gajewski (2005) which is, in fact, made for plurals.

[8] A third alternative theory of homogeneity is the view that homogeneity arises due to semantic underspecification interacting with the Strongest Meaning Hypothesis of Dalrymple et al. (1994). Indeed, Krifka (1996) (see also Križ & Spector (2021) and citations therein) proposes that the grammar does not fix whether a predicate is interpreted universally or existentially when it takes a plurality as an argument. Rather, the Strongest Meaning Hypothesis of Dalrymple et al. (1994) ensures that speakers prefer whichever interpretation results in the strongest meaning. For positive sentences, a universal meaning is stronger than an existential one; for negative sentences, a negated existential is stronger than a negated universal.

Transferring this hypothesis to subatomic homogeneity initially seems straightforward: we would claim that colour terms are lexically ambiguous, and speakers choose to interpret them as strongly as they can within a given sentence. However, this theory quickly runs into the problem of sentence-internal contradictions like (i), discussed at length in chapter 2:

I start by overviewing my controlled-exhaustivity theory of subatomic homogeneity, showing in particular that it does not carry over to plural homogeneity.

5.3.1 The exclusion theory advanced in this thesis

In previous chapters, I made suggestions about the meaning of predicates that can capture the subatomic homogeneity effect, although I did not attempt to make any links to the plural homogeneity effect. In this section, I first overview the proposal for subatomic homogeneity, then show that it straightforwardly cannot be extended to plural homogeneity.

Subatomic homogeneity as the exclusion of related predicates

In chapter 2, I argued that the meaning of colour terms should be understood as involving an Exh operator (see chapter 1) whose domain is composed of the set of colours. I argued, due to the observed meaning of sentences with additive particles and conjunction (more on conjunction below), that colour terms are lexically existential (Harnish 1976; Levinson 1983):[9]

(40) $[\![\text{green}]\!] = \lambda x.\exists y[y \sqsubseteq x \wedge \mathsf{green}(y)]$.

This obtains the meaning of negative sentences immediately:

(41) The flag is not green.
$\approx$ the flag is not green at all

I then argued that colour terms, like many other predicates (including non-summative predicates like *comedy*/*tragedy* or *train*/*plane*) come with an obligatory Exh operator in positive sentences that makes them incompatible with one another:[10]

(42) $[\![\text{Exh}_{\text{ALT}}\ [\text{the flag is green}]]\!]$
$= 1$ iff the flag is green $\wedge$ the flag is not white $\wedge$ the flag is not red $\wedge \ldots$

The meaning of (42) is that the flag has at least one green piece, and does not have a piece that is of any other colour. Given world knowledge that all areas of a surface must have a colour, this is pragmatically strengthened to mean that the flag is entirely green, in the same way that (43) with an overt *only* means that the flag is entirely green.

(43) The flag is only green.

(i) #The white flag is green.

Speakers would choose to interpret the colour terms existentially here, because interpreting them universally is not consistent.

Let me note that this theory is relegated to a footnote only for presentational reasons: my argumentation against it is fundamentally different from the issue I will take with the other alternative theories, which I will show undergenerate in predicting inconsistency in some examples, rather than overgenerating in predicting (i) to be consistent.

[9] Colour adjectives have quite rightly been described as degree predicates. For instance, they can take comparative morphemes—e.g., one tree can be 'greener' than its neighbour. I am not trying to deny this; the existential meaning in (40) could very well arise compositionally from the degree semantics of partial predicates (see Yoon 1996 and Kennedy & McNally 2010).

[10] See chapters 2 and 6 on the syntax of Exh in (42).

No extension of this account to plural homogeneity

Much of this chapter will be spent showing at length that various theories of plural homogeneity cannot be extended to subatomic homogeniety. Seeing that my account of subatomic homogeneity cannot be extended to plural homogeneity is much simpler.

My account crucially relies on world knowledge, so that 'partly green, and no other colour' is strengthened to meaning 'entirely green.' All parts must have a colour, and no part is non-green, so they must all be green. This clearly cannot carry over to plural homogeneity; world knowledge does not dictate anything about the parts of pluralities. Indeed, attempting to capture plural homogeneity through the exclusion of related predicates would look as in (44) (where I am not F-marking *singing* because it is not intonationally (contrastively) focused; see chapter 2 on focus intonation for predicates that control Exh). In (44), I am marking *sing* as 'existential' to create a parallel with my account with colour terms. We can put aside how this arises compositionally (see Bar-Lev's proposal in section 5.3.3); what matters is simply that the non-exhaustified meaning of *the children are singing* is that at least one of them is singing, in the same way that the non-exhaustified meaning of *the flag is green* is that at least one part of the flag is green.

(44) $[\![\text{Exh}_{\text{ALT}}\ [\text{the children are singing}_{\exists}]]\!] = 1$ iff $\begin{cases} \text{there is a child who is singing} \wedge \\ \text{there is no child who is dancing} \wedge \\ \text{there is no child who is talking} \wedge \\ \quad \text{there is no child who is} \ldots \end{cases}$

There are two problems here. First, unlike the subatomic parts of surfaces which *must all* have a colour, it is in fact possible for individuals to do nothing. World knowledge does not make (44) mean that all children must be singing. The assertion is that at least some children are singing and no children are doing anything else—but as far as world knowledge is concerned, there could be children who are *neither* singing *nor* doing anything else, because they could simply do nothing.[11] Second, the entailments about children not doing other things are not in fact intuited. Without contrastive focus on *singing*, (44) is not actually intuited as meaning that the children are not doing other things. They could be both singing and dancing, for example.

As such, the account of 'universal-looking' quantification in subatomic homogeneity that I developed in chapter 2 does not extend to plural homogeneity. It is a theory meant to collapse data like (45a–b), but it does so to the exclusion of explaining data like (45c).

(45) a. This comedy is #(also) a tragedy.
b. The white flag is #(also) green.
c. The children are singing.

There is currently no theory in the literature uniting all these types of examples. In chapter 6, I will claim that what unites these data is that they are all exhaustification effects displaying locality constraints on Exh.

We now turn to considering two theories of homogeneity which both differ from mine in deriving universal quantification in positive sentences, rather than existential quantification together with the exclusion of related predicates.

[11] Exh would clearly not exclude predicates like *breathe* or *exist* (perhaps due to these not being related to *sing* in the right way), in which case (44) would mean that even the children who are singing are neither breathing nor existing.

5.3.2 The first alternative theory: an excluded-middle presupposition

A classic way to derive the homogeneity effect is to postulate an 'excluded-middle presupposition' (EMP), which ensures that predication only results in the assignment of a truth-value if the predicate holds of all or no parts of its argument (Löbner 2000; cf. Schwarzschild 1994, Gajewski 2005). In what follows, I follow Gajewski's (2005) formalization of this approach for pluralities, and suggest two different ways to carry it over to subatomic homogeneity.

The EMP with pluralities

To deal with homogeneity effects, Löbner (2000) hypothesizes that predication introduces a presupposition that the predicate holds of all or no parts of its argument. This is the case for both atomic parts of pluralities and subatomic parts of atoms. The EMP results in truth-value gaps for non-homogeneous individuals or pluralities; predicates are neither true nor false of such arguments.

An influential formalization of this idea comes from Gajewski (2005); his proposal is tailor-made for plural homogeneity, but can be modified to deal with subatomic homogeneity as well (as shown below). Gajewski's proposal is that the EMP is introduced by an obligatory distributivity operator. DIST is defined in (46), where $\sqsubseteq_{\text{AT}}$ refers to atomic parthood; (47) provides a sample LF.

(46) $[\![\text{DIST}]\!] = \lambda P.\ \lambda x : \forall y[y \sqsubseteq_{\text{AT}} x \rightarrow P(y)] \vee \forall y[y \sqsubseteq_{\text{AT}} x \rightarrow \neg P(y)].\ \forall y[y \sqsubseteq_{\text{AT}} x \rightarrow P(y)]$.

(47) [The children] [DIST [sang]].

In the positive case, DIST asserts that the predicate holds of all atoms in the plurality. This is consistent with the first presuppositional disjunct.[12] Given the assertion, the only noticeable effect of the presupposition is that, if only some of the children sang, the sentence would be undefined rather than false.

(48) The children sang.
   a. presupposition: either all of the children sang or none of the children sang
   b. assertion: all the children sang[13]
   → all of the children sang

In the negative case, even if the assertion is only that not all the children sang, the presupposition projects past negation and effectively strengthens this to mean that none of them sang. The assertion is incompatible with the first presuppositional disjunct, so the only remaining possibility is for the second disjunct to hold—that is, for none of the children to have sung.

(49) The children didn't sing.
   a. presupposition: either all of the children sang or none of the children sang
   b. assertion: not all the children sang
   → none of the children sang

Given the all-or-nothing presupposition, it does not matter that the output condition of DIST

[12] I am using the term 'presupposition' to refer to the domain condition on the lambda-expression, and 'assertion' to refer to the output condition.

[13] Nothing would change if the assertion was only that some of the children sang.

is universal rather than existential. Having an existential output condition would yield identical definedness and truth conditions, because DIST is only defined for predicates that hold of all or none of their plural argument's atomic parts.

The EMP for subatomic homogeneity

For Löbner (2000), subatomic homogeneity results from the same EMP as with pluralities. There are two ways to do this: either the presupposition is present in the lexical entry of summative predicates, or language comes with a subatomic $\text{DIST}_{\text{SUBAT}}$ operator in addition to the atomic DIST operator in (46).

The first option would mean that summative predicates like colour terms are lexically only defined for arguments of which they are true of all subatomic parts or no parts at all:

(50) $[\![\text{green}]\!] =$
$\lambda x : \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{green}(y)] \vee \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \neg\mathsf{green}(y)].\ \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{green}(y)].$

Postulating that the EMP is part of *green*'s lexical meaning is motivated by the fact that the presence of this presupposition is regulated lexically, in light of the summative–integrative distinction. In (50), I have specified the parthood relation $\sqsubseteq$ as subatomic; nothing hinges on this. If x is a plurality, this comes out to meaning that all the subatomic parts of all the atoms are green or not green at all, so an 'underspecified' $\sqsubseteq$ relation would have worked just the same.

The lexical entry in (50) straightforwardly captures that *green* is universal in the positive:

(51) The flag is green.
$\rightsquigarrow$ the flag is entirely green

In the negative (52), the assertion and presupposition combine to yield the strong meaning that the flag is not green at all:

(52) The flag is not green.
a. presupposition: the flag is either all green or not green at all
b. assertion: the flag is not all green
$\rightarrow$ the flag is not green at all

Thus, as far as the basic paradigm is concerned, stipulating an EMP as part of summative predicates' lexical entries can in principle explain the subatomic homogeneity effect.

The second option is to claim that there is a subatomic $\text{DIST}_{\text{SUBAT}}$ operator, identical to Gajewski's (46) but with reference to subatomic parts:

(53) $[\![\text{DIST}_{\text{SUBAT}}]\!] =$
$\lambda P.\ \lambda x : \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow P(y)] \vee \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \neg P(y)].\ \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow P(y)].$

Moving the EMP from colour terms' lexical meaning to an operator creates a near-identical result for simple sentences (but as we will see in section 5.5.2, it can make a difference in more complex cases). It will not matter much on this view whether colour terms are lexically existential or universal (I will abbreviate these possible meanings as $\mathsf{green}_{\exists}$ and $\mathsf{green}_{\forall}$).

(54) $[\![\text{green}]\!] =$

a. $\lambda x.\exists y[y \sqsubseteq x \wedge \mathsf{green}(y)]$.
b. $\lambda x.\forall y[y \sqsubseteq x \rightarrow \mathsf{green}(y)]$.

Assuming that $\sqsubseteq$ refers to the arbitrary pieces of atoms (section 5.2.3), we get the following meanings with $\mathsf{green}_\forall$ and $\mathsf{green}_\exists$ respectively.

(55) $[\![\text{DIST}_{\text{SUBAT}}\ \text{green}]\!]$
$= \lambda x : \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{green}_\forall(y)] \vee \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \neg\mathsf{green}_\forall(y)].\ \mathsf{green}_\forall(x)$.

(56) $[\![\text{DIST}_{\text{SUBAT}}\ \text{green}]\!]$
$= \lambda x : \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{green}_\exists(y)] \vee \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \neg\mathsf{green}_\exists(y)].\ \mathsf{green}_\exists(x)$.

These are both good. First consider (55). The positive disjunct in the input condition means that for any piece you choose, that piece is entirely green. On the surface, the second disjunct might seem too weak for negative sentences, since it means that, for any piece you choose, it is not entirely green. But consider what would happen if that piece had any green on it at all—let's say it was half white, half green. In this case, there would be a piece of that piece, namely the green half, which you could have chosen. And that piece *is* entirely green. Therefore, the negative disjunct in the input condition in (55) does in fact lead to negative sentences meaning that the subject is not green at all.

As for (56), the positive disjunct means that all pieces are all green: if there was a piece which was not entirely green, the non-green piece of that piece would itself lack a green piece, contrary to the meaning of this disjunct. As for the negative disjunct, this straightforwardly means that there are no pieces with any green on them.

Thus, whether *green* is universal or existential, as long as we take the $\sqsubseteq$ relation to refer to arbitrary pieces and not to salient parts, we can obtain an all-or-nothing presupposition with $\text{DIST}_{\text{SUBAT}}$.[14] Note that, in the negative, this holds regardless of the relative scope of $\text{DIST}_{\text{SUBAT}}$ and *not* (57)–(58). In (57)–(58), I simply assume existential colour terms for simplicity of presentation. (57) is only defined if the flag is all green or not green at all, and asserts that it is not green.

(57) $[_{v\text{P}}$ not $[_{v\text{P}}$ the flag is $[\text{DIST}_{\text{SUBAT}}$ green$]]]$.

(58), with *not* below $\text{DIST}_{\text{SUBAT}}$, ends up with the same results:

(58) $[_{\text{TP}}$ [the flag]$_1$ $[_{v\text{P-2}}$ $\text{DIST}_{\text{SUBAT}}$ $[_{v\text{P-1}}$ λ_1 $[_{v\text{P}}$ not $[_{v\text{P}}$ t_1 is green$]]]]]$.
a. $[\![v\text{P-1}]\!] = \lambda x.\ \neg\mathsf{green}_\exists(x)$.
b. $[\![v\text{P-2}]\!] = \lambda x : \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \neg\mathsf{green}_\exists(y)] \vee \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{green}_\exists(y)]$.
$\forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \neg\mathsf{green}_\exists(y)]$.

Either way, assuming that $\sqsubseteq$ refers to arbitrary pieces, colour terms with $\text{DIST}_{\text{SUBAT}}$ in negative sentences are only defined if the flag is all green or not green at all, and the asserted meaning is that the flag does not have a green piece. Hence, the sentence is only true if the flag is not green at all.

We have just seen two ways to add an EMP to summative predicates: through their lexical entries, or through a subatomic $\text{DIST}_{\text{SUBAT}}$ operator. Is there a third option, according to which there is a single underspecified DIST operator responsible for both plural and subatomic homogeneity?

[14] I leave out for simplicity of presentation what would happen if $\sqsubseteq$ referred to salient parts.

(59) $[\![\text{DIST}]\!] = \lambda P.\, \lambda x : \forall y[y \sqsubseteq x \rightarrow P(y)] \vee \forall y[y \sqsubseteq x \rightarrow \neg P(y)].\, \forall y[y \sqsubseteq x \rightarrow P(y)]$.

Such a view would come with a difficulty for the distributive predication of integrative predicates:

(60) The women are (not) professors.

(60) involves a plural homogeneity effect; depending on the presence of negation, the sentence means that all or none of the women are professors. On the EMP account of homogeneity, we want a presuppositional DIST operator to capture this. On the other hand, there is no subatomic homogeneity effect in (60) given that *professor* is an integrative predicate. We would *not* want to end up with the meaning that all or none of the subatomic part of the individual women are themselves professors. Thus, it is better to keep separate the EMP intended for the atomic parts of pluralities and the EMP intended for subatomic parts of atoms. The latter is only found with summative predicates, whereas the former is found with both summative and integrative predicates when the subject is a plurality.

The EMP account of homogeneity has come under various types of criticism in the literature (Spector 2013; Križ 2015), both in terms of the core proposal that a presupposition is at work in creating the homogeneity paradigm, and in terms of Gajewski's linking of this presupposition to distributivity. On the latter point, Križ (2015) shows that homogeneity is also observable with non-distributive plural predication, as mentioned in footnote 4. For the sake of argument, I put these criticisms aside, and will turn instead to a new kind of argument against the EMP in section 5.5.2, focusing specifically on subatomic homogeneity. Before doing so, however, I now lay out a second approach to homogeneity.

5.3.3 The second alternative theory: Innocent Inclusion of parts

The second account of homogeneity we consider is based on exhaustivity, giving it more common ground with my controlled-Exh proposal. I will focus on Bar-Lev's (2018; 2021) account (but see also Magri 2014). The basic premise is that the lexical meaning of plurals is existential, immediately capturing the meaning of negative sentences. In positive sentences, Exh strengthens the existential to a universal. Bar-Lev uses the notion of Innocent Inclusion (Bar-Lev & Fox 2017; Bar-Lev 2018) to have Exh assert the truth of ('include') the domain alternatives of the existential plural. I will focus exclusively on the theory Bar-Lev builds for distributive plural homogeneity, which is all that is needed to try to carry over his account to subatomic homogeneity.

Bar-Lev's theory for plural homogeneity

In Bar-Lev's theory, the meaning of plurals is existential: the plain meaning of *the children laughed* is that at least one laughed. In what follows, assume there are two children, Aisha and Ben.

(61) $[\![\text{The kids laughed}]\!] = 1$ iff $\mathsf{laughed}(a) \vee \mathsf{laughed}(b)$

This existential plain meaning comes about from an existential plural operator, ∃-PL:

(62) a. $[\![\exists\text{-PL}]\!] = \lambda D_{\langle et\rangle}.\lambda P_{\langle e,st\rangle}.\lambda x_e.\exists y \in D \cap \mathit{Part}_{\text{AT}}(x)[P(y) = 1]$.
b. $\mathit{Part}_{\text{AT}}(x) = \{y : y \sqsubseteq_{\text{AT}} x\}$ (Bar-Lev 2021:1062)

(61) therefore has the LF in (63), where the domain D is presented as a subscript on $\exists$-PL.

(63) [The kids] [$\exists$-PL$_D$ laughed]. (Bar-Lev 2021:1062)

Assuming that $[\![\text{the kids}]\!] = a \oplus b$ and $D = \{a, b\}$, (63) obtains the meaning in (64), which is equivalent to (61).

(64) $[\![(63)]\!] = 1$ iff $\exists y \in D \cap \mathit{Part}_{\text{AT}}([\![\text{the kids}]\!])[\mathsf{laughed}(y) = 1]$. (Bar-Lev 2021:1062)

Naturally, this immediately obtains the intended meaning for negative sentences, which mean that there is no individual y that (*i*) is in the domain, (*ii*) is part of the denotation of *the kids*, and (*iii*) laughed. Thus, we get the intended meaning that no child laughed.

In the positive, the sentence must be strengthened to obtain the intuited universal meaning. For Bar-Lev (2018, 2021), the alternatives for Exh are obtained by replacing the sentence's domain with subdomains:

(65) ALT = {Aisha laughed $\vee$ Ben laughed, Aisha laughed, Ben laughed}

The subdomain alternatives 'Aisha laughed' and 'Ben laughed' are not innocently excludable. Excluding them would dysfunctionally result in the sentence meaning that Aisha or Ben laughed, but neither Aisha nor Ben laughed. I repeat the definition of Innocent Exclusion from chapter 1:

(66) **Innocent Exlusion procedure**: (Bar-Lev 2021:1066)

a. Take all maximal sets of alternatives that can be assigned false consistently with the prejacent.
b. Only exclude (i.e., assign false to) those alternatives that are members in all such sets—the **Innocently Excludable** (= IE) alternatives.

What is more, the set of alternatives (65) is not closed under conjunction: there is no strong alternative of the form 'Aisha laughed *and* Ben laughed' for Exh to exclude. Thus, exhaustifying (63) with the alternatives in (65) does not result in anything being excluded.

From here, Bar-Lev relies on the notion of Innocent Inclusion. As we saw in chapter 1, this is the idea that Exh *includes* all alternatives that are not excluded and which can be included consistently:

(67) **Innocent Inclusion procedure**: (Bar-Lev 2021:1067)

a. Take all maximal sets of alternatives that can be assigned true consistently with the prejacent and the falsity of all [innocently excluded] alternatives.
b. Only include (i.e., assign true to) those alternatives that are members in all such sets—the **Innocently Includable** (= II) alternatives.

Thus, Exh asserts the alternatives 'Aisha sang' and 'Ben laughed.' This results in the meaning that all the children laughed.

Carrying Bar-Lev's theory over to subatomic homogeneity

Plurals are made up of a finite set of non-overlapping parts, namely atoms. The arbitrary pieces of atoms, on the other hand, are infinite (any piece can be further subdivided into other pieces)

and can overlap (pieces are made up of other pieces). Recall from section 5.2.3 that subatomic homogeneity is a paradigm that touches on arbitrary pieces, not salient parts.

Let's see whether (and how) Bar-Lev's theory for plural homogeneity can be carried over to subatomic homogeneity. We need three things:

(68) a. existential meaning for predicates like *green* (whether obtained lexically or through an operator);
b. the ability to create alternatives for the sentence by restricting the 'domain' (i.e., ignoring some subatomic pieces of the subject); and
c. innocently including all pieces through Exh.

(68a) can in principle be obtained through an existential lexical entry for summative predicates like *green*; (69) is repeated from (40).

(69) $[\![\text{green}]\!] = \lambda x.\exists y[y \sqsubseteq x \wedge \text{green}(y)]$.

This, however, lacks the domain variable that Bar-Lev relies on to create subdomain alternatives (i.e., what we need for (68b)). Instead, let's follow Bar-Lev in postulating an operator taking a domain variable—call it ∃-SG:

(70) a. $[\![\exists\text{-SG}]\!] = \lambda D_{\langle et\rangle}.\lambda P_{\langle e,st\rangle}.\lambda x_e.\exists y \in D \cap Part(x)[P(y) = 1]$.
b. $Part(x) = \{y : y \sqsubseteq x\}$

The quantifier domain D in (70) must be populated not just by individuals (atoms) and pluralities, but also subatomic pieces. This is just as the 'domain of discourse' must contain more than just atoms and pluralities. Link (1983) suggests a domain of discourse for 'portions of matter' in addition to the domain of discourse for individuals, and Marcin Wągiel (p.c.) suggests there could be a general domain of discourse consisting of all types of entities: portions of matter, atomic individuals, pluralities, and clusters of pluralities. As such, there is nothing odd about taking the quantifier domain in (70) to involve subatomic pieces. For presentation, let's assume there are two subatomic pieces to the flag, A and B—this, of course, is a toy model not to be taken seriously, given that atoms are in fact divisible into an infinite number of arbitrary, overlapping pieces.

From here, the idea for a sentence like (71a) is that *green* is an argument of the ∃-SG operator, whose domain argument creates subdomain alternatives (I return to the lexical meaning of *green* in this theory immediately below). Putting aside Exh for a moment, (71a) has the LF in (71b).

(71) a. The flag is green.
b. [$_{vP}$ [$_{DP}$ The flag] [$_{vP}$ is ∃-SG$_D$ green]]

Still following Bar-Lev, (71b) is exhaustified in positive sentences but not negative sentences:

(72) a. [$_{vP}$ Exh$_{ALT}$ [$_{vP}$ [$_{DP}$ the flag] [$_{vP}$ is ∃-SG$_D$ green]]]
b. [$_{vP}$ not [$_{vP}$ [$_{DP}$ the flag] [$_{vP}$ is ∃-SG$_D$ green]]]

This is meant to obtain the homogeneity effect—and indeed, it does, for the simple sentences we are looking at. Let's see how this works.

There are two ways to make (72) work, depending on whether we assume that colour terms are lexically existential or universal. If *green* is existential, the meaning we get from (72) for negative

sentences is that there is no piece with any green on it, as desired.

(73) $[\![(72b)]\!] = 1$ iff $\neg\exists y \in D \cap \mathit{Part}(\iota x[\mathsf{flag}(x)])[\exists z[z \sqsubseteq y \wedge \mathsf{green}(z)]]$.

Prior to exhaustification, the meaning we get in the positive is that there is a piece of the flag which is partly green:

(74) $[\![(71b)]\!] = 1$ iff $\exists y \in D \cap \mathit{Part}(\iota x[\mathsf{flag}(x)])[\exists z[z \sqsubseteq y \wedge \mathsf{green}(z)]]$.

On our toy model where the pieces are A and B, (74) is equivalent to (75):

(75) $\exists x[x \sqsubseteq \mathrm{A} \wedge \mathsf{green}(x)] \vee \exists x[x \sqsubseteq \mathrm{B} \wedge \mathsf{green}(x)]$.

The alternatives triggered by the D variable in (74) are of the form 'subatomic piece A has a green piece,' 'subatomic piece B has a green piece,' and so on:

(76) $\textsc{alt} = \left\{ \begin{array}{c} \text{A is green}_\exists \vee \text{B is green}_\exists, \\ \text{A is green}_\exists, \\ \text{B is green}_\exists \end{array} \right\}$

Of course, the set of alternatives is actually infinite because the subject (the flag) can be cut up in an infinite amount of overlapping pieces of arbitrary sizes.

Following the Innocent Inclusion of all the alternatives in (76), we obtain the meaning that all pieces are partly green. Given that pieces can be subdivided into further pieces, this means that all pieces are in fact entirely green, as described a few times already in this chapter (if there was a piece that was not entirely green, clearly, its non-green subpiece would count as a piece that is not green at all, ensuring that the sentence would be false, as desired). Thus, (72) is compatible with *green* being lexically existential.

On the other hand, (72) would need a small modification to be compatible with universal colour terms. If *green* is universal, the meaning we get in the positive is that there is a piece that is entirely green:

(77) $[\![(71b)]\!] = 1$ iff $\exists y \in D \cap \mathit{Part}(\iota x[\mathsf{flag}(x)])[\forall z[z \sqsubseteq y \rightarrow \mathsf{green}(z)]]$.

In principle, once all pieces are innocently included, we get the meaning that all pieces are entirely green, as desired. The only issue is that to get to this, we first need to ensure that there is no innocently *excludable* subdomain alternative. But the subdomain alternative containing only the piece corresponding to the entire atom *would* in fact be innocently excludable (it is stronger than the prejacent, because it means 'the entire flag is entirely green' rather than 'some piece of the flag is entirely green'). To get around this, we must ensure that the piece corresponding to the entire flag is not in fact part of $\mathit{Part}(\iota x[\mathsf{flag}(x)])$. (78) is a modification of (70b), where $\sqsubseteq$ is replaced by $\sqsubset$.[15]

(78) $\mathit{Part}(x) = \{y : y \sqsubset x\}$

[15]This is not an issue that arises for Bar-Lev's analysis of pluralities, because for Bar-Lev, the plural is necessarily existential. The problem we are considering only arises if colour terms are *universal*, hence creating the strong alternative 'the subatomic piece corresponding to the entire flag is entirely green,' which is stronger than the assertion.

This solves the issue: given (78), there is no subdomain alternative meaning that the flag is entirely green. No alternative is innocently excludable, and all alternatives can therefore be included.

As for negative sentences, if colour terms are lexically universal, the meaning is that there is no piece that is entirely green:

(79) $[\![(72b)]\!] = 1$ iff $\neg\exists y \in D \cap \mathit{Part}(\iota x[\mathsf{flag}(x)])[\forall z[z \sqsubseteq y \rightarrow \mathsf{green}(z)]]$.

This may initially seem too weak (we want no piece to be green *at all*), but it is not. For example, if there was a half-green, half-white piece, the sentence *the flag is not green* would be false, as desired, because the green half of the piece would itself be a piece of the flag.

5.3.4 Interim summary

In this section, I have overviewed two theories of homogeneity, one of which captures the paradigm through an EMP, the other by positing that parts are included through exhaustification. I have shown how both can be made to work for subatomic homogeneity too, at first glance. In the rest of this chapter, I turn to data with *conjoined* summative predicates (in section 5.4), in order to show (in section 5.5) that neither of the theories just outlined can deal with these conjunctions. In contrast, the controlled-Exh account I have given in this thesis can.[16]

5.4 Conjoined predicates with plural vs. atomic subjects

To begin our discussion of conjunction, we first observe two different interpretations for conjoined predicates taking a plural subject (section 5.4.1), namely so-called Boolean and non-Boolean conjunction. These are terms that I will use descriptively; a Boolean conjunction can be described in terms of set-intersection, while a non-Boolean conjunction is any other kind of conjunction. We then turn to conjoined predicates with *atomic* subjects, to see why it may appear appealing to view these through the same lens (section 5.4.2). I show in section 5.4.3, however, that conjunctions predicated of atomic subjects are in fact always Boolean. This means that conjoined summative predicates must be weak even in some positive sentences, raising a new desideratum for theories of subatomic homogeneity (which I will show in section 5.5 that the theories discussed so far, as applied to the subatomic paradigm, do not meet). In section 5.4.4, I wrap up the discussion by showing that the weak intuited meaning of conjoined summative predicates in positive sentences is not the result of pragmatic weakening ('non-maximality'); the data from conjunction truly show that summative predicates can be semantically weak in some positive sentences.

5.4.1 Boolean and non-Boolean conjunction with pluralities

Let's start with a very simple question about plural homogeneity with distributive predicates: what happens to the homogeneity paradigm if a plurality is the argument not of a single predicate, but rather of two conjoined predicates? With some examples, the homogeneity effect is still observed.

[16] The argumentation from conjunction in this chapter could also have been made with the data from additive particles focused on in chapter 2. However, given that there is debate about elementary aspects of additive particles in the literature (see chapters 1 and 3) including whether and how they interact with Exh, I will simply adhere to conjunction in this chapter.

For our purposes, it will suffice to observe the positive side of the homogeneity paradigm. In the examples in (80), predicates which are true of all parts of the plurality *the tourists* in the basic (conjunction-free) sentences in (80a–b) both continue being true of all parts of the plurality when they are conjoined (80c):

(80) a. The tourists wore bathing suits.
≈ all the tourists wore bathing suits
b. The tourists swam.
≈ all the tourists swam
c. The tourists wore bathing suits and swam.
≈ all the tourists wore bathing suits and all the tourists swam

The conjunction in (80c) is 'distributive,' 'intersective,' or 'Boolean': the sentence is true if both (80a) and (80b) are true, i.e. if the tourists are in the intersection of the sets of individuals who wore bathing suits and individuals who swam. Descriptively, the universal quantificational force associated with the positive half of the homogeneity paradigm, observed in (80a) and (80b), effectively percolates to the entire conjunction (80c).

However, this does not hold across-the-board. Many predicates have strong meanings in simple (conjunction-free) sentences but weak meanings when they are conjoined. I will use the predicates *bark* and *crow*, which Krifka (1990) uses as well to make a similar point, on the assumption that some animal species bark (e.g. dogs) and others crow (e.g. roosters), but no single species both barks and crows. Hence, *bark* and *crow* are 'incompatible' predicates in some sense—they have an empty intersection. With a plural subject, the predicates *bark* or *crow* distribute over the plurality, as is now familiar from our discussion of homogeneity effects.

(81) a. The animals barked.
≈ all the animals barked
b. The animals crowed.
≈ all the animals crowed

But crucially, the strong, distributive meanings of *bark* and *crow* disappear with conjunction:

(82) The animals barked and crowed. (cf. Krifka 1990:165)
≈ some of the animals barked, and all of the others crowed

In (82), *barked* and *crowed* are interpreted as weaker than universal: (82) does not mean that every animal barked, and every animal crowed. Interpreting the sentence in this way is forced in this example because, as discussed above, no species both barks and crows.

At the same time, the predicates are stronger than merely being existential, because the sentence still means that all the animals partook in barking or in crowing. This may not be so salient for (82): pragmatically weak (existential) meanings of *bark* and *crow* are particularly salient, because we very easily imagine a discourse scenario for (81) and (82) where the speaker is complaining about noise, so that it does not matter whether all the animals or only some of them engaged in barking/crowing (see section 5.4.4 on non-maximality). But the effect is clearly discernable with other predicates, e.g.:

(83) The children are 7 and 10 years old.

≈ some of the children are 7 years old, and all of the other children are 10 years old

Clearly, (83) is stronger than only meaning that some children are 7 years old and some children are 10 years old. In fact, they are all of one of the two ages.

For Krifka (1990), the 'less than universal, more than existential' meaning of the predicates in (82) is the result of a non-Boolean *and*, i.e. a conjunction which is not set-intersective. While the sentence is in principle ambiguous between a Boolean and non-Boolean conjunction, the world knowledge described above forces a non-Boolean interpretation in this particular example. This can be made particularly clear from the fact that infelicity obtains if we attempt to have the predicates distribute over the plurality explicitly:

(84) #The animals barked and the animals crowed.

This is in contrast to Boolean *and*, naturally, where biclausal paraphrases are possible:

(85) The tourists wore bathing suits and the tourists swam.

Krifka (1990) argues that non-Boolean *and* lexically makes reference to parts of the plural subject, so that one predicate (*bark*) applies to one part, and the other (*crow*) to the other. Hence, (82) is interpreted as in (86). For Krifka (1990), these truth conditions are the result of the meaning of the conjunction, rather than any weakness in the predicates' underlying lexical meaning or the meaning of any plural operator.

(86) $\exists x, x'[\iota y[y \in \max_{\sqsubseteq}(*\mathsf{animal})] = x \oplus x' \wedge \mathsf{bark}(x) \wedge \mathsf{crow}(x')]$ (Krifka 1990:165)

Krifka's purpose is to explain how Boolean and non-Boolean *and* are related, and he goes on to revise some aspects of (86), but he doesn't deviate from the insight that *and* can create a representation like (86) where predicates can apply to different parts of the subject. For the purposes of this chapter, it will suffice to stipulate that there are two meanings for *and* (when it conjoins predicates), one Boolean and one not:[17]

(87) a. $[\![\text{and}_1]\!] = \lambda P.\lambda Q.\lambda x.\ P(x) \wedge Q(x)$.
b. $[\![\text{and}_2]\!] = \lambda P.\lambda Q.\lambda x.\ \exists x', x''[x = x' \oplus x'' \wedge P(x') \wedge Q(x'')]$. (Krifka 1990)

On this view (adopted for simplicity), *and* is lexically ambiguous and speakers choose with which *and* to conjoin predicates based on some notion of naturalness in a given situation (see Poortman 2017 for discussion of how the choice of predicates in a conjunction affects the likelihood of speakers preferring to interpret a conjunction intersectively). Indeed, while some predicates force a non-Boolean reading due to the conceptual incompatibility of two predicates (think of *bark* and *crow*: no animals does both those vocalizations, so the intersection of the sets denoted by these predicates is empty), most predicates are compatible with both readings of *and*. In fact, even (88), repeated from (80c), is compatible with a reading where the tourists are divided between bathing-suit wearers and swimmers (although it is rather difficult to get):

[17] This is probably not right, of course (e.g., Schmitt 2021). For work claiming that there is a single *and* that is underlyingly intersective, see e.g. Winter 2001, Champollion 2016, and Schein 2017; for work claiming that there is a single *and* that is underlyingly non-intersective, see e.g. Krifka 1990, Heycock & Zamparelli 2005, and Schmitt 2013, 2019.

(88) The tourists wore bathing suits and swam.
  a. ≈ all the tourists wore bathing suits and all the tourists swam
  b. ≈ some but not all of the tourists wore bathing suits, and all the other tourists swam

5.4.2 Prima-facie non-Boolean conjunctions with atoms

The distinction between Boolean and non-Boolean conjunction with plurals is often taken to hold with atoms as well. The possibility of a Boolean interpretation is straightforward:

(89) The tourist is drunk and tired.

As for non-Boolean conjunction, Krifka (1990) assumes that conjunctions of summative predicates like (90) are necessarily non-Boolean, because the conjoined predicates are *prima facie* contradictory (91).

(90) a. The flag is green and white. (Krifka 1990)
  b. This is beer and lemonade.

(91) a. #The flag is green and the flag is white.
  b. #This is beer and this is lemonade.

These predicates are 'contradictory' in the sense that summative predicates apply to all the material parts of their argument in positive sentences, and within certain classes of predicates (colours, materials), two cannot be true of a given part at the same time. This is not a logical contradiction but a contradiction brought about by the world knowledge that a given space cannot be taken up by two colours or two materials at the same time. Krifka's (1990) view is that in light of this, it must be that *and*, by quantifying existentially over parts of the subject (87b), 'breaks up' the subject into proper parts in (90), as in (92). In (92), it is tacitly assumed that green and white have universal quantificational force.[18]

(92) $\exists x, x'[\iota y[\mathsf{flag}(y)] = x \oplus x' \wedge \mathsf{green}(x) \wedge \mathsf{white}(x')]$ (Krifka 1990:165)

[18]There is another analysis of the compatibility of colour terms in conjunctions proposed by Lasersohn (1995). He explains (ib) by claiming the colour terms are nouns, as opposed to the adjective in (ia).

(i) a. The flag is white.
  ⇒ white ≈ all white
  b. The flag is white and green.
  ⇒ white ≉ all white

However, the same judgments as in (i) also hold in French, which can morphologically mark the colour terms in (ib) as adjectival (the colour terms can agree with the subject; they are feminine in (ii)).

(ii) a. La chambre est blanche.
  the.F room.F is white.F
  'The room is white.'
  b. La chambre est verte et blanche.
  the.F room.F is green.F and white.F
  'The room is green and white.'

For this reason, Lasersohn's proposal lacks generality.

This all seems entirely well justified. After all, (92), where a non-Boolean *and* makes reference to parts of an atomic individual, is only very minimally different from cases like (86), where a non-Boolean *and* refers to parts of a plurality. However, in the next subsection (section 5.4.3), I will turn to showing that the conjunctions in (90) are actually Boolean, in spite of initial appearances. In fact, conjunction with atomic subjects is never non-Boolean.

The lack of biclausal paraphrases in (91) may seem to be a problem for such a claim. However, looking back to my own analysis for a moment, the lack of biclausal paraphrases is in fact expected from my discussion of the locality constraint on the exhaustification of taxonomic predicates. Indeed, I argued in chapter 2 that there is a rather strict locality constraint on the Exh that strengthens predicates, including lexically weak summative predicates. As such, on my view, (90)–(91) can be understood as follows: the summative predicates in (90) are close enough syntactically for a single Exh to scope above both, but they are too structurally distant from one another in (91). For this reason, they are strengthened irrespective of one another, leading to a contradiction ('the flag is green and no other colour, and the flag is white and no other colour'). See chapter 6 for a proposal on how to model the locality constraint on Exh, which predicts exactly the pattern in (90)–(91). Hence, when I will claim in section 5.4.3 that there is no non-Boolean *and* with atomic subjects, I will assume that the data in (91) can be understood in such a way that they do not actually show that the predicates are lexically incompatible.

5.4.3 Conjunction with atoms is always Boolean

To see that conjunction with an atomic subject is always Boolean, we begin by looking at the colour-term conjunctions in more depth, making two observations. First, there are expressions which can be used to exclude non-Boolean interpretations (specifically *both* and *as well as*), but these can be used with colour conjunctions without creating a contradiction. Thus, a Boolean interpretation must at least be available for (90). Then, I will show that the Boolean interpretation must in fact be the only available interpretation, in light of the way colour conjuncts behave with modifiers that make them explicitly universal (e.g., *entirely*). Krifka's account overgenerates in predicting consistency in such sentences. After this discussion of colour terms, I will turn to other predicates and show once again that Krifka's proposal overgenerates. Indeed, outside the domain of summative predicates, apparently non-Boolean interpretations of conjunctions with atomic subjects like in (90) are straightforwardly unavailable as a matter of course. Finally, I will bring up the data with additive particles focused on throughout much of this thesis again as an independent way to see that colour terms can be 'joined' intersectively.

We begin with colour conjunctions, to see that Krifka's account undergenerates in one way and overgenerates in another.

An explicitly Boolean conjunction

My first argument manipulates the conjunction itself, to see whether Krifka (1990) is right to claim it is always non-Boolean with colour terms. In fact, we will observe that a Boolean interpretation of colour conjunctions must at least be available.

As it turns out, there is a morphologically overt way of controlling whether a given instantiation of *and* is Boolean or not. A Boolean reading of conjunction can be controlled for by using the expressions *both* or *as well as*; these are incompatible with non-Boolean interpretations. At least

when it composes with conjoined predicates (Schwarzschild 1996:149), *both* forces a Boolean interpretation of a conjunction; this is also the case with *as well as* (see comments along these lines by Szabolcsi & Haddican 2004:227, Gawron & Kehler 2004, Szabolcsi 2015:199, and Schmitt 2013, 2021). (93) shows that these 'marked' conjunctions are acceptable with a Boolean interpretation, while (94) shows that they do not allow for a non-Boolean interpretation.

(93) 'BOTH' AND 'AS WELL AS' WITH BOOLEAN CONJUNCTIONS:
   a. The children are both blonde and Swedish.
   b. The children are blonde as well as Swedish.

(94) 'BOTH' AND 'AS WELL AS' WITH NON-BOOLEAN CONJUNCTIONS:
   a. (i) The children are (#both) singing and talking at the same time.
      (ii) Aisha and Ben are (#both) 28 and 31 years old. (adapted from Krifka 1990)
   b. (i) #The children are singing as well as talking at the same time.
      (ii) #Aisha and Ben are 28 as well as 31 years old.

An anonymous *Journal of Semantics* reviewer takes issue with the claim that *both* marks conjunctions as Boolean, focusing on (95). Recall that this datapoint is due to the key piece of world knowledge that no species of animal both barks and crows.

(95) #The animals both barked and crowed.

The reviewer points out elaborations like (96), where (95) is given discourse context and clearly acceptable.

(96) I hate it when the animals keep me awake. Last night, they both barked and crowed!

But this is not a counter-argument to the claim that *both* marks conjunction as Boolean. The context in (96) makes it clear that any minimal amount of barking or crowing 'matters' to the speaker because it contributes to keeping them awake. (96) is a 'team credit' (Lasersohn 1999) reading where the verbs *bark* and *crow* are interpreted existentially (essentially an extreme non-maximality effect making plural predication existential in positive sentences). On the team-credit reading, *The animals barked* means that one or more barked, not that they necessarily all barked. As such, *they both barked and crowed* in (96) means that the animals are in the intersection of pluralities with at least one barker, and pluralities with at least one crower. There is nothing non-Boolean about this. The fact that this sort of context is *required* to make *both* acceptable in (95)/(96) constitutes evidence in favour of the view that *both* marks conjunctions as intersective.

As such, we have at our disposal a set of lexical resources that exclude the non-Boolean interpretation of *and*. If colour conjunctions were necessarily non-Boolean, *both* and *as well as* should be incompatible with them. This is not the case:

(97) a. The flag is both green and white.
   b. The flag is green as well as white.

(97) suggests that the intuited meaning of (90) is actually compatible with a Boolean interpretation of *and*; this is an undergeneration problem for Krifka's account.

Explicitly total conjuncts

The second piece of evidence in favour of a Boolean interpretation of *and* comes from a test which, this time, manipulates not the conjunction but the conjuncts themselves. If the colour terms in conjunction data are really total (universal), we should be able to modify them in ways that bring out totality without this causing inconsistency. Indeed, there are ways to modify predicates that provide partial predicates with a total reading, while also allowing total predicates to keep their total meaning. This happens with modifiers like *completely* and *all*, for example. I'll refer to these as 'totalizers'; they substantially change the interpretation of partial predicates while trivially keeping total predicates total.[19] Hence, the test for colour conjunctions is to modify individual conjuncts with words like *all* or *completely*; if *and* is non-Boolean, nothing should go wrong.

To make sure the test is valid, it will be helpful to step back from colour terms for a moment, and look at predicates that have independently been argued to be partial and total. Take the pair *dirty* and *clean* (Yoon 1996). *Dirty* is felicitously predicated of x as long as some part of x is dirty, while *clean* requires every part of x to be clean.

(98) a. $[\![\text{The flag is dirty}]\!] = 1$ iff $\exists x[x \sqsubseteq \iota y[\mathsf{flag}(y)] \wedge \mathsf{dirty}(x)$.
b. $[\![\text{The flag is clean}]\!] = 1$ iff $\forall x[x \sqsubseteq \iota y[\mathsf{flag}(y)] \rightarrow \mathsf{clean}(x)$.

But when *dirty* is modified with totalizers like *completely* or *all*, the resulting phrase behaves as 'total': all parts of the subject arguments must now be dirty. If the flag under discussion only has a few dirty spots, (98a) is true but both sentences in (99) are false.

(99) a. The flag is all dirty.
b. The flag is completely dirty.

These totalizers can also modify total adjectives, but in this case, it can be hard to pinpoint how the meaning is changed (cf. Rotstein & Winter's (2004:283) comments on "the difficulty that many speakers have in teasing apart the meaning of a total adjective *T* from the meaning of *completely T*"). The difference might have to do with *completely* indicating that totality holds to a particularly stringent standard (Sassoon & Zevakhina 2012), or perhaps its presence marks that the speaker is not speaking loosely (on loose talk, see e.g. Lasersohn 1999 and Hoek 2019). What matters for our purposes is simply that totalizers can felicitously and more or less trivially modify total predicates. Both sentences in (100) mean something quite similar to (98b).

(100) a. The flag is all clean.
b. The flag is completely clean.

Thus, regardless of whether they modify a predicate P that is lexically partial (99) or total (100), *completely* and *all* force the meaning that P is true of all parts of P's argument.

This makes a simple prediction for colour term conjunctions. If the *and* in colour conjunctions can be given a non-Boolean interpretation when the subject is atomic, ensuring that the conjunction is consistent even with total colour terms, it should be possible to make the colour terms' totality explicit via modification by a totalizer. This should bring about nothing more than a small change

[19] See e.g. Rotstein & Winter (2004), Kennedy & McNally (2005), and Sassoon & Zevakhina (2012) for formal work on totalizers like *completely*. The empirical contrasts pointed out in this section stand regardless of how one wants to formalize totalizers' meanings.

parallel to the difference between (98b) and (100). After all, Krifka's proposal is to arrive at the truth conditions in (101), repeated from (92), by predicating *total* colours of parts of the flag. The paraphrase for (101) is that 'part x of the flag is completely green and part x' of the flag is completely white,' so we should be able to overtly add *completely* in the conjuncts.

(101) $\exists x, x'[\iota y[\mathsf{flag}(y)] = x \oplus x' \wedge \mathsf{green}(x) \wedge \mathsf{white}(x')]$

Before checking whether this prediction is true for atomic-subject examples, let's see whether the test works as expected for plural-subject examples with a non-Boolean *and*. It should be possible to totalize the conjuncts; and indeed, as shown in (102), it is acceptable to make conjuncts explicitly total when the subject is plural. The result in an interpretation where the colours are total and each colour only applies to a subset of the plural—a textbook example of non-Boolean conjunction.

(102) a. The flags are completely green and completely white.
b. The flags are all green and all white.

Everything is now in place to see whether a non-Boolean interpretation of conjunction with non-atomic subjects is available. The equivalent of (102) with a singular flag in the subject should be acceptable, if non-Boolean *and* can break up atoms. In fact, totalizing each conjunct with an atomic subject leads to a sharp contradiction:

(103) a. #The flag is completely green and completely white.
b. #The flag is all green and all white.

Both the sentences in (103) mean that the entire flag is simultaneously of both colours. This is unexpected if (101) was the right semantic representation; (103) would be able to mean that one part of the flag is all green, and the other part is all white. The conclusion is that it is not possible for a non-Boolean *and* to refer to parts of atomic individuals. Krifka's account overgenerates: not only is a Boolean interpretation available with colour conjunctions, it is in fact the only possibility.

Many predicate conjunctions lack even prima-facie non-Boolean interpretations

A similar overgeneration problem holds for Krifka's account outside the domain of colour terms. Indeed, the lack of non-Boolean conjunction with atomic subjects is not a special fact of colour adjectives. Consider examples like (104) (cf. Winter 2001:69ff). These are all contradictory, but this is unexpected if atomic subjects could compose with a non-Boolean conjunction. With a non-Boolean conjunction, the examples in (104) would have meanings paraphrasable as 'the forest has a sparse part and a dense part,' 'the table has a wet part and a dry part,' and so on.

(104) a. #The forest is sparse and dense.
b. #The table is wet and dry.
c. #The land is flat and hilly.
d. #The room is hot and cold.

For reasons I cannot explain, not all of these are perfect with pluralities either. For instance, (105) is not clearly much better than (104a).

(105) ??The forests (in Canada) are sparse and dense.

There are relatively clearer contrasts with other examples, however, such as (104d) and (106).

(106) ?The buildings downtown are hot and cold.

Either way, there is nothing special about colours: as a rule, conjoined predicates with an atomic subject are only compatible with a Boolean interpretation.

Additive particles

Another problem for the claim that conjunction with colour terms and atomic subjects is always Boolean is that it makes the prediction that colour terms could not be 'joined' by an additive particle. Additive particles have never been argued to be anything but Boolean. Yet, as seen throughout this thesis, additives make colour terms consistent:

(107) The white flag is #(also) green.

Much of the argumentation in this chapter that will rely on the claim that conjunctions with atomic subjects are Boolean could alternatively be presented using additive particles, which have never been argued not to be Boolean.

5.4.4 Conjoined colours are not weakened through non-maximality

So far, we have observed that colour conjunctions (in fact, all predicate conjunctions) with atomic subjects can only receive a Boolean interpretation. If so, this must mean that under the right circumstances, a weak (non-universal) meaning is possible for colour adjectives in positive sentences.

Does this mean that colour terms are lexically weak? I believe so. But there is an alternative way to understand the consistency of Boolean colour conjunctions. It could be that colour terms are lexically universal, but they undergo pragmatic weakening when they are conjoined, in order for a contradiction to be avoided. In this section, I show that the conjuncts in fact do not undergo weakening. The conclusion will be that summative predicates must be lexically compatible (i.e., existential) given that their Boolean conjunction is intuited as non-contradictory.

So far in this chapter, I have mostly overlooked one important component of homogeneity.[20] Homogeneity effects can allow some exceptions; this is known as 'non-maximality.' For instance, (108) can be judged as true if most but not all of the professors smiled (Križ 2015).

(108) The professors smiled.

Depending on pragmatic factors, non-maximality can go so far as to yield an existential meaning. (109) is from Malamud (2012:4–5).

(109) a. SCENARIO: *Aisha has a large house with over a dozen windows in different rooms. She locks up and leaves to go on a road trip with her friend Ben, forgetting to close just a few of the many windows in various rooms. A few minutes into the ride,*

[20] I thank Nina Haslinger for raising this point.

Ben says, "There is a thunderstorm coming. Is the house going to be OK?" Aisha replies:

b. Oh my, we have to go back—the windows are open!

Given the scenario in (109a), any number of open windows warrants going back home, and (109b) can be interpreted existentially.

In light of this, one could claim that, even if colour conjunctions are necessarily Boolean, this does not necessarily show us anything about the basic meanings of colour terms themselves. After all, they could be used non-maximally in these conjunctions. Perhaps they are underlyingly universal in meaning (in positive sentences), but when they are conjoined with a Boolean *and*, each colour term is pragmatically weakened, in the same way that we observe weak meanings with pluralities in (108) and (109).

To evaluate this, I begin with a particular theory of non-maximality, and show that it cannot in fact take a semantically contradictory assertion and 'pragmatically' make it acceptable. Then I ask more generally what sort of prediction *any* theory would make if it tried to allow the pragmatic weakening of conjuncts in semantically contradictory conjunctions; I suggest that any such theory, even if it worked for colour terms, would overgenerate elsewhere in language.

Before we begin, let me address a possible objection to the non-maximality account of colour terms' existential meaning. The objection would be that, unlike the case of pluralities, there is no need for a special scenario for the existential meaning of colour terms to arise. However, this sort of objection does not stand up. Indeed, conjoined colour terms' existential meaning could be claimed to arise automatically (without any special scenario) because the existential interpretation is the only non-contradictory option available. Faced with the choice between a 'maximal' contradictory interpretation of conjunctions, and a non-maximal non-contradictory interpretation, speakers would choose the non-contradictory, non-maximal parse without requiring a special scenario. This is in contrast with pluralities, where no contradiction is inherent to 'maximal' readings.[21]

Nonetheless, the non-maximality account of colour terms' mutual compatibility does not work. I start by showing that a particular theory of non-maximality does not carry over to the colour conjunction data. Based on the intuition that non-maximal readings are possible when exceptions are pragmatically irrelevant to the issue at hand in a conversation, Malamud (2012) and Križ (2015), following van Rooij (2003), argue that in conversation, the set of possible worlds is partitioned, and the goal of a conversation is to determine which cell the real world is in. When an exception is irrelevant for current purposes, it is because it does not affect which cell of the partition the real world is in. On this approach, in (109), Aisha's statement semantically means that all the windows are open. But the worlds in which all the windows are open are in the same cell as all the worlds where there is even as few as a single window open, because Aisha and Ben must go back home to close the window(s) either way. To represent this, if w_1 is a world in which all the windows are closed, w_2 is a world in which only some windows are open, and w_3 is a world in which all the windows are open, the pragmatic partitioning of worlds is as in (110).

(110) $\{\{w_1\},\{w_2,w_3\}\}$

Identifying the real world as being in the set $\{w_2,w_3\}$ is 'good enough' pragmatically, even if Aisha's utterance is only semantically true in w_3, not w_2.

[21] Still, an important question is why we do not observe the required non-maximality outside of conjunctions.

This account does not carry over to colour term conjunctions like (111), however. Recall that, on the hypothesis that colour terms are lexically strong, and in light of the Boolean status of the conjunction, the semantic meaning of (111) is that the flag is both entirely green and entirely white.

(111) The flag is green and white.

Yet, there are no worlds in which a surface can be entirely of two colours all over. That is, there are no worlds in which (111) is semantically true. The theory of non-maximality just sketched out relies on the idea that non-maximal readings involve there being a world in which the sentence is semantically true, with the real world being in the same cell as that world. If there is no world in which the sentence is semantically true, it is impossible for the real world to be in the same cell as that world. Thus, under this theory, there is generally no way for a semantically contradictory sentence to be made acceptable through non-maximality.

Perhaps this is just a problem with Malamud and Križ's theory. As such, putting aside any actually-existing theory of non-maximality, let's return instead to the intuitive appeal of understanding conjunctions like (111) as non-maximality effects. The idea is quite simply that (111) is semantically contradictory, but due to some pragmatic contradiction-avoidance principle, this problem is solved pragmatically by weakening each conjunct. In fact, even this intuitive, non-formalized version of the non-maximality account immediately runs into a problem: giving this much power to the pragmatics overgenerates significantly. Consider any of the examples in (104), such as (112), repeated from (104b).

(112) #The table is wet and dry.

(112) is semantically contradictory: while *wet* is only existential, *dry* is universal, making the two predicates incompatible (Yoon 1996).[22] The pragmatic weakening account just sketched out informally for colour conjunctions (111) should kick in here as well, and weaken the offending predicate (here only *dry*) to an existential. (112) would mean that the table has a wet part and a dry part. Yet, this pragmatic weakening clearly does not occur, because the sentence is intuited as a contradiction. The same goes for Boolean conjunctions of pluralities as well, which clearly cannot be weakened in a way to avoid a contradiction:

(113) a. #The students are both asleep and awake.
b. #The windows are both open and closed.

In sum, while there is in theory a possibility of viewing Boolean conjunctions of colour terms as non-contradictory only due to pragmatic weakening of the conjuncts, this runs into significant problems. Under Malamud (2012) and Križ's (2015) account, the semantically contradictory nature of the conjunctions means that it is impossible for the actual world to be in the same cell as the set of worlds denoted by the sentence, because there is no such world. More generally, an im-

[22] An even closer correlate to (111) would be (i), where both predicates are universal given the modification of *wet* with *completely*:

(i) #The table is dry and completely wet.

This mirrors formally the conjunction of two colour adjectives in (111), since these adjectives are also supposedly both lexically universal, on the hypothesis that Boolean colour term conjunctions are acceptable due to the weakening of lexically strong colour terms.

portant empirical problem surfaces in examples like (112) and (113), where we observe that, as a matter of fact, semantically contradictory conjunctions are *not* fixed through pragmatic weakening of the conjuncts. This makes it undesirable to make such a claim for colour terms.

5.4.5 Interim conclusion

In this section, we saw that conjoined predicates with atomic subjects, as in (114), are always interpreted as having a Boolean conjunction.

(114) The flag is green and white.

We also saw that the conjoined predicates are not weakened through non-maximality. As such, what we learn from examples like (114) is that the predicates *green* and *white* must be lexically consistent; their intersection is not empty.

This places a new empirical burden on theories of homogeneity. It is not enough to describe the paradigm as involving universal quantification in positive sentences and negated existential quantification in negative sentences, as in (115).

(115) a. The flag is green.
≈ all parts of the flag are green
b. The flag is not green.
≈ no part of the flag is green

Another important part of the paradigm is that, in positive sentences with conjoined predicates (114), colour terms receive an interpretation that is weaker than universal, but also stronger than existential. (114) means that all parts of the flag are green or white, and would be false if the flag had some other colours on it.

5.5 Evaluating the theories of homogeneity

We have just seen that colour conjunctions like (116) are consistent while necessarily being Boolean.

(116) The flag is green and white.

In this section, I show that neither of the two alternative theories of homogeneity considered in this chapter, as applied to the subatomic paradigm, predict the right meaning for (116). In contrast, it is rather straightforward on the view of controlled exhaustification sketched out in chapter 2. I start with my own claim about (116), then move on to each competing theory in turn.

5.5.1 The controlled-Exh account

In chapter 2, based in part on data like (116), I claimed that colours' lexical meaning is not universal but existential. This is forced due to the fact that (116) is consistent, has a Boolean conjunction, and does not involve non-maximality. Thus, we need lexical entries like (117) (Harnish 1976; Levinson 1983).

(117) $[\![\text{green}]\!] = \lambda x.\ \exists y[y \sqsubseteq x \wedge \mathsf{green}(y)]$.

In addition to making conjunctions non-contradictory, the meaning for colour terms posited in (117) immediately explains the meaning of negative sentences, as I overviewed in section 5.3.1.

(118) The flag is not green.
$\rightsquigarrow$ the flag is not green at all

As for positive sentences, I posited an Exh operator taking the set of colour terms as its domain. Putting aside a locality constraint on this Exh operator (see chapters 2 and 6), the meaning of positive sentences is derived as follows:

(119) $[\![\text{Exh}_{\text{ALT}}\ [\text{the flag is green}]]\!]$
$= 1$ iff the flag is green$_\exists \wedge$ the flag is not white$_\exists \wedge$ the flag is not red$_\exists \wedge \ldots$

This means that the flag has a green part, and does not have a part of any other colour. Through some pragmatic strengthening, this comes to mean that the flag is entirely green (all parts must have a colour, after all). As for the 'more than existential, less than universal' meaning of colour conjunctions, if there is a single Exh operator above both colour terms (see chapter 6), this meaning is exactly what is predicted:

(120) $[\![\text{Exh}_{\text{ALT}}\ [\text{the flag is green and white}]]\!]$
$= 1$ iff the flag is green$_\exists \wedge$ the flag is white$_\exists \wedge$ the flag is not red$_\exists \wedge \ldots$

(120) means that the flag has a green part, has a white part, and does not have parts of any other colours. All pieces of the flag must therefore be covered in green or white.

In contrast, the conjunction data pose a problem for the other theories of homogeneity as applied to the subatomic paradigm.

5.5.2 The excluded-middle presupposition

To deal with Boolean conjunctions of colour terms, we need them to have existential lexical meanings. However, this is not compatible with the EMP. In section 5.3.2, I spelled out two possibilities for giving summative predicates an EMP: either it is part of their lexical meaning (121), or it comes from a $\text{DIST}_{\text{SUBAT}}$ operator (122). (121) is repeated from (50) and (122) from (53).

(121) $[\![\text{green}]\!] =$
$\lambda x : \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{green}(y)] \vee \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \neg\mathsf{green}(y)].\ \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{green}(y)]$.

(122) $[\![\text{DIST}_{\text{SUBAT}}]\!] =$
$\lambda P.\ \lambda x : \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow P(y)] \vee \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \neg P(y)].\ \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow P(y)]$.

Attempting to tweak (121) to obtain existential meaning is a non-starter. As mentioned in section 5.3.2, given that the EMP ensures that predication is undefined if the predicate only holds of some material parts, there is effectively no way to create an existential lexical entry. Writing out the output condition in (121) as if it was existential (123) changes nothing. (121) and (123) are equivalent, since both are only defined for arguments that are either entirely green or not green at all.

(123) $[\![\text{green}]\!] =$
$\lambda x : \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \mathsf{green}(y)] \vee \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \neg\mathsf{green}(y)].\ \exists y[y \sqsubseteq_{\text{AP}} x \wedge \mathsf{green}(y)].$

As such, it is not possible to capture the existential meaning that is necessary to make colour terms mutually compatible in Boolean conjunctions while stipulating that their lexical entries have an EMP. We need EMP-less entries like (117).

The alternative proposal in (122), where the EMP comes from an operator, is compatible with EMP-less lexical entries for summative predicate, but it too leads to problems. The two relevant lexical items are repeated in (124)–(125) from (122) and (117).

(124) $[\![\text{DIST}_{\text{SUBAT}}]\!] =$
$\lambda P.\ \lambda x : \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow P(y)] \vee \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \neg P(y)].\ \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow P(y)].$

(125) $[\![\text{green}]\!] = \lambda x.\ \exists y[y \sqsubseteq x \wedge \mathsf{green}(y)].$

As proposed by Gajewski (2005), the purpose of the EMP is mainly to strengthen the meaning of negative sentences; for positive sentences, its only effect is to make non-homogeneous cases undefined rather than false. The proposal embodied by (124)–(125) effectively means using the EMP for the polar opposite of this: the EMP would strengthen positive sentences, and have no effect on negative ones (other than making non-homogeneous cases undefined rather than false).

I already showed in section 5.3.2 that in positive and negative sentences without conjunction, (124) and (125) combine to yield the desired 'all green' and 'not green at all' meanings. What concerns us now is whether such a system would be able to capture the non-contradictory nature of Boolean colour term conjunctions. If $\text{DIST}_{\text{SUBAT}}$ can scope over both colour terms at once like in (126), perhaps it would not create a contradiction.

(126) The flag is [$\text{DIST}_{\text{SUBAT}}$ [green and white]].

Unfortunately for the EMP, this proposal will not, in fact, obtain the right results.

Let's take this step by step. Before the functional application of $\text{DIST}_{\text{SUBAT}}$, the meaning of the conjunction is as in (127) (the subject in (126) is atomic and, therefore, only a Boolean conjunction is possible).

(127) $[\![\text{green and white}]\!] = \lambda x.\ \mathsf{green}_{\exists}(x) \wedge \mathsf{white}_{\exists}(x).$

$\text{DIST}_{\text{SUBAT}}$ then composes with the conjunction, producing (128).

(128) $[\![\text{DIST}_{\text{SUBAT}}\ [\text{green and white}]]\!]$
$= \lambda x : \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow [\mathsf{green}_{\exists}(y) \wedge \mathsf{white}_{\exists}(y)]] \vee \forall y[y \sqsubseteq_{\text{AP}} x \rightarrow \neg[\mathsf{green}_{\exists}(y) \wedge \mathsf{white}_{\exists}(y)]].$
$\mathsf{green}_{\exists}(x) \wedge \mathsf{white}_{\exists}(x).$

There is no contradiction in (128); yet, it is still not a welcome result. The first disjunct requires all pieces of the flag to themselves have both a green and a white piece. As such, however small a piece you choose, it would have to be made up of a white piece and a green piece. And those pieces themselves would also need to be made up of a white piece and a green piece, and so on. Flags certainly don't work this way, nor indeed does any (mental representation of any) object. We would have needed this disjunct of the presupposition to contain a disjunction ('all parts are green *or* white') rather than a conjunction. As for the second disjunct, this is clearly hopeless: there is

no reason why an arbitrary piece of a white and green flag could not have more than one colour.

Recall that I suggested in section 5.2.3 that the subatomic homogeneity paradigm is most simply understood as referring only to arbitrary pieces, and not salient parts; but that my argument was rather weak, since it was only based on a notion of simplicity. Could we make (128) work by understanding $\sqsubseteq$ as referring to salient parts rather than arbitrary pieces? The answer is no; let's see this one disjunct at a time once again.

On the view that $\sqsubseteq$ would refer to salient parts, the first disjunct requires all salient parts of the flag to themselves have both a green and a white part. But this is absurd: as noted above, parts of flags are pragmatically salient parts by virtue of only having one colour. Here too, we would have needed this disjunct to itself be disjunctive ('all parts are green *or* white') rather than a conjunctive. As for the second disjunct, it means that no salient part of the flag has both colours. This is correct for flags (salient parts are made of single colours) but does not strengthen the colour terms' existential meaning; this disjunct is compatible with the presence of other colours on the flag, corresponding to other salient parts. Another problem arises for the second disjunct in the presupposition of (128) when we turn to objects other than flags. The negative disjunct in (128), if understood as referring to parts rather than pieces, works for flags because of an idiosyncrasy about flags: the parts of flags are made up of the colours themselves, and therefore the parts of flags have only one colour each. This is not the case with other objects; recall our rainbow-coloured table from section 5.2.3, for example; a given salient part (e.g., the tabletop) can have more than one colour.

Thus, (128) does not produce the strengthening effect we are trying to obtain, regardless of whether we understand subatomic homogeneity as referring to subatomic pieces or salient parts. The analysis under consideration, where subatomic homogeneity effects arise due to positive sentences being strengthened by an EMP contributed by a $\text{DIST}_{\text{SUBAT}}$ operator, can deliver the right results for both negative and positive non-conjoined cases, as shown in section 5.3.2, but not for colour conjunctions.

5.5.3 Bar-Lev's exhaustivity account of homogeneity

The EMP account of homogeneity cannot capture the meaning of sentences with conjoined summative predicates and atomic subjects. What about Bar-Lev's Innocent Inclusion account?

Let's assume for simplicity some atom (a flag) and some partition of it into two pieces, call them A and B, putting aside for the time being complications from the fact that subatomic pieces, unlike the atomic parts of pluralities, can always be divided into further pieces. If so, the plain (pre-exhaustification) meaning of (129a) is as in (129b). Note that the colour terms have to be lexically existential; otherwise, the conjunctions in the disjuncts would be contradictory.

(129) a. The flag is white and green.
b. $[\![(129a)]\!] = 1$ iff (A is white$_\exists$ $\wedge$ A is green$_\exists$) $\vee$ (B is white$_\exists$ $\wedge$ B is green$_\exists$)

The alternatives are obtained by replacing the domain with subdomains:

(130) $$\text{ALT} = \left\{ \begin{array}{c} \text{A is white and green} \vee \text{B is white and green}, \\ \text{A is white and green}, \\ \text{B is white and green} \end{array} \right\}$$

Now we exhaustify (129). No alternative is excludable. As such, the alternatives are innocently included instead, and we end up with the meaning that A is partly white and partly green, and B is also partly white and partly green.

This is problematic in more than one way. In section 5.2.3, I claimed that the simplest assumption is that subatomic homogeneity refers to the arbitrary pieces of atoms, not the salient parts of atoms. As such, the meaning we obtain by innocently including all alternatives is that all arbitrary pieces of the flag are themselves divisible between a partly green piece and a partly white piece. However small a piece of the flag you choose, it would need to itself have both those colours, all the way down to the tiniest pieces. The flag would effectively have to be some sort of fractal. Even if we tried to get around this problem by claiming that the alternatives are not innocently includable as a result of this problematic outcome, this would create another problem: no strengthening would take place at all, and the sentence would only mean that some piece of the flag is partly green and partly white.

Moreover, nothing improves if we abandon the view of section 5.2.3 and claim that the parts referred to by subatomic homogeneity are not arbitrary pieces, but in fact salient parts. Indeed, the parts of flags are by definition made up of only one colour, but Bar-Lev's proposal creates a state of affairs where each part must have *each* colour asserted in the conjunction.

In conclusion, given the lack of non-Boolean conjunction with atomic subjects, Bar-Lev's account cannot carry over directly to the subatomic homogeneity paradigm. Innocently including arbitrary pieces leads to the predication of each conjunct of each piece. Of course, the fact that Bar-Lev's account does not work for subatomic homogeneity does not mean that it is incorrect for plural homogeneity. In fact, it remains that Bar-Lev's account and mine share an important component: they involve weak lexical meaning together with covert exhaustification in positive sentences, but not in negative sentences. It is possible that this is all that formally unites plural and subatomic homogeneity paradigms.

5.5.4 Interim summary

In this section, I showed that the exhaustivity account of subatomic homogeneity quietly built up in this thesis, according to which the assertion of one predicate excludes other related predicates, can obtain the right results for conjoined summative predicates with atomic subjects. The two other theories of homogeneity under consideration in this chapter cannot do so. The theory that Exh includes subdomain alternatives creates the meaning that both conjoined predicates must be true of all subatomic parts; the theory that a subatomic $\text{DIST}_{\text{SUBAT}}$ operator gives predicates an excluded-middle presupposition runs into a similar problem (the presuppositional disjunction requires all pieces of the flag to be of both colours, or no piece of the flag to be of both colours). I therefore maintain the basic claims of chapter 2: summative (and many other) predicates are lexically weak, and they are strengthened in positives sentences through the exclusion of like predicates.

To avoid the problems pointed to in this section, the EMP approach to subatomic homogeneity would have to claim that the $\text{DIST}_{\text{SUBAT}}$ operator is optional. It would yield no strengthening whatsoever in positive sentences, so that an account like mine would still be necessary. Likewise, Bar-Lev's account is unable to strengthen weak meanings at the subatomic level, so it too undergenerates. Nothing goes wrong for Bar-Lev's claims about plural homogeneity; it simply cannot be extended to subatomic homogeneity.

5.6 Conclusion

In this thesis, I discuss a paradigm touching on the meaning of predicates: many predicates are interpreted as incompatible, unless they are joined by *also* or *and*:

(131) a. #This comedy is a tragedy.
b. This comedy is also a tragedy.
c. This is both a comedy and a tragedy.

(131) involves integrative predicates, which are true of individuals without being true of their subatomic parts. The effect is also found with summative predicates:

(132) a. #The white flag is green.
b. The white flag is also green.
c. The flag is both white and green.

(132) can be described in terms of the quantificational force of the predicate, as universal (132a) or existential (132b–c). Another paradigm where the quantificational force of summative predicates varies between being universal or existential is the contrast between positive and negative sentences (133)—the homogeneity paradigm.

(133) a. The flag is green.
b. The flag is not green.

The theory of controlled exhaustivity developed in this thesis collapses (131) and (132) by claiming that both are the result of predicates being exhaustified vis-à-vis their cohyponyms (or something like this: see chapter 4). Unless they are joined by *also* or *and*, *comedy* comes to mean 'not a tragedy' and *green* comes to mean 'not white.' However, given the rich literature on homogeneity, this analysis for colour terms (132) must be defended against alternatives. Focusing on two theories (one obtaining the homogeneity paradigm through an excluded-middle presupposition, and the other through the Innocent Inclusion of subdomain alternatives), I have shown that a problem arises for these theories in conjunction data (132c). The predicates in (132c) must be existential given the Boolean nature of the conjunction. But neither theory of homogeneity considered in this chapter can capture the weakness of these predicates in these particular positive sentences, predicting (132c) to mean that all parts of the flag are *both white and green*.

On the other hand, if summative predicates are lexically weak and are strengthened through the exclusion of alternative predicates, (132c) can be understood as involving an Exh operator above the conjunction excluding all colours except white and green:

(134) $[\![$The flag is $[\text{Exh}_{\text{ALT}}$ [white and green]]$]\!]$
$=$ 1 iff the flag is $[\text{white}_{\exists}$ & $\text{green}_{\exists}$ & not $\text{red}_{\exists}$ & not $\ldots]$.

In chapter 6, the last chapter of this thesis, I turn more closely to data like (134). In particular, I have claimed in chapter 2 that Exh is subject to a locality requirement with cohyponyms; but in (134), there is at least a small distance between it and each cohyponym (Exh must scope above both at once). Chapter 6 asks what the nature of the locality requirement on Exh is, focusing largely on data from conjunction like (134), and tries to account for it syntactically.

A remaining empirical puzzle for my account of subatomic homogeneity (in fact, any such account) pertains to a difference between the two main types of summative predicates, viz. colour and material terms:

(135) a. The table is green.
b. The table is wood.

While colour terms are readily intuited as weak when joined by *also* (136a), this is not the case with material terms (136b). This is in spite of both otherwise patterning similarly: as we have seen, they are both interpreted as universal when not conjoined, and as fully consistent when conjoined with *and*.[23]

(136) a. The white flag is also green.
b. #The metal table is also wood.

It is not clear why the two classes of summative predicates should pattern differently, and in particular why material terms cannot be joined by an additive. I leave this for future work.

[23]In fact, even the sentence in (136a) is not fully accepted by everyone; some speakers find it more immediately readable as meaning that the fully white flag is somehow 'also' fully green. Still, I have yet to find anyone who fully rejects the sentence or denies that it is significantly better than the counterpart without *also*.

Chapter 6

What makes Exh obligatory and ultra-local

6.1 Introduction

In the previous chapters of this dissertation, I have made the case that taxonomic predicates are subject to a previously undescribed exhaustivity effect. The intuited mutual exclusivity of cohyponyms[1] is in many cases the result of grammatical strengthening, rather than lexical or conceptual meaning. For instance, I have argued that the contradiction intuited in sentences like (1) is not due to the lexical–conceptual meaning of either *comedy* or *tragedy*; in fact, the set-intersection of these concepts' exemplars is non-empty.

(1) #This comedy is a tragedy.

Rather, due to exhaustification, the predicate *comedy* comes to mean 'not a tragedy, not an epic, etc.' rather than only referring to the concept of COMEDY, and likewise for *tragedy*:

(2) ⟦The [Exh_{ALT} comedy] is a [Exh_{ALT} tragedy]⟧
= 1 iff the [comedy & not tragedy & not epic] is a [tragedy & not comedy & not epic].
⇒ contradiction

I have just defended in chapter 5 the view that this is the case not only with integrative predicates like *comedy*, but also with summative predicates like colour terms:

(3) ⟦The [Exh_{ALT} green] flag is [Exh_{ALT} white]⟧
= 1 iff the [$\text{green}_{\exists}$ & not $\text{white}_{\exists}$ & not $\text{red}_{\exists}$] flag is [$\text{white}_{\exists}$ & not $\text{green}_{\exists}$ & not $\text{red}_{\exists}$].
⇒ contradiction

The effect is new to the literature because of the empirical domain where strengthening is posited. Previously, outside of special contexts like intonational focus, the claim that any predicates are subject to exhaustification was essentially limited to predicates taking part in entailment scales, such as {*warm*, *hot*, *boiling*}. But the effect is also novel because of its twin formal characteristics. First, Exh is *obligatory* with taxonomic predicates. If Exh was optional, non-contradictory parses of the above sentences would be available. This is not a new claim by itself (obligatory Exh

[1] It would be more accurate to talk about 'same-jurisdiction predicates' (see chapter 4), but I will carry on with the rhetoric of chapter 2 in this chapter by mostly talking about 'cohyponyms' and 'cohyponymic exclusivity.'

operators are posited by e.g. Magri 2009, Chierchia 2013, and Bade 2016). But in addition to being obligatory, Exh with predicates is necessarily local to them. In (2)–(3), if a global Exh was available, its prejacent would entail both predicates, and no contradiction would be intuited.[2]

(4) $[\![\text{Exh}_{\text{ALT}}\ [\text{the green flag is white}]]\!]$
$= 1$ iff the green$_\exists$ flag is white$_\exists$ $\wedge$ the green$_\exists$ flag is not red$_\exists$.
$\not\Rightarrow$ contradiction

I have referred to these twin properties as the 'controlled' nature of Exh: it is 'controlled' by the alternative-triggering expression, viz. the predicate it associates with, which both requires its presence and constrains its syntactic locus. In this chapter, I sketch out a hypothesis about why Exh is controlled in this way.

There are several steps to asking what results in Exh being controlled, the first of which is to characterize what it means for Exh to necessarily be 'local' to its controller. How should the locality requirement be defined? I suggest that Exh has to be within the projection of the predicate. In particular, I will show that this fares better than alternative constraints which would require that there be no scope-bearing expression between Exh and its controller (a semantic constraint), or that Exh occur in the controller's phase (a slightly looser syntactic constraint). If correct, the locality constraint on Exh is syntactic rather than semantic in nature.

From there, I suggest a syntactic mechanism which accounts for the controlled nature of Exh with predicates. I posit that derivational morphemes (n^0, a^0, etc.) syntactically Agree with Exh. That is, they have a $[u\text{Exh}]$ feature, which probes for an interpretable $[i\text{Exh}]$ feature borne by Exh. For the syntactic computation to converge, derivational morphemes therefore require an Exh operator in their vicinity (Chomsky 2000, 2001, pace Preminger 2011). With the right syntactic assumptions, this captures both the obligatory nature and the locality requirement on Exh with predicates. The upshot of this proposal is that category morphemes are effectively involved in the meaning of their roots in non-trivial ways. While conceptual/lexical space involves overlap between concepts/predicates, category morphemes, when they take roots and integrate them into the syntactic computation, effectively clean up this non-partitioned conceptual space.

I will conclude the chapter by asking whether there are other semantic domains in which Exh is controlled. As we will see, there are analyses of semantic phenomena which have also tacitly put limits on the distance between Exh and the alternative-triggering expression, but without pointing it out. In particular, Bar-Lev's (2018; 2021) analysis of plural homogeneity requires that Exh is both obligatory and necessarily in the same clause as the alternative-triggering expression (the domain variable on his $\exists$-PL$_D$ operator). Thus, the constraints observed on Exh with taxonomic predicates are 'new' in the sense of never having been explicitly discussed, but they may be part of a larger class.

This chapter is organized as follows. I start in section 6.2 by reviewing the data where Exh is necessarily understood as both obligatory and local, as well as semantic areas in which Exh is apparently not. I suggest that for all the data where Exh might appear to be absent (under additives, conjunction, and sentential negation), it can be maintained that Exh is still indeed obligatory and ultra-local. For the negation data, this requires adopting the presuppositional 'Pexh' operator of

[2]Another problem for a global Exh, not shown in (4), is that if the attributive adjective *green* also triggers alternatives (see chapters 2 and 3 for evidence that it does), we would end up with entailments about other flags ('it's not the case that the *red* flag is white').

Bassi et al. (2021). This sets the stage to investigate Exh's controlledness without needing to account for any apparent counterexamples. Turning to this question, I first discuss the locality requirement on Exh in section section 6.3. I argue that it is best understood as meaning that Exh must appear within the XP of the predicate. In section 6.4, I ask why this is, and argue that derivational morphemes have a [*u*Exh] feature which must be valued by an Exh operator (cf. Chierchia (2013), who has an Agree relation involving Exh but in the reverse direction). Finally, section 6.5 points out that Exh is also obligatory and obligatorily local at least in Bar-Lev's (2018; 2021) account of plural homogeneity, suggesting the existence of a broader phenomenon. Section 6.6 concludes.

6.2 When controlled Exh seems to disappear

In chapter 2 (section 2.4.2), I argued that the Exh operator strengthening cohyponyms is always syntactically local to them. I left open how locality is best thought of (how close, and by what measure, must Exh be?), but pointed out that the locality requirement is more strict than simply requiring that Exh appear in the same clause as its associated cohyponym. I gave four types of sentences as evidence for this putative locality requirement: cohyponyms leading to sentence-internal contradictions (5a), cohyponyms under other scope-bearing expressions (5b), cohyponyms co-occurring with another alternative-triggering expression (5c), and cohyponyms in downward-entailing (DE) environments (5d).

(5) a. #The green flag is white.
b. Every flag is green.
c. Some flags are green.
d. If the flag is green, Aisha will sing.

For all these examples, the right meanings (i.e., the meanings displaying cohyponymic exclusivity) can be generated, with the stipulation that the Exh must be very local to the cohyponym. (6) shows this as a proof of concept for each sentence in (5), without yet seriously taking a stance about the exact locus of Exh (see chapter 2 for brief comments on semantic type). Pretend here and throughout that the only colours are green, white, and red.

(6) a. ⟦The [Exh_{ALT} green] flag is [Exh_{ALT} white]⟧
= 1 iff the [$\text{green}_{\exists}$ & not $\text{white}_{\exists}$ & not $\text{red}_{\exists}$] flag is [$\text{white}_{\exists}$ & not $\text{green}_{\exists}$ & not $\text{red}_{\exists}$]
⇒ contradiction

b. ⟦Every flag is [Exh_{ALT} green]⟧
= 1 iff every flag is [$\text{green}_{\exists}$ & not $\text{white}_{\exists}$ & not $\text{red}_{\exists}$]
≈ every flag is entirely green

c. ⟦Exh_{ALT} [some flags are [Exh_{ALT} green]⟧

$$= 1 \text{ iff } \left\{ \begin{array}{l} \text{some flags are } \left(\begin{array}{c} \text{green}_{\exists} \,\& \\ \text{not white}_{\exists} \,\& \\ \text{not red}_{\exists} \end{array} \right) \wedge \\ \neg[\text{all flags are } \left(\begin{array}{c} \text{green}_{\exists} \,\& \\ \text{not white}_{\exists} \,\& \\ \text{not red}_{\exists} \end{array} \right)] \end{array} \right.$$

≈ some but not all flags are entirely green

d. ⟦If the flag is [Exh_{ALT} green], Aisha will sing⟧
= 1 iff Aisha will sing if the flag is [green$_{\exists}$ & not white$_{\exists}$ & not red$_{\exists}$]
≈ Aisha will sing if the flag is entirely green

The 'controlled' nature of Exh with cohyponyms involves Exh being both *required* by the cohyponym and *necessarily syntactically local* to it. This locality requirement is more stringent than just requiring Exh to be in the same clause as its associate, as observed in examples (6a–c).

There are apparent counterexamples to the tenacity of cohyponymic exclusivity, to which we now turn. In section 6.2.1, I suggest that the data with additives and conjunction should not be taken as counterexamples at all to the generalization that cohyponyms require an ultra-local Exh—although we do learn from conjunction that at least some small distance between Exh and its controller must be possible. Then, in section 6.2.2, I turn to sentential negation (and to a lesser degree other negative DE environments), where Exh appears to be entirely absent. I argue that we can avoid viewing this as an exception to Exh's controlledness if we substitute Exh for the presuppositional exhaustivity operator Pexh of Bassi et al. (2021).

6.2.1 Additives and conjunction

The first area where we apparently do not observe cohyponymic exclusivity, central to the argumentation in this thesis, is in the presence of additive particles and conjunction. However, both of these are still compatible with the presence of an Exh operator that is fairly close to the cohyponym, and certainly require an Exh operator somewhere. Indeed, it would not be right to claim that cohyponymic exclusivity 'disappears' in the presence of additives or conjunction, because the asserted predicates are still incompatible with other non-asserted predicates. For example, both sentences in (7) are incompatible with the flag having a red part.

(7) a. The white flag is also green.
b. The flag is white and green.

As such, cohyponymic exclusivity is not entirely removed by *also* or *and* (in contrast to sentential negation; see section 6.2.2). In (7), they allow *white* and *green* to be mutually compatible, while remaining incompatible with other cohyponyms like *red*.

This empirical domain does teach us one important thing about the controlledness of Exh. The conjunction data add a desideratum to theories of controlled exhaustivity, namely that at least some short syntactic distance must be possible between Exh and its controller.

We now take additive particles and conjunction in turn.

Additive particles

For additives like *also*, it is true that cohyponymic exclusivity is weakened so that no contradiction is intuited in pseudo-repetitions (8). But I argued in chapter 3 that Exh is still present in such sentences: *also* simply results in the domain of each Exh operator being restricted, thus avoiding the cohyponymic exclusivity between the asserted cohyponyms.

(8) The white flag is #(also) green.

One potential concern is the status of Exh-X2 (in (8), the Exh associated with *green*). Recall from chapter 3 that Aravind & Hackl (2017) suggest that the relationship between *also* and Exh-X2 is syntactic, rather than being based in domain-restriction. Specifically, *also* must scope beneath Exh-X2, thus adding an entailment to its prejacent. I showed in chapter 3 that *also* scopes no lower than *v*P. Thus, in (8), Exh-X2 would have to be quite far from *green*, on this theory:

(9) $[_{vP}$ Exh-X2$_{ALT}$ $[_{vP}$ *also* $[_{vP}$ the $[_{AP}$ Exh-X1$_{ALT}$ white] flag is green]]].

If correct, it would be necessary for the theory of controlled exhaustivity to capture the non-locality of Exh-X2 with *green* in (9), in contrast to the locality of Exh-X1. We would have to make sure that the theory we build would allow Exh-X2 to scope at *v*P. However, this is not actually a requirement for our theory. Indeed, I suggested in chapter 3 that there is no reason for the relationship between *also* and Exh-X2 not to be based in domain restriction just like the relationship between *also* and Exh-X1. Thus, Exh-X2 in (8) may very well be ultra-local as well:

(10) [The [Exh$_{\{\text{white, } \sout{\text{green}}, \text{ red}\}}$ white] flag] is also [Exh$_{\{\sout{\text{white}}, \text{ green, red}\}}$ green].

I also argued that this is in fact empirically preferable, due to the example in (11):

(11) The partially white flag is #(also) green.

(11) requires *also*, but the contradiction without the additive cannot be due to *white* (which is weakened by *partially*), so it must come from *green*. If Exh-X2 (the Exh associated with *green*) scoped at *v*P, assuming that the subject can be interpreted in Spec-*v*P, Exh-X2's prejacent would entail the whiteness of the flag and therefore not strengthen *green* to exclude *white*. As such, *also* would not be required to avoid a contradiction.

In light of this, I assume that the relationship between *also* and Exh-X2 is based in domain-restriction, not syntactic scope. The data with *also* are therefore perfectly compatible with the suggestion that Exh is always both present with and ultra-local to cohyponyms, and does not place any additional desiderata on the theory of controlled exhaustivity.

Conjunction: some action at a small distance

Cohyponymic exclusivity also weakens with conjunction.

(12) The flag is white and green.

As discussed in chapter 5, one way to understand the non-exclusivity of the asserted cohyponyms in (12) is that they are exhaustified in one fell swoop by an Exh operator scoping immediately above the entire conjunction:

(13) The flag is $[_{AP}$ Exh$_{ALT}$ $[_{AP}$ white and green]].

While Exh can be characterized as obligatory and at least somewhat local, (13) introduces an important desideratum for theories of controlled exhaustivity: the Exh associated with cohyponyms must be able to occur some small distance from them.

There are several points to make about this proposed account of conjunction, which I will take in turn in the rest of this subsection. First, Exh with conjunction is still subject to some locality

constraint. Second, the analysis in (13) is preferable to an alternative analysis with two local Exh operators whose domains are restricted by *and* in the same way that *also* restricts domains. Third, the analysis can be extended to disjunction: *or* can be taken to interact with Exh in the same way as *and*. I take these points in turn.

The first point is that Exh is still subject to some locality constraint. The fact that it is not absolutely as local as possible to each cohyponym does not mean that it can occur at any distance from them. Consider (14a), for example, whose meaning can be captured with an Exh operator scoping below *every* but above the phrase headed by *and* (14b), but not with an Exh so unconstrained as to be able to be situated at the root level (14c).

(14) a. Every flag is white and green.

b. $[\![\text{Every flag is } [\text{Exh}_{\text{ALT}} \text{ [white and green]}]]\!]$

$$= 1 \text{ iff every flag is } \begin{pmatrix} \text{white}_{\exists} \text{ and green}_{\exists} \wedge \\ \neg[\text{red}_{\exists} \text{ and green}_{\exists}] \wedge \\ \neg[\text{white}_{\exists} \text{ and red}_{\exists}] \wedge \\ \neg[\text{red}_{\exists} \text{ and red}_{\exists}] \wedge \\ \neg[\text{red}_{\exists}] \end{pmatrix}$$

$\approx$ every flag is white, green, and no other colour

c. $[\![\text{Exh}_{\text{ALT}} \text{ [every flag is white and green]}]\!]$

$$= 1 \text{ iff } \begin{cases} \text{every flag is white}_{\exists} \text{ and green}_{\exists} \wedge \\ \neg[\text{every flag is white}_{\exists} \text{ and red}_{\exists}] \wedge \\ \neg[\text{every flag is red}_{\exists} \text{ and green}_{\exists}] \wedge \\ \neg[\text{every flag is red}_{\exists} \text{ and red}_{\exists}] \wedge \\ \neg[\text{every flag is red}_{\exists}] \end{cases}$$

$\approx$ every flag is partly white, every flag is partly green, and for all other colours, it is not the case that every flag is partly of that colour

In particular, the truth conditions in (14c) do not rule out the presence of other colours on some flags, as would be required; they only require that other colours not be present on *every* flag. As such, the conjuction data motivate an obligatory local Exh; locality just needs to be defined loosely enough to allow for (13).[3]

The second point to make about (13) is to highlight that this proposal is hypothesizing that the non-exclusivity of conjoined cohyponyms is due to the syntactic position of Exh; it scopes above the entire conjunction, and therefore does not strengthen either conjunct irrespective of the other. As some readers will have noticed, this is quite different from my account of how additive particles weaken cohyponymic exclusivity. Indeed, I argued that *also* 'de-exhaustifies' cohyponyms by pruning the domain of the ultra-local Exh operators. (15a) has the LF in (15b), repeated from (10).

[3]Note that this argument parallels argumentation by Sauerland (2012), reviewed in chapter 1, that the theory of grammatical exhaustivity cannot be replaced by a mix of lexical ambiguity and global pragmatic strengthening. The strengthening we need for (14a) must take scope at an *intermediate* locus (14b) between *and* and the root. This provides an argument against the possibility of attempting to account for the on-and-off exclusivity of cohyponyms through lexical ambiguity (e.g., *comedy* as ambiguous between bona-fide comedies and the union of comedies and tragicomedies).

(15) a. The white flag is also green.
b. [The [Exh$_{\{\text{white, }\sout{\text{green}}\text{, red}\}}$ white] flag] is also [Exh$_{\{\sout{\text{white}}\text{, green, red}\}}$ green].

Why, then, could something similar not be said for conjunction? The claim would be that each conjunct has its own Exh operator, but *and* prunes the domain of each Exh, as in (16). This would mean that our theory of controlled exhaustivity would not need to give any syntactic leeway to Exh: it could be claimed that there is always a controlled Exh immediately above every cohyponym.

(16) The flag is [[Exh$_{\{\text{white, }\sout{\text{green}}\text{, red}\}}$ white] and [Exh$_{\{\sout{\text{white}}\text{, green, red}\}}$ green]].

No contradiction is derived by the LF in (16), without needing to let Exh scope at some distance from cohyponyms.

The issue for this alternative proposal is that it overgenerates; *and* is not capable of the sort of long-distance de-exhaustification that *also* is capable of, as would be expected if it could prune the alternatives of Exh. Indeed, *and* can weaken cohyponymic exclusivity if it conjoins two predicates 'directly,' but not if it conjoins two clauses containing them:

(17) a. The flag is white and green.
b. #The flag is white and it is green.

(18) A: The flag is white.
B: AND green.
B′: #AND it's green.

This contrasts with *also*:

(19) a. The white flag is also green.
b. The flag is white and it's also green.

(20) A: The flag is white.
B: It's also green.

As such, I maintain that despite the similar meanings of *and* and *also* (both being conjunctive, and more specifically *Boolean* conjunctives, in the examples we are looking at—see chapter 5), *and* differs from *also* in that it does not prune alternatives. Instead, it makes cohyponyms mutually compatible by virtue of creating a syntactic constituent containing multiple cohyponyms that Exh can scope above without violating its locality requirement. Of course, this raises again the question of what about *also* gives it the power to prune alternatives (see the tentative remarks in chapter 3).

The last point I want to make about the proposal (13) is how and whether such an LF can also account for the meaning of *disjoined* cohyponyms. Presumably the proposal in (13) should hold across-the-board for coordination, whether conjunction or disjunction. As such, consider (21).

(21) The flag is white or green.

On at least the most salient reading of (21), the speaker is considering that the flag might be entirely white or entirely green. This is not a problem for my analysis. First, recall that the placement of Exh in (13) is the only available way to have Exh local to the colour terms without it creating a contradiction. But in principle, there is no need to rule out the syntax in (22): since (22) is contradictory, speakers would simply choose to interpret the sentence (12) as in (13).

(22) The flag is [[Exh$_{\text{ALT}}$ white] and [Exh$_{\text{ALT}}$ green]].

For disjunction, then, all that is truly needed to obtain the exclusivity of both *or* and the cohyponyms is to allow an ultra-local Exh operator in each disjunct, as in (23).

(23) ⟦The flag is [[Exh$_{\text{ALT}}$ white] or [Exh$_{\text{ALT}}$ green]]⟧ = 1 iff
the flag is [white$_{\exists}$ & not green$_{\exists}$ & not red$_{\exists}$] or [green$_{\exists}$ & not white$_{\exists}$ & not red$_{\exists}$].

This derives the exclusivity of *or* without even needing to negate an alternative with *and*, because the disjuncts are made incompatible through their own ultra-local exhaustification. A second (and I think more principled) way to obtain the judgment for (21) is to posit an ultra-local Exh operator on each disjunct, together with a free, global Exh operator associating with *or*, as in (24). Recall from chapter 2 that *or* is exhaustified freely, rather than controlling Exh; for example, its plain inclusive meaning is observable in DE environments, unless it is focused.

(24) a. Exh$_{\text{ALT-3}}$ [the flag is [[Exh$_{\text{ALT-2}}$ white] or [Exh$_{\text{ALT-1}}$ green]]].

b. (i) ALT-1 = ALT-2 = {white, green, red}

(ii) $\text{ALT-3} = \left\{ \begin{array}{l} \text{the flag is [[Exh}_{\text{ALT-2}}\text{ white] or [Exh}_{\text{ALT-1}}\text{ green]]}, \\ \text{the flag is [[Exh}_{\text{ALT-2}}\text{ white] and [Exh}_{\text{ALT-1}}\text{ green]]} \end{array} \right\}$

c. $\llbracket (24a) \rrbracket = 1$ iff $\left\{ \begin{array}{l} \text{the flag is only white or only green} \wedge \\ \neg[\text{the flag is only white and only green}] \end{array} \right.$

The right meaning is also obtained if there was not an Exh on each disjunct but a single Exh above the entire disjunction, because the exclusiveness of *or* rules out the possibility of the flag being both partly white and partly green. This is crucial because, if the theory of Exh's locality allows a single Exh to scope above two cohyponyms coordinated with *and*, it will presumably also allow it to scope above two cohyponyms coordinated by *or*.

(25) a. Exh$_{\text{ALT-2}}$ [the flag is [Exh$_{\text{ALT-1}}$ [white or green]]]

b. (i) ALT-1 = {white or green, white or red, red or green, white, green, red}

(ii) $\text{ALT-2} = \left\{ \begin{array}{c} \text{the flag is [Exh}_{\text{ALT-1}}\text{ [white or green]]}, \\ \text{the flag is [Exh}_{\text{ALT-1}}\text{ [white and green]]}, \\ \text{the flag is [Exh}_{\text{ALT-1}}\text{ [white]]}, \\ \text{the flag is [Exh}_{\text{ALT-1}}\text{ [green]]} \end{array} \right\}$

c. ⟦Exh$_{\text{ALT-1}}$ [white or green]⟧ = $\lambda x.x$ is white$_{\exists}$ or green$_{\exists}$ and not red$_{\exists}$.

d. $\llbracket (25a) \rrbracket = 1$ iff $\left\{ \begin{array}{l} \text{the flag is white}_{\exists}\text{ or green}_{\exists}\text{ and not red}_{\exists} \wedge \\ \neg[\text{the flag is white}_{\exists}\text{ and green}_{\exists}\text{ and not red}_{\exists}] \end{array} \right.$

I have given both Exh$_{\text{ALT-1}}$ and Exh$_{\text{ALT-2}}$ alternatives involving syntactic pruning (i.e., single disjuncts as alternatives), as predicted by Katzir (2007). In addition to pruning, Exh$_{\text{ALT-1}}$ has alternatives obtained by replacing predicates with their cohyponyms, while Exh$_{\text{ALT-2}}$ has an alternative generated by replacing *or* with *and*. Rather few alternatives are innocently excludable: only *red* for Exh$_{\text{ALT-1}}$ (alternatives of the form 'white or red' are no more excludable than alternatives of the form 'white'),[4] and only the conjunctive alternative for Exh$_{\text{ALT-2}}$. Exh$_{\text{ALT-1}}$ contributes the meaning that the flag does not have any colours other than white or green; Exh$_{\text{ALT-2}}$ contributes the meaning

[4] Of course, if I had written out all the basic colour terms, the idea is that all colours other than *green* or *white* are excluded by Exh$_{\text{ALT-1}}$.

that the flag is not both white and green. As such, the only options are for the flag to be only green or only white, as desired.

This shows that the proposal for conjunction in (13) carries over to disjunction as well. The meaning of disjunctions can be obtained in several ways by having a global Exh associated with *or* and one or two more or less local Exh operators associated with the cohyponyms. Disjunctions are expected to have a single parse available, where each cohyponym is strong and the disjunction is exclusive.

In sum, while the additivity paradigm does not affect the generalization that cohyponyms control Exh in any way, the conjunction paradigm does: it makes it necessary for controlledness to be based on a notion of locality that is flexible enough to handle the small distance between Exh and the cohyponyms in (26), repeated from (13).

(26) The flag is $[_{\text{AP}}\ \text{Exh}_{\text{ALT}}\ [_{\text{AP}}\ \text{white and green}]]$.

6.2.2 Negation and other negative DE contexts

In this section, I discuss DE environments where cohyponymic exclusivity is not maintained; as noted in chapter 2, this is a small minority of DE environments. I argue that, by adopting the Pexh operator of Bassi et al. (2021), the data can be captured without any additional desiderata for theories on the distribution of Exh. On this view, the apparently exceptional DE contexts are not counterexamples to the claim that Exh is always present with and local to cohyponyms.

Three classes of DE environments

Following the discussion in chapter 2, DE contexts can be classed into three categories: environments which, like the antecedent in (5d), maintain cohyponymic exclusivity (27a); sentential negation, where cohyponymic exclusivity is neither maintained nor only weakened, but in fact entirely removed (27b);[5] and a small number of other negative contexts where I have not identified a clear judgment one way or the other (27c).

(27) a. Every green flag was made in China.
⇒ ‘green’ ≈ ‘all green’
b. This flag is not green.
⇒ ‘green’ $\not\approx$ ‘all green’
c. Aisha doubts that the flag is green.
⇒ ‘green’ ≈ ??

I will not take a stance on the third class (27c), and instead start by focusing exclusively on (27a) and (27b). My hope is that (27c) can essentially be collapsed in one or the other category, pending clearer judgments. Given that some DE environments pattern like UE environments is displaying cohyponymic exclusivity, while others (specifically negation) do not, the obvious conclusion is

[5] I noted in chapter 2 that some cohyponymic exclusivity seems to be maintained under sentential negation with predicates like *comedy* or *fork* where there are portmanteaus (*tragicomedy* and *spork*) denoting areas of conceptual overlap. This apparent ‘exception to the exception’ can presumably be dealt with based on the existence of the portmanteau lexeme.

that, unlike traditional exhaustivity effects, whether or not cohyponymic exclusivity is observed depends on polarity rather than the distinction between UE and DE environments.

How does the empirical complication associated with sentential negation (27b) and possibly other negative contexts (27c) fit with my claim that Exh is obligatory and obligatorily ultra-local with cohyponyms? As shown in chapter 2, putting an Exh operator either above or below negation yields bad results for cohyponyms; we see this for the integrative predicate *train* in (28b) (imagine for the sake of argument a sci-fi world where there are train-planes, train-boats, and so on, so that the 'exclusivity' of *train* is not due to world knowledge). In contrast, we obtain the right results if there is simply no Exh whatsoever in the LF (28a). For simplicity of presentation, I am ignoring the controlled exhaustification of *vehicle* (the reader can pretend the sentence is just 'this is a train') and using a not very local Exh for *train* in (28b-ii); nothing hinges on this for this example.

(28) a. $[\![\text{not [this vehicle is a train]}]\!] = 1$ iff $\neg\mathsf{train}(v)$.

b. (i) $[\![\text{Exh}_{\text{ALT}}\ \text{[not [this vehicle is a train]]}]\!] = 1$ iff
$\neg\mathsf{train}(v) \wedge \neg\neg\mathsf{plane}(v) \wedge \neg\neg\mathsf{bicycle}(v)$
$\equiv \neg\mathsf{train}(v) \wedge \mathsf{plane}(v) \wedge \mathsf{bicycle}(v)$.
$\approx$ this vehicle is an exemplar of all vehicles except trains

(ii) $[\![\text{not [Exh}_{\text{ALT}}\ \text{[this vehicle is a train]]}]\!] = 1$ iff
$\neg[\mathsf{train}(v) \wedge \neg\mathsf{plane}(v) \wedge \neg\mathsf{bicycle}(v)]$.
$\approx$ it is not the case that this vehicle is only an exemplar of trains

The same goes for summative predicates. In (29) and the rest of this chapter, I abbreviate formulae of the form '$\exists y[y \sqsubseteq x \wedge (\text{colour})(x)]$' as '$(\text{colour})_{\exists}(x)$.'

(29) a. $[\![\text{not [the flag is green]}]\!] = 1$ iff $\neg\mathsf{green}_{\exists}(f)$.

b. (i) $[\![\text{Exh}_{\text{ALT}}\ \text{[not [the flag is green]]}]\!] = 1$ iff
$\neg\mathsf{green}_{\exists}(f) \wedge \neg\neg\mathsf{red}_{\exists}(f) \wedge \neg\neg\mathsf{white}_{\exists}(f)$.
$\equiv \neg\mathsf{green}_{\exists}(f) \wedge \mathsf{red}_{\exists}(f) \wedge \mathsf{white}_{\exists}(f)$.
$\approx$ the flag has parts of all colours except green

(ii) $[\![\text{not [Exh}_{\text{ALT}}\ \text{[the flag is green]]}]\!] = 1$ iff
$\neg[\mathsf{green}_{\exists}(f) \wedge \neg\mathsf{red}_{\exists}(f) \wedge \neg\mathsf{white}_{\exists}(f)]$.
$\approx$ the flag is not only green

As it turns out, the negation data can be put aside entirely, if we adopt the 'presuppositional' Exh of Bassi et al. (2021).[6] That is, with the right definition of Exh, the data are in fact fully compatible with the claim that Exh is always obligatory and ultra-local with cohyponyms. As such, before building a theory of Exh's controlledness, our first substantial task in this chapter is to see how the proposal by Bassi et al. (2021) makes it possible to put aside the negation data as fully compatible with the presence of an ultra-local Exh.

A presuppositional exhaustivity operator

Bassi et al. (2021) propose a theory of exhaustivity which would capture the apparent lack of exhaustivity below negation not due to negation being a DE operator, but because (*i*) negation is an operator reversing truth and falsity conditions, and (*ii*) presuppositions are independently known

[6] I thank Keny Chatain for pointing this out to me.

to project past negation. On these authors' proposal, it is possible to capture that negation does not pattern with other DE contexts, as desired.

They propose that Exh does not *assert* the exclusion of alternatives, as standardly assumed (e.g., Chierchia et al. 2012), but rather *presupposes* it. Exh is thus the mirror image of *only*. This presuppositional 'Pexh' operator is defined as in (30), where $Excl(\phi)$ is the set of *innocently excludable* alternatives of ϕ, and R is the set of *relevant* alternatives (*grosso modo*; see their (1c)).

(30) $$[\![\text{Pexh}(\phi)]\!] = \begin{cases} 1, \text{if } [\![\phi]\!] = 1 \wedge \bigwedge([\![\psi]\!] = 0) : \psi \in Excl(\phi) \wedge [\![\psi]\!] \in R \\ 0, \text{if } [\![\phi]\!] = 0 \\ \#, \text{otherwise} \end{cases}$$

(Bassi et al. 2021:9)

(30) defines Pexh as altering truth conditions but not falsity conditions. According to Strong Kleene trivalent logic (e.g., George 2008; Fox 2013; Winter 2019), the presupposition of a sentence ϕ is "the disjunction of its truth and falsity conditions," i.e. the conditions in which ϕ has a truth value (Bassi et al. 2021:9). As such, Pexh(ϕ) asserts rather than presupposing ϕ: it is true in the truth conditions but false in the falsity conditions, so it is not presupposed. But—limiting our attention for now to excludable *stronger* alternatives—the alternatives are necessarily false in either the truth or falsity conditions of Pexh(ϕ) according to (30). In the truth conditions, they must be false by the definition of Pexh. In the falsity conditions, they must be false because they are stronger than ϕ, and ϕ is false by the definition of Pexh. Therefore, if Pexh(ϕ) has a truth value, the alternatives of ϕ are false. This falsity is present in both the truth and falsity conditions, and therefore presupposed.[7]

Bassi et al. (2021:10) illustrate this in action as follows:

(31) Aisha found some of her marbles.

a. $$[\![\text{Pexh}((31))]\!] = \begin{cases} 1, \text{if Aisha found some but not all of her marbles} \\ 0, \text{if Aisha found none of her marbles} \\ \#, \text{otherwise} \end{cases}$$

b. Predicted presupposition:
Aisha found none of her marbles $\vee$ Aisha found some but not all of her marbles
$\equiv$ Aisha didn't find all her marbles.

What is crucial for us is that on the definition of Pexh in (30), while ϕ and Pexh(ϕ) differ in *truth* conditions (as long as ϕ has relevant excludable alternatives), they do not differ in their *falsity* conditions. $[\![\text{Pexh}(\phi)]\!] = 0$ if $[\![\phi]\!] = 0$.

As noted by Bassi et al. (2021), what is counter-intuitive about this proposal is the view that a sentence like (31) *presupposes* that Aisha didn't find all of her marbles, rather than asserting it. Clearly, this is not a presupposition that must be entailed by the common ground upon the utterance of the sentence. For Bassi et al. (2021), the presuppositions triggered by Pexh are globally accommodated: as long as they are *consistent with* the common ground (i.e., not necessarily *entailed* by it), listeners adjust the common ground for it to entail the presupposed content.[8] Bassi et al. (2021) point out that such global accommodation is not rare in language, as can be observed in the prejacent of *only* (32a) and the presupposition of possessives (32b):

[7] We will return below to alternatives that are neither stronger nor weaker than Pexh's prejacent. In their case, the presupposition that is predicted is conditional: if ϕ is true, then its alternatives are false.

[8] See these authors' paper for discussion, including the 'Hey, wait a minute!' test for presuppositions.

(32) a. SCENARIO*: Aisha doesn't know whether Carrie or anyone other than Ben was at the party.*
A: Who went to the party?
B: Only Carrie was there!
b. SCENARIO*: Aisha and Ben both believe that Aisha doesn't know whether Ben has a sister.*
A: Why are you leaving?
B: I have to pick up my sister at the airport.
(Bassi et al. 2021:34)

Pexh does not affect truth conditions under negation

The fact that Pexh does not affect falsity conditions makes it possible to make Pexh fail to affect *truth conditions* in certain environments, including at least sentential negation (Bassi et al. 2021:§3). Assume a standard trivalent semantics for sentential negation:

(33) $$[\![\text{not}(\phi)]\!] = \begin{cases} 1, \text{if } [\![\phi]\!] = 0 \\ 0, \text{if } [\![\phi]\!] = 1 \\ \#, \text{otherwise} \end{cases}$$

As noted above, $[\![\text{Pexh}(\phi)]\!] = 0$ if $[\![\phi]\!] = 0$. As such, adding Pexh between negation and ϕ, while affecting the falsity conditions of the sentence, will not affect its truth conditions:

(34) $$[\![\text{not}(\text{Pexh}(\phi))]\!] = \begin{cases} 1, \text{if } [\![\text{Pexh}(\phi)]\!] = 0 \equiv \text{if } [\![\phi]\!] = 0 \\ 0, \text{if } [\![\text{Pexh}(\phi)]\!] = 1 \\ \#, \text{otherwise} \end{cases}$$

Of course, given that Pexh does have an impact on the falsity conditions of 'not(Pexh(ϕ)),' it does yield a presupposition, thereby affecting the meaning of the sentence in some sense. Like all presuppositions, this presupposition projects past negation. But at least in the case of scalar alternative-triggering expressions, the meaning of the presupposition is in fact entailed by the assertion, making it undetectable and unproblematic. To see this, consider an example with the scalar element *or* (cf. Bassi et al. 2021:14). Under negation, there is no exclusion of an alternative with *and*:

(35) Aisha didn't talk to Ben or Carrie.
$\not\rightsquigarrow \neg$[Aisha didn't talk to Ben and Carrie] $\equiv$ Aisha talked to Ben and Carrie.

With Pexh, nothing goes wrong if we claim that (35) has a Pexh operator below negation:

(36) not [Pexh [Aisha talked to Ben or Carrie]]

(36) obtains the following meaning (where (37b) simply reverses the truth and falsity conditions of (37a)):

(37) a. $$[\![\text{Pexh [A talked to B or C]}]\!] = \begin{cases} 1, \text{if A talked to B or C} \wedge \neg[\text{A talked to B and C}] \\ 0, \text{if } \neg[\text{A talked to B or C}] \\ \#, \text{otherwise} \end{cases}$$

b. $[\![\text{not}[\text{Pexh}[\text{A talked to B or C}]]]\!] = \begin{cases} 1, \text{if } \neg[\text{A talked to B or C}] \\ 0, \text{if A talked to B or C} \wedge \neg[\text{A talked to B and C}] \\ \#, \text{otherwise} \end{cases}$

c. Presupposition of (37b):
$\neg[\text{A talked to B or C}] \vee [\text{A talked to B or C} \wedge \neg[\text{A talked to B and C}]]$
$\equiv \neg[\text{A talked to B and C}]$

The presupposition in (37c) projects past negation, yielding a global inference which is not only globally accommodated, but is also in fact entailed by the assertion. The assertion is that Aisha talked to neither B nor C, while the presupposition is that she did not talk to both. As such, it is "not detected *qua* presupposition" (Bassi et al. 2021:14); adding Pexh below negation has no empirical effect.

Applying Pexh to cohyponyms

We have just observed that replacing Exh with Pexh makes it possible to claim that there is a Pexh operator below sentential negation in sentences with weak scalar alternative-triggering expressions like *or* or *some*. What about expressions that trigger alternatives which are not ordered by entailment? Consider the answer in (38).

(38) A: Who sang? (Bassi et al. 2021:15, fn. 20)
B: Aisha$_F$ sang.

In this case, the presupposition is a conditional: if Aisha sang, Ben and Carrie did not sing.

(39) a. $[\![\text{Pexh } [\text{Aisha}_F \text{ sang}]]\!] = \begin{cases} 1, \text{if Aisha sang} \wedge \neg[\text{Ben sang}] \wedge \neg[\text{Carrie sang}] \\ 0, \text{if } \neg[\text{Aisha sang}] \\ \#, \text{otherwise} \end{cases}$

b. Presupposition of (39a):
$[\text{Aisha sang} \wedge \neg[\text{Ben sang}] \wedge \neg[\text{Carrie sang}]] \vee \neg[\text{Aisha sang}]$
$\equiv \text{Aisha sang} \rightarrow \neg[\text{Ben sang} \vee \text{Carrie sang}]$

This conditional presupposition is globally accommodated.[9]

Now let's apply this theory to cohyponyms. Cohyponyms are different from other expressions triggering non-stronger alternatives because there is a claim in my chapter 2 that they impose a constraint on Exh that other expressions do not. So far, I have essentially had to stipulate that cohyponyms are not exhaustified under negation, as a counterexample to their otherwise obligatory and ultra-local exhaustification. But if we change Exh to Pexh in our theory, there is no longer any need to stipulate the lack of exhaustivity operator in negative sentences. Let's start with the following LF:[10]

(40) not [Pexh$_{\text{ALT}}$ [this vehicle is a train]]

The meaning of (40) is the following:

[9] As noted by Bassi et al. (2021:15, fn. 20), the presupposition is in fact entailed by the strong assertion that only Aisha sang, because the conditional is true if Aisha did not sing.

[10] I put aside the 'ultra-locality' of Exh/Pexh for simplicity of presentation in this section, but an ultra-local Pexh taking only the noun *train* as its complement produces identical results.

(41) a. $[\![\text{Pexh}_{\text{ALT}}\ [\text{This vehicle is a train}]]\!] = \begin{cases} 1, \text{if } \mathsf{train}(v) \wedge \neg\mathsf{plane}(v) \wedge \neg\mathsf{bicycle}(v) \\ 0, \text{if } \neg\mathsf{train}(v) \\ \#, \text{otherwise} \end{cases}$

b. $[\![\text{not } [\text{Pexh}_{\text{ALT}}\ [\text{This vehicle is a train}]]]\!] = \begin{cases} 1, \text{if } \neg\mathsf{train}(v) \\ 0, \text{if } \mathsf{train}(v) \wedge \neg\mathsf{plane}(v) \wedge \neg\mathsf{bicycle}(v) \\ \#, \text{otherwise} \end{cases}$

c. Presupposition of (41b):
$\neg\mathsf{train}(v) \vee [\mathsf{train}(v) \wedge \neg\mathsf{plane}(v) \wedge \neg\mathsf{bicycle}(v)]$
$\equiv \mathsf{train}(v) \rightarrow \neg[\mathsf{plane}(v) \vee \mathsf{bicycle}(v)]$

This is a perfectly acceptable result. The sentence (40) presupposes that the vehicle is either not a train at all or a train which does not meet the qualifications for being an exemplar of other vehicles, and it asserts that in fact, it is not a train at all. This leaves completely open what type of vehicle the subject is, as long as it is not an exemplar of trains. While this presupposition is not necessary to capture the intuited meaning, it does not create a problem. There is no need to entertain the possibility of an LF with Pexh *above* negation, since I posit on independent grounds that Pexh must always be local anyway.

I have just adopted Pexh as an easy way to capture the sentential-negation data with cohyponyms without needing to derive any exceptions to cohyponyms' control of exhaustivity. This fends off overgeneration threats from both Exh $>$ *not* and *not* $>$ Exh. In the next subsections, I will make sure that this does not cause any unwanted problems. I will deal with *doubt*-type DE environments (environments which are negative but do not involve sentential negation), check that nothing goes wrong with non-negative (*if*-type) DE environments, and discuss how Pexh fits with the fact (mentioned in chapter 2) that cohyponymic exclusivity *is* possible under *not* with focus intonation.

Returning to *doubt*-type DE environments

I noted above that some negative-flavoured DE environments lead to unclear judgments about cohyponymic exclusivity; (42) is repeated from (27c).

(42) Aisha doubts that the flag is green.
$\Rightarrow$ 'green' $\approx$??

Of course, one constant difficulty with such judgments is to determine what is the 'basic' meaning and what is due to pragmatic effects. In particular, it could be that *green* is 'really' universal under *doubt*, and that to the extent that existential meanings are possible, this is a non-maximality effect (see chapter 5). Let's assume that there is only one real semantic meaning, which in the case of colour terms is either existential or universal. With Pexh, either possibility can be captured with an ultra-local exhaustivity operator, depending on the lexical entry for *doubt*. If *doubt* refers to the *falsity conditions* of its complement (43a), no strengthening is predicted to be observable in the complement (*green* would be existential); if *doubt* refers to *truth* conditions (43b), then strengthening is predicted to be observable.

(43) a. $[\![\text{doubt}]\!] = \lambda p.\lambda x.x$ thinks that it is probably the case that $p = 0$. (option 1)
$\Rightarrow$ Pexh has no observable effect under *doubt*

b. $[\![\text{doubt}]\!] = \lambda p.\lambda x.x$ thinks that it is not likely the case that $p = 1$. (option 2)
⇒ Pexh has an observable effect under *doubt*

Whichever is the right definition, non-controlled exhaustivity effects like the exclusive meaning of *or* are still expected not to show up under *doubt*, not because of whether it refers to truth or falsity conditions but because nothing would force a Pexh operator below *doubt* in the first place for such alternative-triggering expressions.

Pexh with non-negative DE operators

As I discussed in chapter 2, most DE environments, including the antecedents of conditionals and the restrictor of universals, do display cohyponymic exclusivity. (44) is repeated from (27a).

(44) Every green flag was made in China.
⇒ 'green' ≈ 'all green'

This has the syntax in (45), with exhaustivity computed below *every* (i.e., Pexh is adjointed to AP or NP). (45) is simplified in part due to not showing the exhaustification of *flag*.

(45) $[_{vP}\ [_{QP}\ \text{every}\ [_{NP}\ [_{AP}\ \text{Pexh}_{ALT}\ \text{green}]\ [_{NP}\ \text{flag}]]]_1\ [_{vP}\ \text{was made}\ t_1\ \text{in China}]]$.

If *every* is defined as in (46), then no operator here makes reference to falsity conditions.[11]

(46) $[\![\text{every}]\!] = \lambda P.\lambda Q.\forall x[P(x) \rightarrow Q(x)]$.

Thus, changing Exh to Pexh makes no difference.

Pexh with focused cohyponyms under sentential negation

I noted in chapter 2 that cohyponymic exclusivity can surface under sentential negation when the cohyponym is focused:

(47) a. #The flag of Ireland is not green, it's green, white, and orange.
b. The flag of Ireland is not GREEN, it's GREEN, WHITE, AND ORANGE.

This could appear to be a problem for the Pexh analysis. (47b) could have been obtained by placing an Exh operator below the negation, but a Pexh operator below negation is vacuous, as we have just discussed at length.

Bassi et al. (2021:§3.2) deal with this exact problem for the free exhaustification of *some*:

(48) Aisha didn't talk to Ben OR Carrie … she talked to BOTH! (Bassi et al. 2021:15)

They suggest to deal with this as an instance of presupposition cancellation or local accommodation, akin to (49):

[11] If *every* was defined as in (i), then an issue would arise, of course:

(i) $[\![\text{every}]\!] = \lambda P.\lambda Q.\forall x[\neg P(x) \vee Q(x)]$.

(49) Aisha didn't STOP smoking … she never used to! (Bassi et al. 2021:15)

They use the $\mathscr{A}$ operator of Beaver & Krahmer (2001) to model the local accommodation of presuppositions; this operator ensures that conditions which would be undefined in a trivalent semantics result in falsity instead. Bassi et al. (2021) therefore argue that non-vacuous exhaustivity under negation is due to the presence of $\mathscr{A}$:

(50) not [$\mathscr{A}$ [Pexh$_{\text{ALT}}$ [Aisha talked to Ben or Carrie]]]

a. $[\![\text{Pexh}_{\text{ALT}}\ [\text{A talked to B or C}]]\!]$
$= \begin{cases} 1, \text{if } [\mathsf{talk}(a,b) \vee \mathsf{talk}(a,c)] \wedge \neg[\mathsf{talk}(a,b) \wedge \mathsf{talk}(a,c)] \\ 0, \text{if } \neg[\mathsf{talk}(a,b) \vee \mathsf{talk}(a,c)] \\ \#, \text{otherwise} \end{cases}$

b. $[\![\mathscr{A}\ [\text{Pexh}_{\text{ALT}}\ [\text{A talked to B or C}]]]\!]$
$= \begin{cases} 1, \text{if } [\mathsf{talk}(a,b) \vee \mathsf{talk}(a,c)] \wedge \neg[\mathsf{talk}(a,b) \wedge \mathsf{talk}(a,c)] \\ 0, \text{if } \neg[\mathsf{talk}(a,b) \vee \mathsf{talk}(a,c)] \vee \text{otherwise} \end{cases}$
$\equiv \begin{cases} 1, \text{if } [\mathsf{talk}(a,b) \vee \mathsf{talk}(a,c)] \wedge \neg[\mathsf{talk}(a,b) \wedge \mathsf{talk}(a,c)] \\ 0, \text{otherwise} \end{cases}$

c. $[\![\text{not } [\mathscr{A}\ [\text{Pexh}_{\text{ALT}}\ [\text{Aisha talked to Ben or Carrie}]]]]\!]$
$= \begin{cases} 1, \text{otherwise} \\ 0, \text{if } [\mathsf{talk}(a,b) \vee \mathsf{talk}(a,c)] \wedge \neg[\mathsf{talk}(a,b) \wedge \mathsf{talk}(a,c)] \end{cases}$

(50c) is true as long as Aisha did not talk to Ben and not Carrie, and she did not talk to Carrie and not Ben—i.e., she talked to neither or both.

The same principle can explain (47b).

6.2.3 Interim summary

I began this section with data, repeated from chapter 2, suggesting that the Exh operator leading to cohyponymic exclusivity is non-optional and always ultra-local to the cohyponym. I then turned to two types of data that could be taken as counterexamples to this. The first was the interaction between cohyponymic exclusivity and additive particles and conjunction. If additive particles serve to restrict the domain of Exh rather than removing Exh or necessarily scoping below Exh, the additive data do not serve as a counterexample to Exh's controlledness, and in fact teach us nothing about it. As for conjunctions of cohyponyms, they too can be understood as necessarily requiring Exh, although they *do* teach us something about controlledness, viz. Exh must have the ability to scope at some distance from cohyponyms (a point we return to immediately in section 6.3).

The second type of data was sentential negation. While problematic meanings are obtained if sentential negation were to co-occur with the standard Exh operator of Chierchia et al. (2012), this apparent obstacle to the claim that cohyponyms always require a local Exh can be overcome by claiming that Exh is in fact presuppositional, as argued by Bassi et al. (2021) for independent reasons. I therefore adopt these authors' Pexh operator. This allows for a simpler generalization about Exh's behaviour with cohyponyms. Note that, after this section, I will return to writing 'Exh' rather than 'Pexh.'

The use of Pexh for controlled exhaustification is more than a passive adoption of Bassi et al.'s (2021) theory; it also constitutes an argument in favour of it, because controlled exhaustification is an area of grammar where negation and non-negative DE environments pattern differently. This

can be straightforwardly handled by Pexh, but remains a mystery with Exh. The different behaviour of negation and other DE operators is a new kind of argument in favour of Pexh, adding to the argumentation provided by Bassi et al. (2021); these authors discussion the distinction between positive and negative sentences but do not compare negative sentences to other DE contexts—sensibly, given that it is not necessary to make this distinction for non-controlled exhaustification.

In sum, as long as locality is understood loosely enough to allow for a single Pexh to scope above a pair of conjoined cohyponyms, we can make the strong general claim that there is always a Pexh operator present ultra-locally to cohyponyms. Even when there is no empirical evidence favouring this (as is the case for the data from sentential negation), a simple generalization is better than a complicated one. We now turn to characterizing (section 6.3) and deriving (section 6.4) this locality constraint.

6.3 The locality requirement on Exh

So far, we have observed that Exh with cohyponyms is both obligatory and subject to a locality requirement. In this section, we turn to understanding what the nature of this locality requirement it. How local to the alternative-triggering expression does a controlled Exh operator have to be, and by what measure? The proposal I defend in this section is that the locality requirement is that Exh must be within its controller's XP—a syntactic rather than semantic constraint. Looking ahead, I will model what in the syntax causes such a locality requirement in section 6.4.

Let's return to the four environments in which we can observe the locality requirement on Exh. (51) is repeated from (6).

(51) a. $[\![$The [Exh_{ALT} green] flag is [Exh_{ALT} white]$]\!]$
= 1 iff the [$\text{green}_{\exists}$ & not $\text{white}_{\exists}$ & not $\text{red}_{\exists}$] flag is [$\text{white}_{\exists}$ & not $\text{green}_{\exists}$ & not $\text{red}_{\exists}$]
$\Rightarrow$ contradiction

b. $[\![$Every flag is [Exh_{ALT} green]$]\!]$
= 1 iff every flag is [$\text{green}_{\exists}$ & not $\text{white}_{\exists}$ & not $\text{red}_{\exists}$]
$\approx$ every flag is entirely green

c. $[\![\text{Exh}_{\text{ALT}}$ [some flags are [Exh_{ALT} green]$]\!]$
$$= 1 \text{ iff } \left\{ \begin{array}{l} \text{some flags are } \left(\begin{array}{c} \text{green}_{\exists} \,\& \\ \text{not white}_{\exists} \,\& \\ \text{not red}_{\exists} \end{array} \right) \wedge \\ \neg[\text{all flags are } \left(\begin{array}{c} \text{green}_{\exists} \,\& \\ \text{not white}_{\exists} \,\& \\ \text{not red}_{\exists} \end{array} \right)] \end{array} \right.$$
$\approx$ some but not all flags are entirely green

d. $[\![$If the flag is [Exh_{ALT} green], Aisha will sing$]\!]$
= 1 iff Aisha will sing if the flag is [$\text{green}_{\exists}$ & not $\text{white}_{\exists}$ & not $\text{red}_{\exists}$]
$\approx$ if the flag is entirely green, Aisha will sing

As should be obvious, my goal in this section is to come up with a *single* generalization of Exh's locality that predicts *all* the data, rather than coming up with ad-hoc stipulations example-by-example. To obtain the right generalization, let me also add another pair of examples to the examples in (51), namely the minimal pair in (52) with attributive adjectives. The observation is that

two colour terms can modify a noun if they are conjoined, but not if they are stacked one above the other.

(52) a. The white and green flag is torn up.
b. #The white green flag is torn up.

As argued in section 6.2.1, given that (52a) is non-contradictory and that *and* does not have the ability to prune Exh's alternatives, there must be a single Exh above both colour terms, as in the LF in (53a).[12] Moreover, given that (52b) is contradictory, each colour term must have its own Exh, as in (53b).

(53) a. The $[_{AP}$ Exh_{ALT} $[_{AP}$ white and green$]]$ flag is torn up.
b. #The $[_{AP}$ Exh_{ALT} white$]$ $[_{AP}$ Exh_{ALT} green$]$ flag is torn up.

We now attempt some different locality requirements for Exh: a semantic requirement (section 6.3.1), a syntactic requirement based in Chomsky's (2000; 2001) notion of phases (section 6.3.2), and a syntactic requirement based around the construct of maximal projections (section 6.3.3). Finally, I briefly discuss productive compounding in some central European languages in section 6.3.4, showing that the data is compatible with the syntactic constraint I adopt (which is not to say that they are incompatible with other competing hypotheses).

6.3.1 A semantic locality requirement

Perhaps the most intuitively appealing kind of locality constraint on Exh would be semantic in nature; after all, Exh is an operator whose existence is posited for strictly semantic reasons. A semantic locality requirement on Exh would presumably be as follows:

(54) **A semantic account of controlled exhaustivity:**
No scope-bearing operator can intervene between Exh and its controller.

This immediately explains all four examples in (51). In (51b), Exh must scope below *every*; in (51c), it must scope under *some*; and in (51d), it must scope below *if*. As for (51a), repeated as (55), the Exh associated with *green* must scope below *the*, creating the contradiction.

(55) #The green flag is white.

This makes (54) a promising hypothesis for the data in (51). But what about our minimal pair in (52)? This is where the semantic hypothesis goes wrong. First, (55) makes the prediction that (52a) would be contradictory. In this example, we need Exh to scope over the whole conjunction; but that would mean that the scope-bearing expression *and* intervenes between Exh and the second conjunct. What is more, (54) might also predict that (52b) should be acceptable, as in (56), where

[12] An alternative LF with a global Exh (i) must be ruled out.

(i) Exh_{ALT} [the white and green flag is torn up].

This LF would involve the exclusion of alternatives like 'The blue flag is torn up,' making inferences about altogether different flags.

white and *green* are under a single Exh operator and no scope-bearing element intervenes between Exh and the two controllers.

(56) The $[_{NP}$ Exh_{ALT} $[_{NP}$ $[_{AP}$ white$]$ $[_{NP}$ $[_{AP}$ green$]$ $[_{N}$ flag$]]]]$ is torn up.

No contradiction is predicted on the syntax in (56), contrary to the intuited meaning. To be sure, (54) could be amended to clarify that each controller must have its own Exh, in which case (52b) would indeed be contradictory. Either way, we have already identified a problem for (52a).

Requirements on the syntactic locus of Exh are usually framed in semantic terms; in particular, Exh is dispreferred in downward-entailing environments (e.g., Chierchia et al. 2012, Fox & Spector 2018). This is quite sensible given that Exh is an operator with semantic effects. At the same time, given that we are dealing with the grammaticalized view of exhaustivity rather than a neo-Gricean approach to strengthening, it is not obvious that semantic preferences or requirements should be the only kind of requirement on Exh. After all, if Exh is an operator that is present in the syntax, there is no *a priori* reason that it could not be subject to syntactic requirements like other operators. I now turn to two ways to model the locality requirement on Exh syntactically.

6.3.2 A phase-syntactic locality requirement

The first syntactic locality requirement that we will consider would force a fair amount of locality on Exh with cohyponyms, while also giving it enough flexibility to scope above entire conjunction phrases, unlike the above semantic requirement. The idea is to use phase theory (Chomsky 2001, 2008, a.o.) to consider the possibility that controlled Exh must be in the same syntactic domain as its controller.[13] In phase-theoretic syntax, some syntactic constituents—in particular *v*P and CP (Chomsky 2001) and DP—are postulated to constitute 'phases' in syntactic derivation. Phases involve particular heads (*v*, C, and D) 'spelling-out' (i.e., sending to PF) their complements once all their associated operations have been performed. This is meant to account, among other things, for successive-cyclicity in movement. Expressions must pass through the specifier of the phase-head to avoid being sent off to PF before reaching their final destination (e.g., a root Spec-CP for *wh*-expressions). Terminologically, *v* is a phase head; it spells-out the VP, which is the 'spell-out domain'; the 'phase' is the entire *v*P. Likewise with C (a phase head) and TP (its spell-out domain) and CP (the phase), and D/NP/DP.

The naive idea behind using phase theory is that phases are small enough syntactic constituents that they could require some locality between Exh and its controller, without being so strict as to create problems with conjunction data. Note that we start by assuming that AP is not a phase, because if it is one, claiming for colour adjectives that Exh must be in their *phase* would be the same as claiming that Exh must be in their *maximal projection*, which is the hypothesis we will turn to after this one.

The first phase-based hypothesis we consider is the following:

(57) **A first phase-syntactic account of controlled exhaustivity:**
Exh must be in the same phase as the controller.

[13] See Erlewine 2017 for another kind of locality requirement on overt *only* in some languages making use of phase theory. The goal and outcome of Erlewine's discussion is fundamentally different; his aim is not to make *only* necessarily appear in the same phase as its associate.

(57) amounts to claiming that Exh must be anywhere in the lowest DP, *v*P, or CP dominating the controller. A possible way to think about (57) is to rely on Chomsky's notion of lexical arrays. On Chomsky's view, the lexical items from which a sentence is generated are first taken from the lexicon and put into a workspace. But they are not all stored in a single undivided set; rather, they are organized in subsets of lexical items, with each subset corresponding to the lexical items to be merged in the same phase. If (57) is correct, we could stipulate that when building lexical arrays, language ensures that there is an Exh operator for each cohyponym in the array.

Let's go through the data in (51) before turning to the minimal pair in (52) again. (51a) is predicted to be a contradiction only due to the meaning of *green*: its Exh must be within the DP *the green flag*. The Exh associated with *white* could scope above *green* if *the green flag* is interpreted low in Spec-*v*P, so that it would not cause a contradiction.[14] (57) might also produce the right result for the conditional in (51d), but it depends on whether unaccusative verbs (like *to be*) head *v*Ps that are phases (Legate 2002, 2003) or not (Chomsky 2001). If unaccusative *v*Ps are phases, the farthest away from *green* that Exh can be located is at *v*P within the antecedent (*v* is a phase head, so Exh must occur within that *v*P); this derives that Exh will exclude colours other than green, as desired:

(58) $[_{CP}$ If $[_{TP}$ T $[_{vP}$ Exh_{ALT} $[_{vP}$ $[_{DP}$ the flag] $[_{vP}$ is green]]]]], Aisha will sing.

On the other hand, if unaccusative *v*Ps are not phases, Exh will be able to as high as being adjoined to CP:

(59) $[_{CP}$ Exh_{ALT} $[_{CP}$ if $[_{TP}$ T $[_{vP}$ $[_{DP}$ the flag] $[_{vP}$ is green]]]]], Aisha will sing.

Exh in (59) only has the antecedent as its prejacent, but it has the entire antecedent, including *if*. This is might be undesirable; I put the matter aside.

On the hypothesis in (57), (51b) and (51c) are more straightfowardly problematic. These sentences involve quantifiers; assuming that quantifier raising (QR) involves adjunction to *v*P (as done in, e.g., Hirsch 2017) and no higher, there is nothing in (57) preventing the Exh associated with the cohyponym from scoping above the quantifier, rather than ultra-locally to the cohyponym:

(60) $[_{vP}$ Exh_{ALT} $[_{vP}$ $[_{QP}$ every flag] $[_{vP}$ is green]]]

As such, we do not obtain the locality constraint for (51b) and (51c) on the phasal hypothesis in (57).

What about the minimal pair in (61), repeated from (52)?

(61) a. The white and green flag is torn up.
b. #The white green flag is torn up.

On the one hand, the conjoined colour terms in (61a) are correctly predicted to be compatible. The smallest phase the colour terms are in is the DP, so these cohyponyms only require the presence of an Exh operator within that DP. The Exh operator can scope high enough to have both colour terms

[14] See chapter 2 for evidence that this is in fact not right; in fact, each adjective needs to be strengthened so as to exclude the other.

in its prejacent; this ensures that neither will contradict the other.[15] On the other hand, (61b) is predicted to be consistent rather than contradictory. Here too, the smallest phase the cohyponyms are in is the DP, so the Exh operator(s) could scope high enough to have both cohyponyms in its scope. This is shown in (62), repeated from (56).

(62) The $[_{NP}$ Exh_{ALT} $[_{NP}$ $[_{AP}$ white] $[_{NP}$ $[_{AP}$ green] $[_{N}$ flag]]]] is torn up.

As such, the hypothesis in (57) overgenerates by predicting (61b) to be consistent.

Most of the problems just identified for the hypothesis in (57) could be circumvented by modifying the phase account of locality to claim that Exh must occur in the same spell-out domain as the expression-triggering expression—that is, *below* the phase head:[16]

(63) **A second phase-syntactic account of controlled exhaustivity:**
Exh must be in the same spell-out domain as the controller.

This would force Exh to occur below v^0 and therefore below the QP *every flag* in (60), and below the C^0 *if* regardless of the phase-status of unaccusative *v* in (51d)/(58)/(59), for example. But this change would still fail to capture the difference between the two sentences in (52)/(61): in either case, Exh would be forced to scope below D^0, but possibly above both colour terms as in (62).

All this discussion assumes that AP is not a phase. On the other hand, if AP *is* a phase, the results of the hypothesis in (57) are the same as if we simply claimed that controllers require an Exh in their maximal projection. This is the hypothesis we turn to now.

6.3.3 A phrase-syntactic locality requirement

We turn to the final locality requirement:

(64) **A phrase-syntactic account of controlled exhaustivity:**
Exh must be in the maximal projection of the controller.

This immediately obtains all four examples in (51); as the reader can see in those examples, while I did not label the phrases, I put the Exh operators in places that are compatible with them being adjoined to the colour term's AP.

What about the minimal pair with conjoined and stacked colour terms in (52)/(61)? The incompatibility of the stacked colour terms in (61b) follows immediately from (64), because each is in its own AP:

(65) a. ⟦the $[Exh_{ALT}$ white] $[Exh_{ALT}$ green] flag⟧
≈ the [$white_{\exists}$ & not $green_{\exists}$ & not $red_{\exists}$] [$green_{\exists}$ & not $white_{\exists}$ & not $red_{\exists}$] flag
b. *the $[_{NP}$ Exh_{ALT} $[_{NP}$ white green flag]]

As for (61a), I suggest that it can also be understood to follow from (64) with the right assumptions about the syntax of conjunction. A classic analysis of conjoined phrases is that the entire phrase

[15] It is immaterial whether there is one Exh operator for each cohyponym, or a single Exh for both: either way, the single Exh operator or the two Exh operators can scope high enough to have both colour terms in the prejacent, thus ensuring that neither is so strengthened as to exclude the other.

[16] It's not clear at all what would lead to such a constraint, at least on the T-model of linguistic architecture proposed by Chomsky (2001), where spell-out is a purely morphophonological process.

inherits the category of its conjuncts (see e.g. Bresnan 1976, Williams 1978, Gazdar 1981, Sag et al. 1985, McCloskey 1986, Burton & Grimshaw 1992, McNally 1992, Chaves 2012, Williams 2015:49): conjoining two DPs yields a DP, conjoining two APs yields an AP, and so on:

(66)

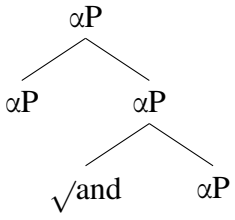

This is an appealing hypothesis for two reasons. First, if the purpose of stipulating grammatical categories is to capture facts about which types of expressions have the same paradigmatic distribution in sentences, it would be odd to state that conjoined phrases do not have the same category as their conjuncts. For instance, DPs can occur in subject position, and the phrase resulting from conjoining two DPs can as well:

(67) a. [$_{DP}$ The boy] ate an apple.
b. [$_{\alpha P}$ [$_{DP}$ The boy] and [$_{DP}$ his teacher]] ate an apple.

This follows immediately if the phrase labelled as αP in (67b) is in fact itself a DP. The second appeal to this theory is that two constituents of different categories can generally not be conjoined:

(68) a. *The scene in the movie and that I wrote was in Chicago.
(cf. Sag et al. 1985:118, who cite Chomsky 1957)
b. Aisha walked slowly (*and) to the store.

(68a) attempts and fails to conjoin a PP and a relative CP, and (68b) an AdvP and a PP. Focusing on (68a), this can be understood as a failure in labelling the phrase *in the movie and that I wrote*, whose conjuncts do not share a category and which therefore cannot be labelled as either a PP or a CP.

To be sure, it is more common nowadays to view conjunctions as having their own category; the phrase is a 'ConjP' or '&P' (e.g., Munn 1993, Johannessen 1996, 1998, Camacho 2003, Zhang 2010). The insight is intended to be that *and* is a head (it is a simplex lexical item taking a complement and specifier—the conjuncts), and phrases inherit categories from heads rather than from the heads' complements or specifiers.[17] As such, for these reasons internal to X′-theory, the claim is that *and* has a category Conj, and this is what projects for the entire conjunction phrase:

[17]While Chomsky's (2013; 2015) recent work on labelling complicates this picture, it still stands for Chomsky that a head-complement structure projects the label of the head. However, it is not clear what the label of the entire conjunction phrase would be on his analysis. The daughters of the highest node are the left conjunct XP and the ConjP containing *and* and its complement; neither of these are heads, so they must share a feature in order for labelling to proceed.

(69)

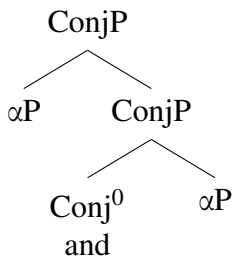

This proposal requires ignoring what I understand as the main purpose of stipulating the existence of categories (and indeed labels), which is to capture that certain expressions pattern together syntactically to the exclusion of others. On the proposal in (69), the simplest prediction is that all ConjPs would pattern the same way (rather than conjoined DPs patterning like DPs and conjoined APs like APs), and that conjoined phrases would behave differently from their individuals conjuncts (that is, a conjunction of DPs, being a ConjP rather than a DP, would not pattern like a non-conjoined DP). One could brush aside these concerns by claiming that the syntax can 'see into' the ConjP, effectively ignoring its status as a ConjP, so that the syntax is instead affected by the category of its complement and specifier. But this raises the question of what the purpose of stipulating the Conj category is supposed to be other than to circumvent the X′-theory–internal problem of category projection. More problematically, it raises the question of why other labels cannot be ignored in this way. For instance, a DP with an AP inside of it never patterns like an AP; it always patterns like a DP. Why would ConjP be different? These issues do not arise on the view that the category of conjoined phrases is inherited from the conjuncts.

What is more, even among authors who suggest a ConjP label for conjoined phrases, it is not uncommon to posit inheritance of other features from conjuncts to the entire phrase (e.g., Kiss 2012, Franks & Willer-Gold 2014). If so, it is not a big step to claim that the category of the conjuncts is inherited as well.

To be sure, it has been noted that not all conjunctions involve conjuncts of the same category. Sag et al. (1985:117–118) note the following examples where they claim that constituents of different categories are conjoined (the right-hand parentheticals are slightly modernized from Sag et al.'s, but are not meant to represent my own view of the category of the conjunctions):

(70)
a. Aisha is a Republican and proud of it. (DP and AP)
b. Aisha is healthy and of sound mind. (AP and PP)
c. That was a rude remark and in very bad taste. (DP and PP)
d. I am hoping to get an invitation and optimistic about my chances. (VP and AP)
e. I am both expecting to get the job and of the opinion that it is a desirable one. (VP and PP)
f. Aisha was awarded the Golden Fleece Award and very upset about it. (VP and AP)

This is apparently problematic for my proposed syntax (66): if the two conjuncts have a different category, the entire phrase cannot simultaneously inherit the categories of both. The easiest answer is to claim that the examples in (70) are all conjoined *v*Ps with elided material. If so, the examples

in (70) do not pose a problem to the generalization that conjunction must be of like categories.[18]

Summing up, I assume that conjoined phrases inherit the category of each conjunct, as in (71), repeated from (66). This could occur because *and* is a root and lacks a category, for example. The labelling algorithm would therefore copy the labels of the phrases rather than *and* (see Chomsky 1995a on roots' inability to project labels).

(71)

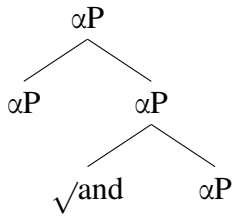

Let's return to our data pertaining to the exhaustification of cohyponyms, and in particular the minimal pair in (72), repeated from (52) and (61).

(72) a. The white and green flag is torn up.
b. #The white green flag is torn up.

What the syntax of conjunction proposed in (71) means is that we can capture the difference between (72a) and (72b) from the observation that the two colour terms are in a single AP in (72a), but in two separate APs in (72b). On the hypothesis that cohyponyms require an Exh inside their XP (64), we therefore predict that each colour term is exhaustified by itself in (72b) (as pointed out above), but that (72a) would have the option of there being a single Exh above both colour terms. Indeed, for conjunctions, Exh has some flexibility, because each colour term only needs to be exhaustified by some Exh operator somewhere in their maximal projection. On the syntax of conjunction just proposed, the maximal projection is the entire conjoined phrase. As such, the syntax in (73) is available, as needed.

(73) ⟦The flag is $[_{\text{AP}}$ Exh_{ALT} $[_{\text{AP}}$ green and white$]]$⟧
= 1 iff the flag is [green$_{\exists}$ & white$_{\exists}$ & not red$_{\exists}$].

Of course, (74) is syntactically available too. This is fine, as long as the non-contradictory (73) is available as well: speakers will choose the non-contradictory parse.

(74) ⟦The flag is $[_{\text{AP}}$ Exh_{ALT} green$]$ and $[_{\text{AP}}$ Exh_{ALT} white$]$⟧
= 1 iff the flag is [green$_{\exists}$ & not white$_{\exists}$ & not red$_{\exists}$] and [white$_{\exists}$ & not green$_{\exists}$ & not red$_{\exists}$].
⇒ contradiction

Accepting the hypothesis in (64) means that, by extrapolation, even in a simple sentence like (75a), where no difference in meaning is predicted to be observable based on whether Exh is local or global, Exh must in fact be local to the cohyponym (75b).

[18] Another approach would be to accept that the examples in (70) have different categories, but to point out that all that is really needed for my account of controlled Exh to work is the *possibility* of conjunction phrases inheriting the category of their conjuncts (66). It may be that conjunction phrases inherit as much matching information from their conjuncts as possible; if both conjuncts share a category, the conjunction phrase inherits it, but if they do not, the phrase remains acategorial. This wouldn't cause any problems for the theory of controlled exhaustivity, because all phrases that control Exh have (as far as I can tell) expressions of the same category as their alternatives.

(75) a. The flag is green.
b. ⟦The flag is [$_{AP}$ Exh$_{ALT}$ green]⟧
= 1 iff the flag is [green$_{\exists}$ & not white$_{\exists}$ & not red$_{\exists}$ & . . .]

Also note that this hypothesis interacts slightly with different possibilities about the syntax of copular sentences (see chapter 1). If copular sentences start out with the subject inside the predicate's XP (76), then pseudo-repetitions on this hypothesis are only contradictory due to the meaning of X1.

(76)

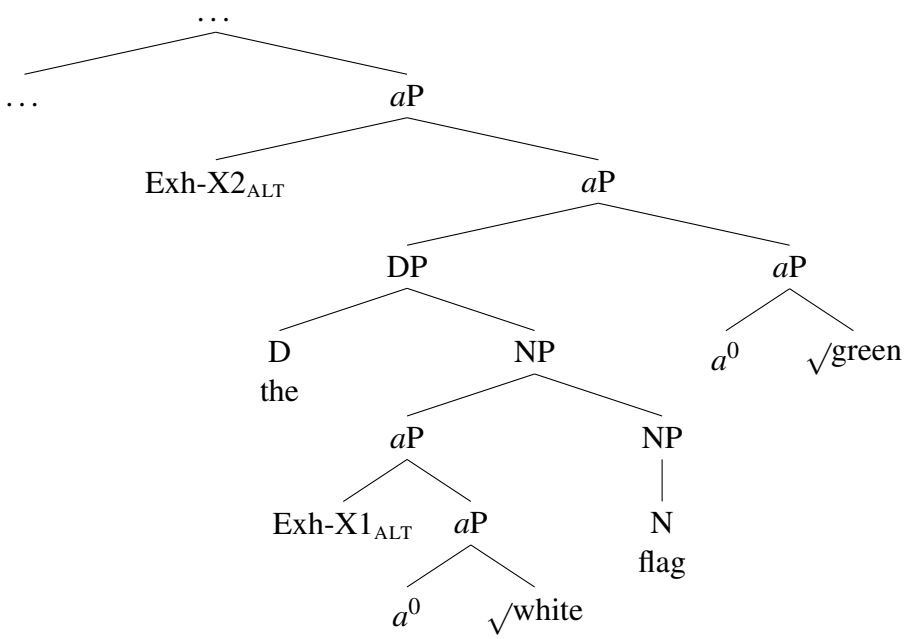

In (76), Exh-X2 does not create a contradiction because its prejacent entails both colour terms; the contradiction is only due to Exh-X1. See chapter 2 for the suggestion that this is not right; both Exh operators must strengthen their associated cohyponym without taking the other into account. As such, I must claim that subjects in copular sentences do not start at as low as the XP of the predicate.

In summary, my claim in this section is that the locality requirement on Exh is that an Exh operator must occur within the XP headed by the cohyponym. In section 6.4, I will seek to derive this constraint. Before then, let me point out that my locality requirement is also compatible with a slightly broader empirical outlook.

6.3.4 Compound colour terms

Some central European languages (at least German and Polish; I will focus on the former) productively obtain non-exclusive meanings for summative cohyponyms through compounding rather than conjunction. This fits straightforwardly with the generalization that cohyponyms allow their controlled Exh anywhere in their XP. My purpose in this section is simply to present the data and point out that it aligns with my generalization.

Compounding instead of conjunction is specifically found with summative predicates—colour terms and materials. In predicative position, both compounding and conjunction are available,

although conjunction is marked.[19]

(77) a. Die Flagge ist grün-weiß.
the flag is green-white
'The flag is green and white.'
b. Die Flagge ist grün und weiß.
the flag is green and white
'The flag is green and white.'

On the other hand, in attributive position, conjunction is not only marked but in fact very degraded:

(78) a. Die grün-weiße Flagge ist zerrissen.
the green-white flag is torn
'The green and white flag is torn.'
b. ??Die grüne und weiße Flagge ist zerrissen.
the green and white flag is torn
'The green and white flag is torn.'

Evidence that these colour adjectives are truly compounded (rather than stacked) comes from the fact that they must share a single agreement suffix at the right edge of the compound:

(79) die weiß-grün.e Flagge
the white-green.AGR flag
'the white and green flag'

Attempting to have each adjective agree with the head noun is not morphosyntactically impossible but leads to semantic oddness suggesting that such adjectives are in fact stacked:

(80) #die weiß.e grün.e Flagge
the white.AGR green.AGR flag
[comment: "feels much like '#the white green flag'"]

This sort of compounding is fully productive and indeed the default way to 'conjoin' colour terms.[20]

These data are completely aligned with the view that Exh must be within the XP of a cohyponym. Compounding is the merger of two heads to form a complex head. In (79), *weiß* and *grün* are both a^0 elements merging to form another a^0 (which then agrees with the head noun).

[19] I thank Bernhard Schwarz and Viola Schmitt for the German data.

[20] Spanish is another language with some compounding, although the process is less productive (Luis Alonso-Ovalle, p.c.).

(i) l-a bander-a roj-i-blanc-a
the-FEM flag-FEM red-COMPOUND-white-FEM
'the red and white flag'

The compounding is not only less productive but also rarely used even in the case of possible compounds. A Google search (May 8, 2022) for the conjoined phrase `"la bandera roja y blanca"` provides 44,900 results, compared to only 4,080 for `"la bandera rojiblanca"`.

(81)

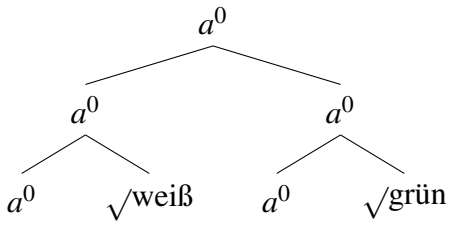

These can therefore be exhaustified in one fell swoop by a single Exh (82); this meets both cohyponyms' requirement that there be an Exh within their *a*P.

(82)

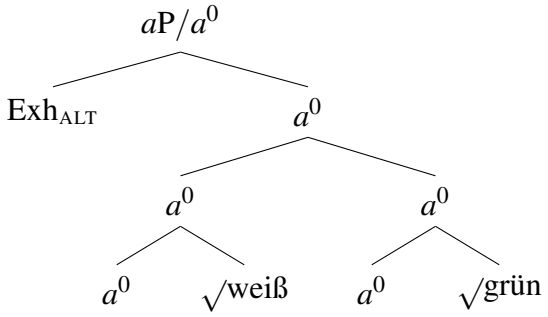

6.4 Roots, category, and the conceptual system

We have just defined the locality requirement on controlled Exh as follows: controllers (cohyponyms) require an Exh operator in their maximal projection. Having established this generalization, I now turn to asking where this locality requirement comes from. I suggest that Exh's obligatoriness and locality with cohyponyms is due to derivational morphemes requiring Exh, specifically because they must Agree with it.[21]

I lay out this hypothesis in section 6.4.1, focusing on cohyponyms and in particular adjectives. Then, in section 6.4.2, I make some comments about expressions that are strengthened without controlling Exh, and expressions which control Exh without being cohyponymic.

6.4.1 Derivational morphemes Agree with Exh

Why is there obligatorily an Exh operator in cohyponyms' XP? I suggest a path forward based on the observation that there is another grammatical process, from the domain of morphology, which has the same twin properties as controlled Exh, being both obligatory and always occurring locally to the element it 'associates' with (selects). This is derivational morphology: functional morphemes (n^0, v^0, a^0, etc.) that assign grammatical category to roots (e.g., Marantz 1997, 2001, 2007; Harley & Noyer 1997; Arad 2003, 2005; Borer 2014; Harley 2014). As such, I propose to collapse the obligatoriness and locality of Exh vis-à-vis its controlling predicate with the obligatoriness and locality of derivational morphemes vis-à-vis roots. Exh is controlled due to derivational

[21] A simpler version of this hypothesis would simply state that cohyponyms themselves Agree with Exh, without distinguishing between roots and category morphemes. This would cause some difficulties around feature percolation in the argumentation below, but might be salvageable. Either way, the distinction between roots and derivational morphemes is independently needed in morphosyntax, so there is nothing cumbersome about making use of it here.

morphology: it is not cohyponyms per se, but derivational morphemes that require Exh within their projection.

In addition to being obligatory and local to roots, the idea of using derivational morphology to explain cohyponyms' exclusivity fits nicely with what derivational morphemes and roots are. Indeed, roots are the part of language that "add conceptual meaning to the structures built by syntax" (De Belder & van Craenenbroeck 2015), and they thereby "serve as the interface between [language] and the broader cognitive system" (Harley 2014:248). Derivational morphemes deal directly with linguistic material (roots) that interfaces with (lexicalized) concepts, and serve as the bridge between these concepts and functional or logical elements in language.[22] What the current hypothesis adds to this is the idea that the derivational morphemes also exhaustify the root they select (that is, require that it be exhaustified through the presence of an Exh operator), as in (83) for the adjective *green*:

(83) $[_{a\text{P}}\ \text{Exh}_{\text{ALT}}\ [_{a^0}\ a^0\ \sqrt{\text{green}}]]$

Stepping back, what this means is in effect that derivational morphemes serve to partition the lexical or conceptual domains that the roots they select come from.

Let's see how this can work in practice. Recall that we need to formalize the need for Exh in (83) in a way that provides enough flexibility to also generate (84), where some distance separates Exh from a^0:

(84)

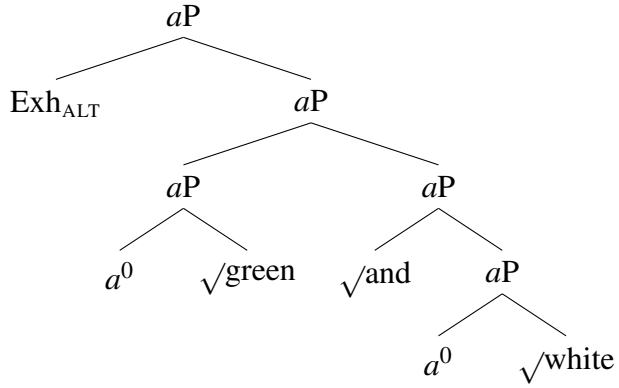

How should the relationship between Exh and derivational morphemes be formalized to allow the flexibility in (84)? In (84), Exh has some distance from the root selected by a^0, and there is a single Exh for two a^0 heads. These facts would make it immediately unappealing to claim that derivational morphemes *select* an Exh operator. If a^0 was selecting Exh, there would be one Exh for each a^0, contrary to fact. On the other hand, if derivational morphemes merely Agree (Chomsky 2001) with Exh, we can get the flexibility we need, as I will show. As such, my claim

[22] Boeckx (2011) provides another interesting proposal giving an important status to derivational morphemes in the interplay between language and thought. Indeed, recent work (Spelke 2000, 2003; Spelke & Kinzler 2006, 2009; Kinzler & Spelke 2007; Ott 2009) has suggested that, rather than accounting for humans' cognitive capacities by claiming that human minds are made up of different building blocks ('core knowledge systems') from other animal minds, what is special about humans is instead that we are able to create connections between these core knowledge systems in ways that other species cannot. Boeckx (2011) proposes to explain the demodularization of concepts in the human mind as the result of roots (concepts) being made mergeable through derivational morphology. While intriguing, a challenge for this idea (at least as stated) is that the roots must have already been mergeable if they are able to be merged with the derivational morpheme.

is that derivational morphemes Agree with Exh; they have an uninterpretable/unvalued [*u*Exh] feature.

I will show how this works for conjunctions shortly below. First, let me note that this is not the first proposal claiming that there is agreement involved in the distribution of Exh; Chierchia (2013) models the obligatory presence of Exh with negative polarity items and other alternative-bearing expressions like *or* through agreement as well.[23] Our accounts differ significantly, however. On my account, Exh is the *target* of agreement. There is a probe on derivational morphemes which must find an Exh to Agree with; since there is no upward probing (as I am about to discuss), Exh is syntactically constrained to being in an XP no higher than the derivational morpheme's. Chierchia has agreement the other way around. For him, the probe is on Exh, and it searches for alternative-triggering expressions (Chierchia 2013:§2.5). This means that Exh is not syntactically constrained, as long as it c-commands the alternative-triggering expression. As such, while we both use the mechanism of Agree, our proposals are mirror images of the other's. I emphasize that this is not necessarily problematic for either of us; we are trying to explain different empirical instances of exhaustification, in particular free vs. controlled exhaustivity effects.

In addition to Chierchia's proposal, the idea of agreement with Exh also has a close cousin in the proposal by Lee (2004) and Hirsch (2017) that there is agreement with *only*. They suggest this due to overt *only* (in Korean and English, respectively) surfacing in the 'wrong' place for the semantics. They therefore suggest an [ONLY] feature allowing the overt instantiation of *only* to merely be the reflex of agreement with a semantically substantial covert *only*. Another similar proposal is due to Horvath (2007), who has an 'exhaustive identification' ([EI]) feature she uses to distinguish between the syntactic behaviour of *only* and *even* in Hungarian.

Let's see how the claim that derivational morphemes Agree with Exh works in practice to capture both the obligatoriness of Exh and its locality to its controller, with a focus on the conjunction data. What we'll need in order to capture the flexibility in (84) is to allow for not only heads to Agree, but also labels (or just labels; see in particular Béjar & Rezac 2009:48–49).[24] That is, the [*u*Exh] feature of a^0, if not dealt with at the merger of a^0, will percolate upwards, such that Agree takes place between an *a*P label and Exh.

To obtain this, we first let the [*u*Exh] feature percolate to the right level. In (85), percolation is shown with a dotted arrow.

[23] I thank Hedde Zeijlstra and Aron Hirsch for pointing me in this direction.

[24] See also Sells 1985; Lieber 1989; Pollard & Sag 1994; Kobele 2005; Tsarfaty & Sima'an 2007; Kiss 2012; Franks & Willer-Gold 2014 on feature percolation, as well as Chomsky's (1995a; 2013; 2015) work on labelling and bare phrase structure.

(85)

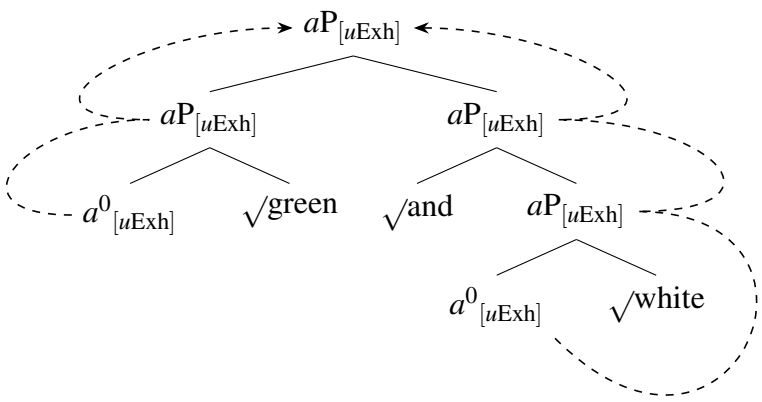

The [*u*Exh] features on both a^0 morphemes are not dealt with immediately upon the merger of these heads, so they percolate upwards as part of the projection of the label. After this occurs, we merge Exh and have the label $a\text{P}_{[u\text{Exh}]}$ Agree with it. In (86), Agree is shown with a solid arrow.[25]

(86)

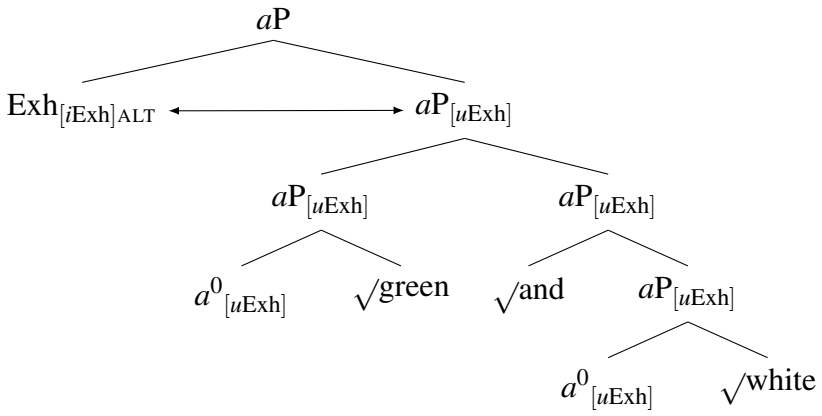

The feature valuation is then passed down to the two a^0 heads.

Given the proposed syntax, there is nothing forcing Exh to be in this particular position within the *a*P; as discussed around example (74), other parses with a lower Exh (one in each conjunct) would be syntactically possible but semantically contradictory, and therefore not chosen. Crucially, the proposal that Exh is required due to [*u*Exh] on derivational morphemes rules out any parse where the only Exh is higher than a cohyponym's XP, given the lack of upward Agree (Chomsky 2000, 2001). The [*u*Exh] feature can percolate up to *a*P and no higher, from where it cannot Agree with an Exh in some higher projection as in (87) (which is starred):

[25] Alternatively, it could be the topmost *a*P label that Agrees downward with Exh (cf. Béjar & Rezac 2009).

(87) *

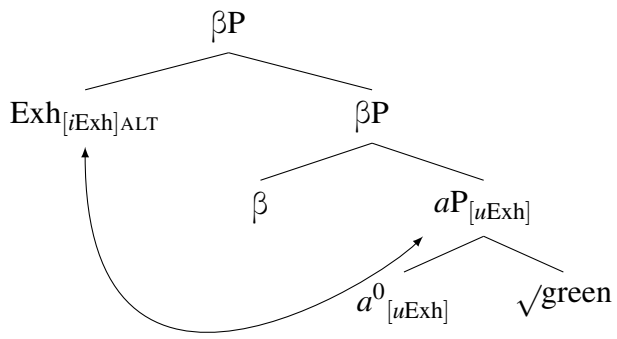

$a\text{P}_{[u\text{Exh}]}$ cannot probe upward for its [*u*Exh] feature to be valued by the [*i*Exh] on an Exh operator in a higher projection.

In summary, Exh is obligatory with and obligatorily local to predicates because derivational morphemes Agree with it. This captures that Exh is obligatory because its absence would create the syntactic problem that a [*u*Exh] probe lacks a goal to Agree with. Moreover, it captures that Exh must be present within the predicate's XP because there is no upward Agree (Chomsky 2001): the goal found by [*u*Exh] must be present within the syntactic domain that [*u*Exh] can percolate to, which corresponds to the maximal projection of the head initially bringing [*u*Exh] into the derivation.

This account rests on several independently motivated assumptions. First, the [*u*Exh] feature, like other features, can percolate from the head X to the labels XP, but no higher. Second, there is no upward Agree; thus, for [*u*Exh] to Agree successfully, there has be an Exh within the XP of the head that introduces [*u*Exh] into the derivation. Third, Agree is a non-fallible operation: it must take place successfully in order for a syntactic derivation to converge (Chomsky 2000, 2001, pace Preminger 2011). Thus, all [*u*Exh] features must be met by the presence of an Exh operator. Fourth and last, features from conjuncts can percolate to the entire conjunction phrase, presumably together with category labels (as discussed in section 6.3).

6.4.2 Extending the proposal beyond simplex adjectives

So far in this chapter, I have only discussed controlled Exh with syntactically simplex predicates whose alternatives are not ordered by entailment; and so far in this section, I have only discussed adjectives. How does the proposal work for other predicates? In this section, I discuss nouns, scalar predicates (predicates that form entailment scales) like *warm* and *hot*, and complex phrases like PPs and VPs. I will show that the results of my hypothesis are mixed. I conclude the section with an advantage of the proposal for data with an overt *only*.

Simplex nominal controllers

What the theory that derivational morphemes Agree with Exh means for nominal controllers is that Exh is (weakly) predicted to necessarily occur below the determiner. On the standard view, phrases like *the dog* or *a dog* are DPs, not NPs/*n*Ps (Abney 1987). If Exh is required due to n^0 bearing a [*u*Exh] feature, then Exh must occur low in the phrase:

(88) a. $[_{\text{DP}}\ \text{a}\ [_{n\text{P}}\ \text{Exh}_{\text{ALT}}\ [_{n\text{P}}\ n^0{}_{[u\text{Exh}]}\ \text{dog}]]]$

b. $[_{DP}$ the $[_{nP}$ Exh_{ALT} $[_{nP}$ $n^0{}_{[uExh]}$ dog$]]]$

This is not horrible on the surface: (88b) is presumably the only option anyway—having Exh above *the* would give it a set of *e*-type individuals as its prejacent, where Exh would presumably not be able to do anything (it's not clear what 'excluding' an individual would mean).

Let's turn to conjunctions. On the theory we are considering, nominal controllers require Exh in their *n*P. This seems right for definites. Definites involving conjoined nominals are not perfectly acceptable, but to the extent that they are, they involve *n*P rather than DP conjunctions (that is, the definite is not repeated in each conjunct):

(89) a. ?The comedy and tragedy was a good play.
b. #*The comedy and the tragedy was a good play.

Indefinites are less promising for the hypothesis, however. Conjunctions involve repetition of the determiner:

(90) This play is both a comedy and a tragedy.

What we have in (90) is a conjunction of two DPs:

(91) $[_{DP}$ $[_{DP}$ a comedy$]$ $[_{DP}$ and $[_{DP}$ a tragedy$]]]$

If nominal controllers necessarily host an Exh within their *n*P, (90) is really as in (92):

(92) $[_{DP}$ $[_{DP}$ a $[_{nP}$ Exh_{ALT} n^0 comedy$]]$ $[_{DP}$ and $[_{DP}$ a $[_{nP}$ Exh_{ALT} n^0 tragedy$]]]]$

This is problematic; a contradiction is incorrectly predicted. One possible avenue is to treat *a* as syntactically lower than *the*; another is to deny the DP hypothesis of Abney (1987) altogether. Indeed, some work claims that nominal phrases are indeed NPs as classically proposed, rather than DPs (e.g., Payne 1993, Bruening 2009, 2020, 2022, Bruening et al. 2018).[26]

An argument against the DP hypothesis from Bruening (2009) comes from selection. When verbs take clauses as complements, they may select for particular formal features on material on clausal heads above the verb (e.g. finiteness, mood) (93); this can be understood as verbs selecting for C and different selectors having different requirements for what kind of C it must be.

(93) a. FINITE VS. NON-FINITE SELECTION: (Bruening 2009:28)
(i) Aisha wants the world to be flat.
(ii) *Aisha wants that the world is flat.
b. SUBJUNCTIVE VS. INDICATIVE SELECTION:
(i) Aisha asked that the answer {be, *is} two.
(ii) Aisha thinks that the answer {*be, is} two.

On the other hand, when verbs take nominals as complements, we never find that they have selectional requirements for any features associated with determiners. In particular, there are no verbs requiring the selection of a definite or indefinite nominal; Bruening exemplifies this via the made-up verb *stread* which would necessarily select for indefiniteness.

[26] I thank Justin Royer for pointing me to Bruening's work on this.

(94) NON-EXISTENT INDEFINITE SELECTION: (Bruening 2009:28)
a. Aisha is streading a book.
b. *Aisha is streading the book.

In section 6.4.1, I suggested that *and* might not project a category due to it being acategorial; there must be a category for projection to proceed. If determiners lack a category too, it follows by the same token that they do not project either. Nominals would then be *n*Ps rather than DPs; if Exh is required within the maximal projection of *n*, it is expected that Exh could occur above the determiner.[27] Note that I will keep labelling nominals as DPs in this thesis.

Scalar predicates

Scalar predicates constitute another area of grammar that my hypothesis makes predictions about. Scalar predicates (predicates that do not take part in an entailment scale) do not control exhaustivity (see chapter 2):

(95) a. Some of the warm plates are hot.
$\Rightarrow$ *warm* $\not\approx$ warm and not hot
b. If the plate is warm, ...
$\Rightarrow$ *warm* $\not\approx$ warm and not hot

Yet, as it stands, my proposal would lead us to expect that these predicates do control Exh: there is a derivational morpheme a^0 for both *warm* and *hot*, so each would be ultra-locally exhaustified. This may be circumventable with the right lexical entries or the right approach to degree semantics, but at present I can merely note that this problem exists.

Syntactically complex controllers

As discussed in chapter 4, some controllers are syntactically complex:

(96) a. The book **about cats** is #(also) **about bicycles**.
b. The train crash that **killed my teacher** #(also) **killed a surgeon**.

(96a) and (96b) might be analyzed as P and *v* (respectively) Agreeing with Exh. The [*u*Exh] feature on P would be a pure stipulation, while the feature on *v* would at least be an expected piece of the generalization that bearing a [*u*Exh] feature is a property of derivational morphemes. On the other hand, a challenge for P and *v* controllers is that they have complex complements; as such, we must find a way to prevent the agreement from taking place with an Exh operator that is too low to create the effect in (96). Recall from chapter 4 that the Exh operators associated with the bolded phrases in (96) come *in addition to* the Exh operators strengthening the cohyponyms they take as complements. (96a) has at least the following Exh operators (numbered for ease of reference):

[27] I commented above that we would not want Exh to occur above *the*; but this could be ruled out on semantic grounds, because Exh cannot only take an individual as its prejacent.

(97)

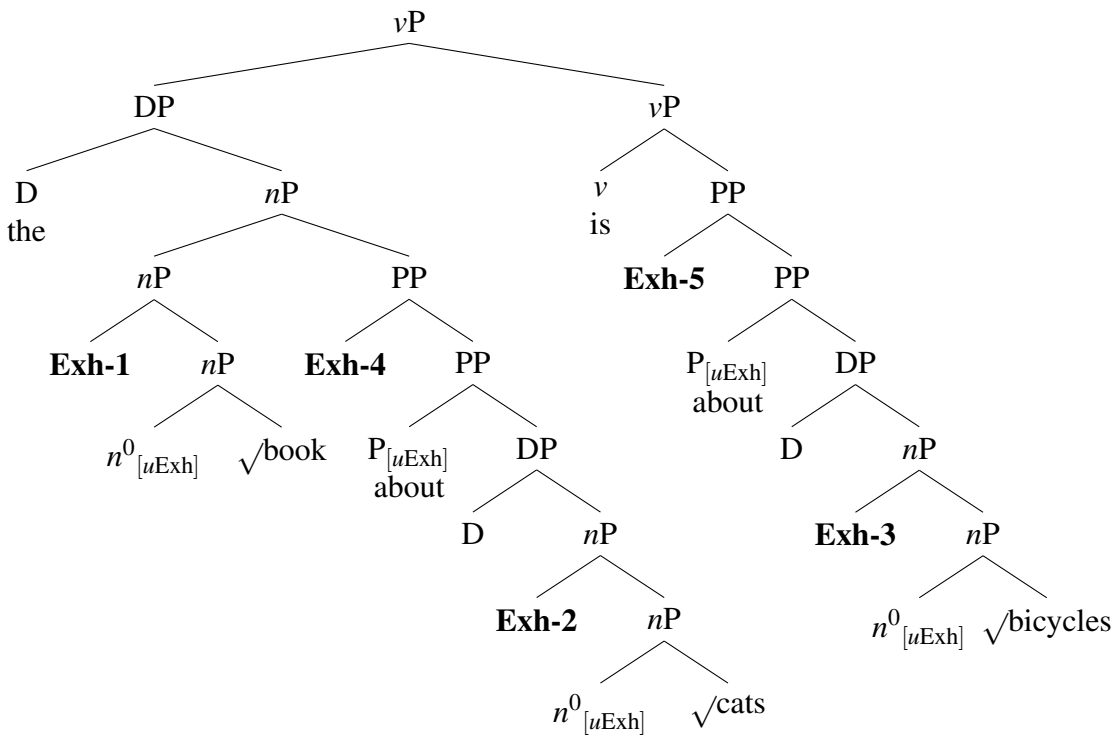

That is, the book is only a book (Exh-1), the cats are only cats (Exh-2), and the bicycles are only bicycles (Exh-3); *and* 'about cats' is the only topic of the book (Exh-4), and so is 'about bicycles' (Exh-5), hence the contradiction in the sentence. The locality of the Exh operators 1–3 is easily understood as emerging from n^0 morphemes Agreeing with them, forcing them to occur within their respective *n*Ps. But can my Agree story account for the fact that Exh operators 4 and 5 are both required and necessarily local to their associate PPs?

Let's focus on the PP *about cats*. Of course, as stated above, we begin by assuming that the P *about* has a [*u*Exh] feature. This feature cannot be satisfied by anything higher than the highest PP label. There *is* an Exh operator, Exh-2 in (97), c-commanded by the P. However, Exh-2 has already been Agreed with by the n^0 nominalizing *cats*, rendering it inactive for agreement by P. As such, P needs its own Exh. What is tricky is to ensure that this Exh operator is where it needs to be, namely *above* the preposition (to take alternatives that are bigger than just the nominal, i.e. alternatives of the form 'about DP'). In principle, P could still agree with an Exh in its c-command domain (downward agreement is always possible), as long as it is distinct from the Exh agreed with by n^0. One possibility is shown in (98).

(98)

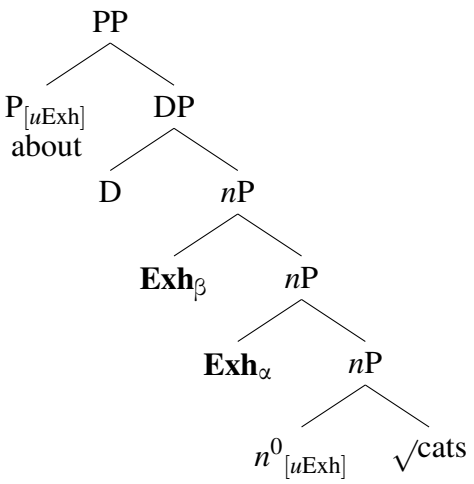

In (98), n^0 can Agree with Exh_α, and P can Agree with Exh_β. Syntactically, this converges. But there is no Exh in (98) that strengthens *about cats* to exclude other *about*-PPs.

One way to rule out (98), and force the Exh operator Agreed with by P to have P in its prejacent, would be to ban the stacking of Exh operators seen in (98) on independent grounds. Such a claim would be a little radical since stacking Exh operators has been common practice since Fox's (2007) analysis of Free Choice. But at least for Free Choice, it has recently been suggested by Bar-Lev & Fox (2017) to replace the stacked-Exh analysis with an analysis involving a single Exh capable of including alternatives in addition to excluding them (the analysis is summarised in chapter 1). It may therefore be possible to defend the view that Exh operators cannot be stacked immediately on top of one another, without needing to bulldoze through the already-existing literature. I leave this open as a tentative way to bar structures like (98) and maintain the agreement-based modelling of controlled Exh.

As a related point worth clarifying, back in the domain of simplex controllers, we must also rule out structures such as (99), where Exh is head-adjoined to n^0 before the resulting complex n^0 head merges with *cats*:

(99)

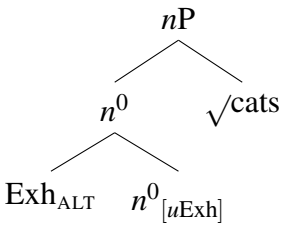

(99) would not exhaustify *cat* as needed. It might be possible to rule out the structure on purely syntactic grounds. Otherwise, it can be ruled out for semantic reasons: Exh cannot have a semantically vacuous prejacent.

Finally, in chapter 4, I suggested that DPs might come with null heads assigning them a thematic role, which would then be 'matched' by a clausal thematic head like Voice. The proposal is quite tentative; it is due to extrapolation from the data with PPs, and the desire to claim that all thematic uniqueness effects arise at least in part from controlled exhaustivity. If the proposal is accepted, then we must claim that these DP-internal thematic heads bear a [*u*Exh] feature as well.

[*u*Exh] and overt *only*

I end this section with another instance where Exh does not appear with cohyponyms, not discussed so far in this thesis. This is in the prejacent of *only*:[28]

(100) The flag is only green.
≈ the flag has no colour other than green

(100) presupposes that the flag has a green part and asserts that for all other colours, there is no part of that colour. If there was an Exh below *only*, however, *only*'s prejacent would entail that the flag is green and no other colour. (100) would therefore presuppose that the flag has no colour other than green, making *only*'s assertive contribution vacuous. This issue is avoided if there is simply no Exh in (100). On the theory of Exh's obligatoriness developed in this chapter, we can capture that Exh is not required in (100) by claiming that the [*u*Exh] feature on the a^0 head adjectivizing the root *green* can be satisfied through agreement with *only*. That is, *only* bears the same [*i*Exh] feature as Exh itself, making Exh non-required in its presence.

If *only* can replace Exh due to it satisfying the Agree relation in the same way, it is expected to be subject to the same locality constraints as Exh. This seems to be the case. In an example like (101) (the equivalent of a free exhaustivity effect—contrastive focus—but with an overt *only*), *only* can either surface in the same clause as the focused constituent or in an embedding clause of the form *it is the case that* . . . :

(101) a. Aisha only saw [the children]$_F$.
b. It is only the case that Aisha saw [the children]$_F$.

At least one meaning available to (101b) is that Aisha saw the children and no one else. In contrast, in examples that have been analyzed as involving controlled Exh, it is no longer the case that an *only* can surface at a distance from the focused controller:

(102) a. The flag is only [green]$_F$.
b. ??It is only the case that the flag is [green]$_F$.

Intuively, to me, (102b) is odd precisely because *green* already means 'entirely/only green' within *only*'s prejacent. The sentence is paraphrasable as '??it is only the case that the flag is entirely green,' which is hard to interpret because we do not know what the alternatives are meant to be. (102b) does *not* simply presuppose that the flag has a green part and assert that it does not have a part of any other colour, as is the case in (102a). I therefore suggest that (102b) has an Exh in addition to *only*, in contrast to (102a):

(103) It is only the case that the flag is [Exh$_{\text{ALT}}$ green]$_F$.

This can be understood as follows: the a^0 taking *green* as a root has a [*u*Exh] feature, which must Agree with an expression bearing [*i*Exh]; since *only* is syntactically higher than the *a*P, it is not a candidate for agreement in (102b) (in contrast to (102a)). There must therefore be an Exh in the *a*P in (102b) to satisfy the [*u*Exh] feature of a^0.

[28] I thank Aron Hirsch and Keny Chatain for pointing this out to me.

Interim summary

In this subsection, I have discussed how the hypothesis that derivational morphemes Agree with Exh, laid out for adjectives in section 6.4.1, plays out outside the domain of adjectives. I have shown some mixed results. The proposal can deal with nominals if determiners do not project, so that nominal phrases are *n*Ps rather than DPs (e.g., Bruening 2009). It can also deal with complex controllers like PPs and *v*Ps on some assumptions (namely, that there is an independent ban on stacked Exh operators). On the other hand, the proposal as stated predicts that scalar predicates should control Exh; they do not. This might be solvable through the right alternatives or semantic meaning for such predicates, but I do not currently have a concrete suggestion. Finally, I showed that the proposal can account for the fact that overt *only* does not co-occur with Exh if *only* has the same [*i*Exh] feature as Exh does; this is a welcome result.

6.5 Obligatory and local Exh beyond predicates

The above sums up my hypothesizing about what leads to the controlled nature of Exh with predicates: derivational morphemes (as well as prepositions) agree with Exh. Moving on, I now turn to the question of whether the sort of data focused on in this thesis is the only area where controlled exhaustivity (or something like it) is observed, or whether we are looking at part of a bigger pattern. I will limit my discussion to pointing out that there is another proposal in the literature which also requires exhaustivity to be both obligatory and local. This is Bar-Lev's (2018; 2021) analysis of homogeneity, which I already discussed in chapter 5. Bar-Lev does not posit such a locality constraint, but in fact, I will show that it is necessary to do so. Intriguingly, the empirical data discussed by Bar-Lev overlaps slightly with mine. Indeed, he discusses plural homogeneity effects, while part of the pattern I discuss involves subatomic homogeneity (the colour term data). Of course, the subatomic homogeneity pattern is only a subset of the effects for which I have posited controlled exhaustivity; most of the data cannot be described as a homogeneity effect. As such, there is only limited overlap between my empirical concern and Bar-Lev's. The existence of such overlap should still be taken seriously.

In this section, I first review once again Bar-Lev's theory (section 6.5.1), which uses Exh to create the homogeneity paradigm. Then, in section 6.5.2, I show that the positive side of homogeneity requires the stipulation that Exh is both obligatory with and necessarily local to the plural operator. Finally, in section 6.5.3, I point out that while no such stipulation is necessary for the negative side of homogeneity, nothing goes wrong if Exh is still subject to such a constraint there. As such, the conclusion is that the simplest way to capture all the data (on Bar-Lev's exhaustivity analysis of homogeneity) is to posit that Exh is generally obligatory with and local to plural morphology. I therefore claim that the plural operator has a [*u*Exh] feature.

6.5.1 A second summary of Bar-Lev's theory

Recall that Bar-Lev uses Exh's ability to include alternatives to derive homogeneity. I overview the theory here once again, focusing on the simpler case of distributive homogeneity. Bar-Lev's proposal is that the basic meaning of plurals is existential (104), due to an existential plural operator $\exists$-PL (105).

(104) $[\![\text{The kids laughed}]\!] = 1$ iff $\mathsf{laughed}(a) \vee \mathsf{laughed}(b)$

(105) a. $[\![\exists\text{-PL}]\!] = \lambda D_{\langle et\rangle}.\lambda P_{\langle e,st\rangle}.\lambda x_e.\exists y \in D \cap Part_{AT}(x)[P(y) = 1]$.
b. $Part_{AT}(x) = \{y : y \sqsubseteq_{AT} x\}$ (Bar-Lev 2021:1062)

(104) has the LF in (106), where the domain D is presented as a subscript on $\exists$-PL.

(106) [The kids] [$\exists$-PL$_D$ laughed]. (Bar-Lev 2021:1062)

Assuming that $[\![\text{the kids}]\!] = a \oplus b$ and $D = \{a, b\}$, (106) obtains the meaning in (107), which is equivalent to (104).

(107) $[\![(106)]\!] = 1$ iff $\exists y \in D \cap Part_{AT}([\![\text{the kids}]\!])[\mathsf{laughed}(y) = 1]$. (Bar-Lev 2021:1062)

Due to the existential meaning of (104), negating it obtains the right meaning for negative sentences without any further work:

(108) The kids didn't laugh.
$\approx$ none of the kids laughed

As such, no Exh is necessary for the meaning of negative sentences. In fact, on the standard definition of Exh, Exh would create problems if it occurred below negation; we return to this in section 6.5.3.

As for positive sentences, we need to strengthen (107) to create a universal. For Bar-Lev, the sentence's alternatives are subdomain alternatives:

(109) ALT $= \{$Aisha laughed $\vee$ Ben laughed, Aisha laughed, Ben laughed$\}$

None of these alternatives are innocently excludable, but they are all innocently includable. Exh therefore yields the meaning that all the children in the domain laughed:

(110) $[\![\text{Exh}_{\text{ALT}}\ [(106)]]\!] = 1$ iff
(Aisha laughed $\vee$ Ben laughed) $\wedge$ Aisha laughed $\wedge$ Ben laughed
$\equiv$ Aisha and Ben laughed.

6.5.2 Obligatory and local Exh in positive sentences

Clearly, on Bar-Lev's theory, Exh is obligatory in positive sentences. After all, if it was optional, plural sentences would optionally be interpreted as existential. This is something it has in common with the exhaustivity effect observed with cohyponyms. But plurals have something else in common with cohyponyms: their strengthening is necessarily computed locally to them (with 'locally' here meaning 'in the same clause').

Indeed, plurals' universal meaning in positive sentences is observed far beyond simple sentences like (104); in particular, it is observed even in downward-entailing (DE) environments. Križ (2015) discusses positive clauses under DE operators as sometimes involving universal meanings for plurals (111a), and sometimes involving existential meaning (111b):

(111) a. If you solve the problems, you will pass the exam. (Križ 2015:27)
$\approx$ 'If you solve **all** the problems, you will pass the exam.'

b. If you touch the statues, you will be asked to leave.
≈ 'If you touch **any** of the statues, you will be asked to leave.'

Križ therefore describes the positive side of homogeneity as optional in DE contexts. However, I claim that this apparent optionality is not due to the DE environment. In particular, the same judgments as in (111) hold even in UE contexts (cf. Winter & Scha 2015 on *touch*):

(112) a. A: Do you think you passed the exam?
B: Yes, I solved the problems.
≈ 'I solved **all** the problems.'
b. A: Why do you think you were asked to leave?
B: I touched the statues.
≈ 'I touched **at least some** of the statues.'

Since the existential meaning of the plural in (111b)/(112b) is observable in both UE and DE environments, we are clearly dealing with a run-of-the-mill non-maximality effect, rather than the result of the DE context in (111b). It is perfectly intuitive that non-maximality should be involved: the sentences invoke the scenario of a museum or the like, where visitors are asked to leave regardless of whether they touch one, multiple, or all of the statues, because touching any artefact is forbidden. As such, it is precisely the sort of sentence where we expect to observe a non-maximality effect so strong as to allow for plurals to be merely existential (e.g., Malamud 2012, Križ 2015). Outside of scenarios licensing non-maximality, the observation is that plurals are interpreted universally in positive clauses embedded in DE environments (111a).

As such, if Bar-Lev is right to claim that homogeneity is the result of an Exh operator, we are looking at an Exh operator which necessarily occurs under DE operators:

(113) a. If [Exh_{ALT} [you solve the problems]], you will pass the exam.
b. *Exh_{ALT} [if you solve the problems, you will pass the exam].

Indeed, the meanings from these two LFs are as in (114). Pretend there are two problems, A and B.

(114) a. $[\![(113a)]\!] = 1$ iff
you will pass the exam if [you solve problem A $\wedge$ you solve problem B].
b. $[\![(113b)]\!] = 1$ iff
[you will pass the exam if you solve problem A] $\wedge$ [you will pass the exam if you solve problem B].

(113a) is universal as desired, while (113b) is only existential. If Bar-Lev's use of exhaustivity to obtain the positive side of the homogeneity paradigm is the right approach, we learn that plural homogeneity effects involve an Exh operator that is both obligatory and necessarily local to (in the same clause as) the plural.

As written in (113a), the Exh operator is not 'ultra-local' to the alternative-triggering expression (the existential plural operator $\exists$-PL, or more particularly the *D* variable it takes). It is written out as taking an entire clause as its prejacent, and nothing goes wrong. But if we decide to push forward with the parallels between the exhaustification of predicates and of plurals (both involve obligatory and more or less local Exh operators), it might be better to claim that Exh is in fact even

more local to $\exists$-PL$_D$ than it is shown in (113a). I now show as a proof of concept that this is a possible claim, but note that the claim is only desirable for the theory-internal goal of collapsing as much as possible the constraint on Exh with pluralities and predicates.

Let's go back to the sentence in (115), repeated from (104):

(115) The kids laughed.

Except for the Exh operator, (115) has the LF in (116), repeated from (106).

(116) [The kids] [$\exists$-PL$_D$ laughed].

The alternative-triggering expression is $\exists$-PL$_D$. Can we get the right meaning with an Exh operator that is 'ultra-local' to this alternative-triggering expression?

On the one hand, a problem would arise if the LF was as in (117).

(117) [The kids] [Exh$_{\text{ALT}}$ [$\exists$-PL$_D$ laughed]].

Exh's prejacent in (117) has the meaning in (118); I assume again that the domain D is $\{a,b\}$.

(118) $[\![\exists\text{-PL}_D \text{ laughed}]\!] = \lambda x_e.\exists y \in \{a,b\} \cap \mathit{Part}_{AT}(x)[\mathsf{laughed}(y) = 1]$.

The problem here is that the intersection between $\{a,b\}$ and $\mathit{Part}_{AT}(x)$ cannot be determined before the *e*-type argument is provided. That is, Exh's prejacent must include the external argument *the kids*, since its part-structure is a *condicio sine qua non* for meaning to be properly computed.

This issue can be avoided if Bar-Lev's $\exists$-PL$_D$ is hosted by v^0. While $\exists$-PL$_D$ is syntactically lower than *the kids* on such a proposal, Exh is only so constrained as to be in the same maximal projection as $\exists$-PL$_D$; this projection is *v*P. It could therefore be adjoined to *v*P *above* the subject *the kids*:

(119)

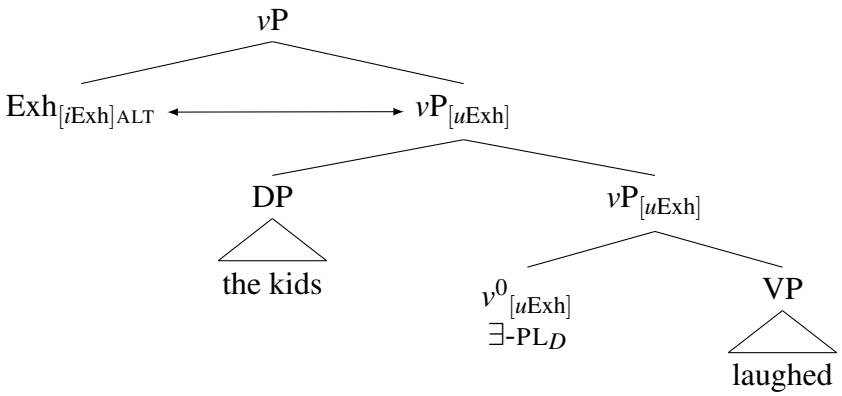

From here, Exh has the entire clause in its prejacent, so it can proceed as described by Bar-Lev (2021). But it is constrained to taking no more than the clause, deriving the locality constraint.

In sum, nothing goes wrong if we radically collapse the constraints on Exh (obligatoriness and locality) that are found with plurals and with predicates. Bar-Lev's (2018; 2021) theory of plural homogeneity *requires* the stipulation that the Exh associating with alternatives from $\exists$-PL$_D$ be obligatory and local to (in the same clause as) $\exists$-PL$_D$. But the theory is also *compatible* with the view that Exh is in fact obligatorily *ultra-local* to the alternative-triggering expression $\exists$-PL$_D$. On the latter view, the syntactic distribution of Exh associating with $\exists$-PL$_D$ can be fully collapsed

with the distribution of the Exh associating with cohyponyms.

6.5.3 Exh in negative sentences with pluralities

I noted above that Bar-Lev's theory is compatible with negative sentences not being exhaustified at all. (120) is repeated from (108).

(120) The kids didn't laugh.
≈ none of the kids laughed

Since the basic meaning of *the kids laugh* is existential, negating this existential statement yields the meaning observed in (120).

A curiosity for Bar-Lev's account is why the negation cannot scope above the Exh operator that is obligatorily present in positive sentences, as in (121).

(121) $[\![\text{not } [\text{Exh}_{\text{ALT}} \text{ [the kids laughed]}]]\!] = 1$ iff $\neg$[all the kids laughed].

(121) is the wrong meaning, so the LF apparently needs to be ruled out. But in fact, this is the same issue discussed for predicates in section 6.2.2,[29] where Exh also *prima facie* had to be stipulated not to occur under negation. I suggest the same solution: if the Exh operator in (121) is exchanged for the presuppositional Pexh of Bassi et al. (2021) (augmented with the Innocent Inclusion property), there is no need to rule out the syntax of (121):

(122) a. $[\![\text{Pexh}_{\text{ALT}} \text{ [the kids laughed]}]\!] = \begin{cases} 1, \text{if A and B laughed} \\ 0, \text{if neither A nor B laughed} \\ \#, \text{otherwise} \end{cases}$

b. $[\![\text{not } [\text{Pexh}_{\text{ALT}} \text{ [the kids laughed]}]]\!] = \begin{cases} 1, \text{if neither A nor B laughed} \\ 0, \text{if A and B laughed} \\ \#, \text{otherwise} \end{cases}$

c. Presupposition of (122b):
A and B laughed $\vee$ neither A nor B laughed
$\equiv$ A laughed $\leftrightarrow$ B laughed

Thus, it is possible to obtain the right truth conditions with a Pexh operator below negation; there is no need to stipulate that (P)exh cannot occur below negation. Note that, for what it's worth, we have also accidentally obtained the all-or-nothing presupposition of Löbner (2000), Gajewski (2005), et al. in (122c).

6.5.4 Interim conclusion

In this section, I pointed out that controlled Exh might be found in an empirical domain quite different from the meaning of predicates, namely plural homogeneity effects. If Bar-Lev (2018; 2021) is right that the positive side of the plural homogeneity paradigm is the result of an Exh operator innocently including subdomain alternatives, two constraints have to be stipulated for Exh:

[29] There is one difference, viz. that Exh > *not* would probably be harmless for Bar-Lev's account of plurals, whereas as shown in section 6.2.2 and chapter 2, it would cause problems for predicates.

it is obligatory, and necessarily local to (at least in the same clause as) the alternative-triggering expression $\exists$-PL$_D$. I have also shown that nothing goes wrong if we claim that Exh is in fact *always ultra-local* to the alternative-triggering $\exists$-PL$_D$ operator, as long as we adopt the presuppositional Pexh of Bassi et al. (2021) to deal with the negation data. Thus, the parallel between the syntax of Exh with predicates and with plurals might be perfect. In line with the theory developed in section 6.4, I therefore suggest that the plural operator Agrees with Exh. (123) is repeated from (119).

(123)

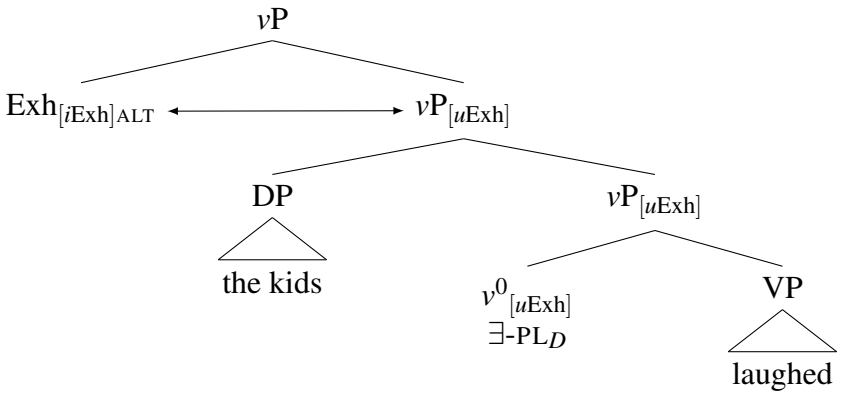

The difference between the exhaustification of plurals and the exhaustification of predicates is the nature of the alternatives that Exh takes as its domain, and whether, as a result of the nature of the alternatives, Exh ends up including or excluding them.

Returning to the question from chapter 5 about whether subatomic and plural homogeneity can be theoretically collapsed on my controlled-Exh account of subatomic homogeneity, the answer is yes—just not entirely. Both involve weak lexical meaning that is exhaustified ultra-locally in positive sentences; they differ as just stated.

6.6 Conclusion

As discussed throughout this thesis, taxonomic predicates display an exclusivity effect, which I have argued is the result of an Exh operator. But such an account comes with an unusual-looking explanandum, as laid out in chapter 2: this Exh would have to be both obligatory and subject to a strict locality requirement. In this chapter, I have suggested that the locality requirement is best understood as the Exh operator needing to be in the same maximal projection as the predicate that requires it. I dealt with data that initially appear not to have an Exh operator (specifically sentential negation and perhaps other negative-flavoured DE environments) by claiming that the Exh operator is in fact the presuppositional Pexh operator of Bassi et al. (2021). To understand why predicates would require an Exh in their maximal projection, I turned to another part of grammar that is obligatory and local with predicates, viz. derivational morphology. I argued that derivational morphemes Agree with Exh, forcing one to appear within their maximal projection. The [*u*Exh] feature on derivational morphemes must find an Exh operator for the derivation to converge, and this Exh cannot be higher than the morpheme's maximal projection; there is no upward agreement. This derives both the obligatory nature of Exh and its locality constraint.

While this is, as far as I'm aware, the first explicit discussion of a twin obligatory-presence and locality constraint on Exh, there is in fact at least one other proposal in the literature that leads to

these same two properties for Exh. This is the exhaustivity account of plural homogeneity suggested by Bar-Lev (2018, 2021). I showed that for Bar-Lev's account to work, Exh (in fact, Pexh) must be taken to be not just obligatory, but also necessarily local to the alternative-triggering expression (viz. $\exists$-PL$_D$). The controlledness of Exh with predicates and with plural number morphology can be collapsed as a single pattern. As a direction for future research, it should be understood what unites predicates and number, causing both of these, but not many other alternative-triggering expressions, to control Exh. As a start, this can be modelled through a [uExh] feature on derivational morphemes and the head hosting $\exists$-PL$_D$ (not to mention P), but one wonders what governs the distribution of this feature.

Returning to the agreement between derivational morphemes and Exh, the take-away from this hypothesis is that derivational morphology is directly involved in shaping our intuitions about the meaning of roots. When derivational morphemes take a root (roughly corresponding to a concept, in the basic case) and make it linguistically usable by giving it a category, they also strengthen the meaning of the concept so as to exclude other concepts. The effect of this is as if derivational morphemes were cleaning up the nonlinguistic conceptual space by ensuring that predicates' meanings do not overlap. As we have seen, this can come at the cost of 'unnecessarily' creating contradictory sentences like (124).

(124) #This comedy is a tragedy.

Chapter 7

Conclusion

7.1 Thesis summary

What is the distribution of exhaustification in human language? There are at least two components to this question. First, what is the set of expressions whose meaning involves exhaustification? Second, what are the constraints (if any) on the syntactic loci from which this exhaustification can be computed? In this thesis, I contributed to answering both these questions first by showing that many more expressions involve exhaustification than previously thought, and by showing that they constrain the distribution of Exh in a way not previously discussed in the literature.

Taxonomic predicates, usually taken to be essentially unanalysable units in formal semantics, are routinely exhaustified in natural-language sentences. They are interpreted as strong, i.e. mutually exclusive with other predicates sharing a particular natural class with them. What counts as their natural class can vary sentence by sentence, a point I will return to below. These predicates are interpreted as mutually incompatible due to strengthening in all environments, except when they are joined through conjunction or an additive particle, or when they are in a negative sentence. While these latter environments let us observe that taxonomic predicates' underlying meaning is weak (so that the strong intuited meanings must be the result of strengthening), they are also all environments where it is possible (and in some cases necessary) to posit that the predicates are indeed exhaustified, at least if we adopt the presuppositional Pexh operator of Bassi et al. (2021). The view that the meanings of taxonomic predicates are in part the result of the exclusion of other predicates is in effect a modified revival of the insight of de Saussure (2011[1916]) that predicates delimit one another, and do so in the same way as logical vocabulary items delimit one another (e.g., tense and number). My proposal differs from de Saussure's (2011[1916]) insight in two ways: I take the mutual delimitation of predicates to occur in the grammar rather than the lexicon, and I take the class of predicates that are excluded to be context-dependent rather than fixed. Still, this thesis essentially motivates an updated, 'neo-structuralist' semantics.

Discussing the exhaustification of taxonomic predicates also creates a path forward to understanding the second question above: what is the syntactic distribution of Exh in sentences? Something quite new comes out from the data with taxonomic predicates: they 'control' Exh, meaning that they not only require its presence in the sentence, but also have a say in its syntactic position. I argued that Exh must be within the XP of the predicate, and I modelled this through an Agree relation between derivational morphemes and Exh. Derivational morphemes are effectively involved

in partitioning the conceptual spaces that the concepts denoted by the roots they select come from. Derivational morphemes are involved substantially in affecting the meaning of roots.

The signatures of controlled exhaustivity effects (the creation of clause-internal contradictions, interaction with *and* and additives, disappearance under sentential negation) are found beyond taxonomic predicates. I mostly focused on PPs. From the observation that sentences do not allow more than one PP with the same preposition, unless they are joined by *and* or *also*, I argued more generally that thematic uniqueness effects are the result at least of a controlled Exh operator. With DP arguments, independent syntactic constraints may be at play in creating thematic uniqueness as well. In addition to taxonomic predicates and prepositions, I showed that analyses of plurality that claim that plurals undergo strengthening in positive sentences (Magri 2014; Bar-Lev 2018, 2021) also overlap at least in part with the claims about Exh's distribution with taxonomic predicates and prepositions. The Exh leading to the strong meaning of plural predication must be obligatory and reasonably local to (at least in the same clause as) the plural operator. Finally, I also showed that some clause-like constituents containing contrastive topics necessarily involve local exhaustification as well, as in (1).

(1) The train crash that hurt Aisha #(also) hurt Ben.

While such examples have been far from the main focus of this thesis, they offer a clear empirical bridge with the discourses normally discussed in the literature on obligatory additive particles (e.g., Bade 2016), which also make use of contrastive topics:

(2) The train crash hurt Aisha. It #(also) hurt Ben.

Future work should clarify how the effect in (1), where at least one Exh operator must have only one of the two VPs in its prejacent, relates to examples like (2), and in particular Bade's (2016) claim, due to data with sentential conjunctions, that the Exh operators in (2) are free.

The fact that taxonomic predicates and other expressions are systematically strengthened in sentences raises the question of which expressions Exh excludes. I have proposed to give theoretical status to what seems like an obvious descriptive notion, namely that different expressions contribute different kinds of information to sentences. They have different 'jurisdictions' in meaning. I claimed that expressions are alternatives for controlled Exh if they share a jurisdiction. For example, all form-denoting concrete nouns are alternatives to one another for controlled Exh, to the exclusion of other types of nouns. While most of the thesis has discussed the exhaustification of predicates in terms of cohyponymy, this is really too narrow a notion. The notion of jurisdictions can also be extended beyond simplex predicates to expressions like PPs and thematic uniqueness effects more generally. I argued that apparent counterexamples to thematic uniqueness may be understood through the notion of jurisdiction if other factors like the figure–ground distinction can be taken as sufficient for two constituents not to share a jurisdiction.

Finally, throughout the thesis, I more or less quietly built up a theory of subatomic homogeneity. I collapsed the exhaustification of summative predicates like colour terms, whose meanings can be described as a subatomic homogeneity effects, with the exhaustification of integrative predicates. This is well motivated from these predicates' identical interactions with additives, conjunction, and sentential negation. On the view developed here, subatomic homogeneity arises from weak lexical meaning together with exhaustification (which is non-trivial in positive sentences, but trivial in negative sentences, due to the semantics of the Pexh operator of Bassi et al. (2021)).

Summative predicates behave like all other taxonomic predicates in being exhaustified to exclude same-jurisdiction alternatives. A sentence like (3) therefore means that the flag is 'only' green; it is not grey, purple, blue, and so on.

(3) The flag is green.

In contrast to other theories of homogeneity, there is nothing linguistic here entailing that all parts of the flag are green; this is due to world knowledge that all parts must have some colour. This approach is well motivated from the felicitous paraphrase of (3) in (4).

(4) The flag is only green.

However, this analysis is not extendable to plural homogeneity; clearly, world knowledge and exclusion of alternative predicates does not suffice to make an existential plural operator universal. The flip side of this is that, as I showed, theories of homogeneity made for pluralities do not extend to subatomic homogeneity, either. Specifically, they cannot handle conjunctions of summative predicates with atomic subjects. Roughly speaking, they result in (5) meaning that all pieces of the flag are both white and green.

(5) The flag is white and green.

In chapter 6, I therefore adopted Bar-Lev's (2018; 2021) exhaustivity approach to plural homogeneity, as an addition to my theory of subatomic homogeneity. This collapses the plural and subatomic homogeneity paradigms insofar as they involve local and obligatory exhaustification, while differing in the nature of the alternatives taken by Exh (and as a result of this, whether Exh excludes or includes them).

In the rest of this concluding chapter, I discuss some empirical data suggesting that there may be still more to learn about the meaning of taxonomic predicates from additive particles.

7.2 Still more to learn about predication from additives

In this thesis, I have used additive particles (among other logical expressions) to learn about the meaning of content predicates. We might think that this is all that additives have to teach us about predication, but in fact, it seems there is still more to learn about predication from additives.

Most work on additive particles in English use *too*, not *also*. In this thesis, I moved to *also* without comment. There is nothing suspect about this; all the data in the literature using *too* could just have well used *also*. The only exception is from Göbel (2019), who notes some differences between *also* and *too*, in particular that *also* is sometimes acceptable where *too* is not:

(6) SCENARIO: Aisha and Ben are watching *Dexter*, a TV-show about a psychopath who satisfies his homicidal urges by killing criminals. In one scene, he donates some money to charity before killing someone. While Ben detests Dexter's actions, Aisha sympathizes with him. (Göbel 2019:281)

A: Dexter is such a good person!
B: He **also** MURDERS PEOPLE.
B′: #He MURDERS PEOPLE **too**.

In fact, the reason I used *also* in this thesis is because there are also differences between *also* and *too* in the domain of clause-internal additivity with predicates. Not all such sentences allow *too*. Let's start with an optional additivity effect:

(7) a. The morning star is also the evening star.
b. ??The morning star is the evening star too.

What is puzzling about (7b) is that, when the additivity isn't sentence-internal, *too* is acceptable:

(8) This is the morning star. It's the evening star, too.

In contrast, some sentence-internal additivity effects *do* allow *too* (9). This prevents us from making the simple claim that *too* is simply incompatible with sentence-internal additivity.

(9) a. (i) My birthday is also your birthday.
(ii) My birthday is your birthday too.
b. (i) This film is also a book. *(i.e., a story exists as a film and as a novel)*
(ii) This film is a book too.

What is more, sentences that disallow *too* can be paraphrased with *simultaneously*, while sentences that allow *too* cannot. (10) disallows *too* but not *simultaneously*:

(10) a. This fork is **also** a spoon.
b. #This fork is a spoon **too**.
c. This spoon is **simultaneously** a spoon.

Meanwhile, the sentences in (11) allow *too* (9) but not *simultaneously*:

(11) a. #My birthday is **simultaneously** your birthday.
b. #This film is **simultaneously** a book.

It is not clear how much this generalization can be maintained in this simple form. For instance, I am not sure about the *morning star* example (7), which disallows *too* but is far from perfect with *simultaneously*:

(12) ?The morning star is simultaneously the evening star.

Putting aside the *morning star* example, there are a few ways one might think about this contrast. At first glance, we could think of it as concerning whether there is a single referent or two. 'This film is a book too' could be taken as allowing *too* and disallowing *simultaneously* because there are two referents (the story that exists as a film, and the (same) story that exists as a book), in contrast to (10) where there is a single referent (the spork). But that would not hold of birthdays if *my birthday* simply refers to a date in the calendar.

Perhaps a more promising avenue is to point to whether there is a causal or temporal relation between two predicates holding of a given individual. The fact that 'my birthday' and 'your birthday' are on the same day is an accident, and the existence of the story qua film and the story qua novel is roughly speaking unrelated (one can exist without the other). For the spork in (10), however, it simultaneously has the essence of a fork and spoon as a core part of its existence from start to finish. Of course, this is just a description of the data (and may be far from correct), and raises

difficult questions about why *too* but not *also* would be sensitive to this distinction.

In conclusion, it appears that the role of exhaustification in the intuited meanings of predicates is not the only way that additives shed light on predication. Apparently, *too* and *also* interact with subtle aspects of predication, not fully understood at present.

Bibliography

Abney, Steven Paul. 1987. The English noun phrase in its sentential aspect. Doctoral dissertation, MIT, Cambridge, MA.

Alxatib, Sam, & Jeff Pelletier. 2011. On the psychology of truth-gaps. In *Vagueness in communication*, ed. Rick Nouwen, Robert van Rooij, Uli Sauerland, & Hans-Christian Schmitz, 13–36. Berlin: Springer.

Amsili, Pascal, & Claire Beyssade. 2006. Compulsory presupposition in discourse. In *Proceedings of the workshop on constraints in discourse*, ed. Candy Sidner, John Harpur, Anton Benz, & Peter Kühnlein, 5–10. Maynooth: National University of Ireland.

Arad, Maya. 2003. Locality constraints on the interpretation of roots: The case of Hebrew denominal verbs. *Natural Language & Linguistic Theory* 21:737–779.

Arad, Maya. 2005. *Roots and patterns: Hebrew morpho-syntax*. Dordrecht: Kluwer Academic Publishers.

Aravind, Athulya, & Martin Hackl. 2017. Against a unified treatment of obligatory presupposition trigger effects. In *Proceedings of SALT 27*, ed. Dan Burgdorf, Jacob Collard, Sireemas Maspong, & Brynhildur Stefánsdóttir, 173–190. Washington, DC: The Linguistic Society of America.

Asher, Nicholas, & Alex Lascarides. 1998. The semantics and pragmatics of presupposition. *Journal of Semantics* 15:239–299.

Bade, Nadine. 2014. Obligatory implicatures and the presupposition of "too". In *Proceedings of Sinn und Bedeutung 18*, ed. Urtzi Etxeberria, Anamaria Fălăuș, Aritz Irurtzun, & Bryan Leferman, 42–59.

Bade, Nadine. 2016. Obligatory presupposition triggers in discourse: Empirical foundations of the theories *Maximize Presupposition* and *Obligatory Implicatures*. Doctoral dissertation, University of Tübingen, Tübingen.

Bade, Nadine. 2021. On the scope and nature of *Maximise Presupposition*. *Language and Linguistics Compass* 1–16.

Bade, Nadine, & Agata Renans. 2021. A cross-linguistic view on the obligatory insertion of additive particles—Maximize Presupposition vs. Obligatory Implicatures. *Glossa* 6:1–24.

Bade, Nadine, & Sonja Tiemann. 2016. Obligatory triggers under negation. In *Proceedings of Sinn und Bedeutung 20*, ed. Nadine Bade, Polina Berezovskaya, & Anthea Schöller, 109–126.

Bar-Lev, Moshe E. 2018. Free Choice, homogeneity, and Innocent Inclusion. Doctoral dissertation, The Hebrew University of Jerusalem, Jerusalem.

Bar-Lev, Moshe E. 2021. An implicature account of Homogeneity and Non-maximality. *Linguistics and Philosophy* 44:1–53.

Bar-Lev, Moshe E., & Danny Fox. 2017. Universal Free Choice and Innocent Inclusion. In *Proceedings of SALT 27*, ed. Dan Burgdorf, Jacob Collard, Sireemas Maspong, & Brynhildur Stefánsdóttir, 95–115.

Bar-Lev, Moshe E., & Daniel Margulis. 2014. Hebrew *kol*: a universal quantifier as an undercover existential. In *Proceedings of Sinn und Bedeutung 18*, ed. Urtzi Etxeberria, Anamaria Fălăuș, Aritz Irurtzun, & Bryan Leferman, 60–76.

Bassi, Itai, & Moshe E. Bar-Lev. 2018. A unified existential semantics for bare conditionals. In *Proceedings of Sinn und Bedeutung 21*, ed. Robert Truswell, Chris Cummins, Caroline Heycock, Brian Rabern, & Hannah Rohde, 125–142.

Bassi, Itai, Guillermo Del Pinal, & Uli Sauerland. 2021. Presuppositional exhaustification. *Semantics and Pragmatics* 14. Early access version.

Beaver, David, & Emiel Krahmer. 2001. A partial account of presupposition projection. *Journal of Logic, Language and Information* 10:147–182.

Beaver, David, Craige Roberts, Mandy Simons, & Judith Tonhauser. 2017. Questions Under Discussion: where information structure meets projective content. *Annual Review of Linguistics* 3:265–84.

Beaver, David Ian. 1997. Presupposition. In *Handbook of logic and language*, ed. Johan van Benthem & Alice ter Meulen, 939–1008. Amsterdam: Elsevier.

Béjar, Susana, & Milan Rezac. 2009. Cyclic Agree. *Linguistic Inquiry* 40:35–73.

Benz, Anton, & Katja Jasinskaja. 2017. Questions under discussion: From sentence to discourse. *Discourse Processes* 54:177–186.

Blake, Frank R. 1930. A semantic analysis of case. In *Curme volume of linguistic studies (language monographs 7)*, ed. James T. Hatfield, Werner Leopold, & A. J. Friedrich Zieglschmid, 34–49. Baltimore: Linguistic Society of America.

Blutner, Reinhard. 2002. Lexical semantics and pragmatics. In *Linguistische Berichte Sonderheft 10. Semantics*, ed. Fritz Hamm & Thomas Ede Zimmermann, 27–58. Hamburg: Helmut Buske Verlag.

Blutner, Reinhard. 2004. Pragmatics and the lexicon. In *The handbook of pragmatics*, ed. Laurence R. Horn & Gregory Ward, 488–514. Malden, MA: Blackwell.

Boeckx, Cedric. 2011. Some reflections on Darwin's Problem in the context of Cartesian biolinguistics. In *The biolinguistic enterprise: New perspectives on the evolution and nature of the human language faculty*, ed. Anna Maria di Sciullo & Cedric Boeckx, 42–64. Oxford: Oxford University Press.

Borer, Hagit. 2014. The category of roots. In *The syntax of roots and the roots of syntax*, ed. Artemis Alexiadou, Hagit Borer, & Florian Schäfer, 112–148. Oxford: Oxford University Press.

Bowler, Margit. 2014. Conjunction and disjunction in a language without 'and'. In *Proceedings of the 24th Semantics and Linguistic Theory Conference (SALT)*, ed. Todd Snider, Sarah D'Antonio, & Mia Weigand, 137–155. LSA and CLC.

Bresnan, Joan. 1982. The passive in lexical theory. In *The mental representation of grammatical relations*, ed. Joan Bresnan, 3–86. Cambridge, MA: MIT Press.

Bresnan, Joan W. 1976. On the form and functioning of transformations. *Linguistic Inquiry* 7:3–40.

Bruening, Benjamin. 2009. Selectional asymmetries between CP and DP suggest that the DP Hypothesis is wrong. *University of Pennsylvania Working Papers in Linguistics* 15:27–35.

Bruening, Benjamin. 2020. The head of the nominal is N, not D: N-to-D movement, hybrid agreement, and conventionalized expressions. *Glossa* 5:1–19. Article 15.

Bruening, Benjamin. 2022. Hebrew nominals do not require functional structure above the NP. *Glossa* 7:1–45.

Bruening, Benjamin, Xuyen Dinh, & Lan Kim. 2018. Selection, idioms, and the structure of nominal phrases with and without classifiers. *Glossa* 3:1–46. Article 42.

Büring, Daniel. 2016. (Contrastive) topic. In *The Oxford handbook of information structure*, ed. Caroline Féry & Shinichiro Ishihara, 64–85. Oxford: Oxford University Press.

Burton, Strang, & Jane Grimshaw. 1992. Coordination and VP-internal subjects. *Linguistic Inquiry* 23:305–313.

Camacho, José. 2003. *The structure of coordination: Conjunction and agreement phenomena in Spanish and other languages.*. Dordrecht: Kluwer Academic Publishers.

Cann, Ronnie. 2011. Sense relations. In *Semantics: An international handbook of natural language meaning*, ed. Claudia Maienborn, Klaus von Heusinger, & Paul Portner, volume 1, 456–479. Berlin: de Gruyter Mouton.

Carlson, Greg. 1984. On the role of thematic roles in linguistic theory. *Linguistics* 22:259–279.

Carlson, Greg. 1998. Thematic roles and the individuation of events. In *Events and grammar*, ed. Susan Rothstein, 35–51. Dordrecht: Kluwer Academic Publishers.

Carnap, Rudolf. 1947. *Meaning and necessity*. Chicago: University of Chicago Press.

Carnap, Rudolf. 1955. Meaning and synonymy in natural languages. *Philosophical Studies* 6:33–47.

Carnap, Rudolf. 1963. Replies and systematic expositions. In *The philosophy of Rudolf Carnap*, ed. Paul Arthur Schilpp, 859–888. LaSalle, IL: Open Court.

Champollion, Lucas. 2016. Ten men and women got married today: Noun coordination and the intersective theory of conjunction. *Journal of Semantics* 33:561–622.

Chaves, Rui P. 2012. On the grammar of extraction and coordination. *Natural Language & Linguistic Theory* 30:465–512.

Chemla, Emmanuel. 2008. An epistemic step for anti-presuppositions. *Journal of Semantics* 25:141–173.

Chierchia, Gennaro. 2004. Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In *Structures and beyond: The cartography of syntactic structures*, ed. Adriana Belletti, volume 3, 39–103. Oxford: Oxford University Press.

Chierchia, Gennaro. 2006. Broaden your views: Implicatures of domain widening and the 'logicality' of language. *Linguistic Inquiry* 37:535–590.

Chierchia, Gennaro. 2013. *Logic in grammar: polarity, free choice, and intervention*. Oxford: Oxford University Press.

Chierchia, Gennaro, Danny Fox, & Benjamin Spector. 2012. Scalar implicatures as a grammatical phenomenon. In *Semantics: An international handbook of natural language meaning*, ed. Claudia Maienborn, Klaus von Heusinger, & Paul Portner, volume 3, 2297–2331. Berlin: De Gruyter Mouton.

Chomsky, Noam. 1957. *Syntactic structures*. The Hague: Mouton.

Chomsky, Noam. 1981. *Lectures on government and binding*. Dordrecht: Foris.

Chomsky, Noam. 1995a. Bare phrase structure. In *Evolution and revolution in linguistic theory*, ed. Héctor Campos & Paula Kempchinsky, 51–109. Washington DC: Georgetown University Press.

Chomsky, Noam. 1995b. *The minimalist program*. Cambridge, MA: MIT Press.

Chomsky, Noam. 2000. Minimalist inquiries: The framework. In *Step by step: Essays on minimalist syntax in honor of Howard Lasnik*, ed. Roger Martin, David Michaels, Juan Uriagereka, & Samuel Jay Keyser, 89–155. Cambridge, MA: MIT Press.

Chomsky, Noam. 2001. Derivation by phase. In *Ken Hale: A life in language*, ed. Michael Kenstowicz, 1–52. Cambridge, MA: MIT Press.

Chomsky, Noam. 2008. On phases. In *Foundational issues in linguistic theory: Essays in honor of Jean-Roger Vergnaud*, ed. Robert Freidin, Carlos P. Otero, & Maria Luisa Zubizarreta, 133–166. Cambridge, MA: MIT Press.

Chomsky, Noam. 2013. Problems of projection. *Lingua* 130:33–49.

Chomsky, Noam. 2015. Problems of projection: Extensions. In *Structures, strategies and beyond: Studies in honour of Adriana Belletti*, ed. Elisa Di Domenico, Cornelia Hamann, & Simona Matteini, 1–16. Amsterdam: John Benjamins.

Contreras, Heles. 1987. Small clauses in Spanish and English. *Natural Language & Linguistic Theory* 5:225–243.

Coppock, Elizabeth, & Lucas Champollion. in progress. *Invitation to formal semantics*.

Dalrymple, Mary, Makoto Kanazawa, Sam Mchombo, & Stanley Peters. 1994. What do reciprocals mean? In *Proceedings of SALT 4*, ed. Mandy Harvey & Lynn Santelmann, 61–78. Ithaca, NY: Cornell University.

Davidson, Donald. 1976. The logical form of action sentences. In *The logic of decision and action*, ed. Nicholas Rescher, 81–95. Pittsburgh, PA: University of Pittsburgh Press.

De Belder, Marijke, & Jeroen van Craenenbroeck. 2015. How to merge a root. *Linguistic Inquiry* 46:625–655.

Dowty, David. 1991. Thematic proto-roles and argument selection. *Language* 67:547–619.

Einstein, Albert. 1971. Letter 84. In *The Born–Einstein letters: correspondence between Albert Einstein and Max and Hedwig Born from 1916 to 1955, with commentaries by Max Born*, 157–158. London: Macmillan. Translated by Irene Born.

Erlewine, Michael Yoshitaka. 2017. Vietnamese focus particles and derivation by phase. *Journal of East Asian Linguistics* 26:325–349.

Fauconnier, Gilles. 1975. Pragmatic scales and logical structure. *Linguistic Inquiry* 6:353–375.

Fillmore, Charles J. 1968. The case for Case. In *Universals in linguistic theory*, ed. Emmon Bach & Robert T. Harms, 1–90. New York: Holt: Rinehart & Winston.

von Fintel, Kai. 2001. Conditional strengthening: a case study in implicature. MIT manuscript. Available at: `http://web.mit.edu/fintel/fintel-2001-condstrength.pdf`.

Fox, Danny. 2007. Free choice and the theory of scalar implicatures. In *Presupposition and implicature in compositional semantics*, ed. Uli Sauerland & Penka Stateva, 71–120. Basingstoke, Hampshire (UK): Palgrave Macmillan.

Fox, Danny. 2013. Presupposition projection from quantificational sentences: trivalence, local accommodation, and presupposition strengthening. In *From grammar to meaning: The spontaneous logicality of language*, ed. Ivano Caponigro & Carlo Cecchetto, volume 2: Formal developments, new findings, and challenges, 201–232. Cambridge: Cambridge University Press.

Fox, Danny, & Martin Hackl. 2006. The universal density of measurement. *Linguistics and Philosophy* 29:537–586.

Fox, Danny, & Roni Katzir. 2011. On the characterization of alternatives. *Natural Language Semantics* 19:87–107.

Fox, Danny, & Benjamin Spector. 2018. Economy and embedded exhaustification. *Natural Language Semantics* 26:1–50.

van Fraassen, Bas C. 1967. Meaning relations among predicates. *Noûs* 1:161–179.

Francis, Naomi. 2019. Imperatives under *even*. In *Proceedings of the North East Linguistics Society 49*.

Francis, Naomi. 2020. Licensing free choice *any* on an existential semantics for imperatives. In *Proceedings of Sinn und Bedeutung 24*, ed. Michael Franke, Nikola Kompa, Mingya Liu, Jutta L. Mueller, & Juliane Schwab, volume 1, 193–204.

Franks, Steven, & Jana Willer-Gold. 2014. Agreement strategies with conjoined subjects in Croatian. In *New insights into Slavic linguistics*, ed. Jacek Witkoś & Sylwester Jaworski, 91–115. Frankfurt am Main: Peter Lang.

Gajewski, Jon. 2005. Neg-Raising: Polarity and presupposition. Doctoral dissertation, MIT, Cambridge, MA.

Gärdenfors, Peter. 2000. *Conceptual spaces: The geometry of thought*. Cambridge, MA: MIT Press.

Gawron, Jean Mark, & Andrew Kehler. 2004. The semantics of respective readings, conjunction, and filler-gap dependencies. *Linguistics and Philosophy* 27:169–207.

Gazdar, Gerald. 1979. *Pragmatics: implicature, presupposition, and logical form*. New York: Academic Press.

Gazdar, Gerald. 1981. Unbounded dependencies and coordinate structure. *Linguistic Inquiry* 12:155–184.

Geis, Michael L., & Arnold M. Zwicky. 1971. On invited inferences. *Linguistic Inquiry* 2:561–566.

George, Benjamin. 2008. Presupposition repairs: a static, trivalent approach to predicting projection. Master's thesis, UCLA, Los Angeles.

Geurts, Bart. 2010. *Quantity implicatures*. Cambridge: Cambridge University Press.

Geurts, Bart, & Rob van der Sandt. 2004. Interpreting focus. *Theoretical Linguistics* 30:1–44.

Göbel, Alexander. 2019. Additives pitching in: L*+H signals ordered Focus alternatives. In *Proceedings of SALT 29*, 279–299.

Goodman, Nelson. 1949. On likeness of meaning. *Analysis* 10:1–7.

Goodman, Nelson. 1953. On some differences about meaning. *Analysis* 13:90–96.

Goodman, Nelson. 1968. *Languages of art: An approach to a theory of symbols*. Indianapolis: Bobbs-Merrill.

Green, Georgia M. 1973. The lexical expression of emphatic conjunction: theoretical implications. *Foundations of Language* 10:197–248.

Grice, H. Paul. 1975. Logic and conversation. In *Syntax and semantics*, ed. Cole Peter & Jerry L. Morgan, volume 3: Speech acts, 41–58. New York: Academic Press.

Grice, H. Paul. 1989. *Studies in the way of words*. Cambridge, MA: Harvard University Press.

Gruber, Jeffrey. 1965. Studies in lexical relations. Doctoral dissertation, MIT, Cambridge, MA.

Guilfoyle, Eithne, Henrietta Hung, & Lisa Travis. 1992. Spec of IP and Spec of VP: Two subjects in Austronesian languages. *Natural Language & Linguistic Theory* 10:375–414.

Hampton, James A. 2006. Concepts as prototypes. In *Psychology of learning and motivation*, ed. Brian H. Ross, volume 46, 79–113. New York: Academic Press.

Harley, Heidi. 2014. On the identity of roots. *Theoretical Linguistics* 40:225–276.

Harley, Heidi. 2017. The 'bundling' hypothesis and the disparate function of little v. In *The Verbal Domain*, ed. Roberta D'Alessandro, Irene Franco, & Ángel J. Gallego, 3–28. Oxford: Oxford University Press.

Harley, Heidi, & Rolf Noyer. 1997. Mixed nominalization, short verb movement and object shift in English. In *Proceedings of NELS 28*, ed. Pius N. Tamanji & Kiyomi Kusumoto, 143–157. Amherst, MA: GLSA, University of Massachusetts, Amherst.

Harnish, Robert. 1976. Logical form and implicature. In *An integrated theory of linguistic ability*, ed. Thomas G. Bever, Jerrold J. Katz, & D. Terence Langendoen, 313–391. New York: Thomas Cromwell.

Heim, Irene. 1991. Artikel und Definitheit [Articles and definiteness]. In *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, ed. Arnim von Stechow & Dieter Wunderlich, 487–535. Berlin: de Gruyter.

Heim, Irene. 1992. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics* 9:183–221.

Heim, Irene, & Angelika Kratzer. 1998. *Semantics in generative grammar*. Malden, MA: Blackwell.

Herburger, Elena. 2015. Conditional perfection: the truth and the whole truth. In *Proceedings of SALT 25*, 615–635.

Heycock, Caroline, & Roberto Zamparelli. 2005. Friends and colleagues: Plurality, coordination, and the structure of DP. *Natural Language Semantics* 13:201–270.

Hirsch, Aron. 2017. An inflexible semantics for cross-categorial operators. Doctoral dissertation, MIT, Cambridge, MA.

Hoek, Daniel. 2019. Loose talk, scale presuppositions and QUD. In *Proceedings of the 22nd Amsterdam Colloquium*, ed. Julian J. Schlöder, Dean McHugh, & Floris Roelofsen, 171–180.

Horn, Laurence. 1972. On the semantic properties of logical operators in English. Doctoral dissertation, UCLA, Los Angeles.

Horn, Laurence. 1985. Metalinguistic negation and pragmatic ambiguity. *Language* 61:121–174.

Horn, Laurence. 2000. From *if* to *iff*: conditional perfection as pragmatic strengthening. *Journal of Pragmatics* 32:289–326.

Horn, Laurence. 2001[1989]. *A natural history of negation*. Stanford: CSLI Publications.

Hornstein, Norbert. 1999. Movement and control. *Linguistic Inquiry* 30:69–96.

Hornstein, Norbert, Jairo Nunes, & Kleanthes K. Grohmann. 2005. *Understanding minimalism*. Cambridge: Cambridge University Press.

Horvath, Julia. 2007. Separating "focus movement" from focus. In *Phrasal and clausal architecture*, ed. Simin Karimi, Vida Samiian, & Wendy K. Wilkins, 108–145. Amsterdam: John Benjamins.

Hurford, James R. 1974. Exclusive or inclusive disjunction. *Foundations of Language* 11:409–411.

Jackendoff, Ray. 1972. *Semantic interpretation in generative grammar*. Cambridge, MA: MIT Press.

Johannessen, Janne Bondi. 1996. Partial agreement and coordination. *Linguistic Inquiry* 27:661–676.

Johannessen, Janne Bondi. 1998. *Coordination*. Oxford: Oxford University Press.

Kaplan, Jeff. 1984. Obligatory *too* in English. *Language* 60:510–518.

Karttunen, Lauri, & Stanley Peters. 1979. Conventional implicature. In *Syntax and semantics*, ed. Choon-Kyu Oh & David Dineen, volume 11: Presupposition, 1–56. New York: Academic Press.

Katzir, Roni. 2007. Structurally-defined alternatives. *Linguistics and Philosophy* 30:669–690.

Katzir, Roni. 2013. A note on contrast. *Natural Language Semantics* 21:333–343.

Katzir, Roni. 2014. On the roles of markedness and contradiction in the use of alternatives. In *Semantics, pragmatics, and the case of scalar implicatures*, ed. Salvatore Pistoia Reda, 40–71. Basingstoke, Hampshire (UK): Palgrave Macmillan.

Kennedy, Christopher, & Louise McNally. 2005. Scale structure, degree modification and the semantics of gradable predicates. *Language* 81:345–381.

Kennedy, Christopher, & Louise McNally. 2010. Color, context, and compositionality. *Synthese* 174:79–98.

Kinzler, Katherine D., & Elizabeth S. Spelke. 2007. Core systems in human cognition. *Progress in Brain Research* 164:257–264.

Kiss, Katalin É. 2012. Patterns of agreement with coordinate noun phrases in Hungarian. *Natural Language & Linguistic Theory* 30:1027–1060.

Kitagawa, Yoshihisa. 1986. Subjects in Japanese and English. Doctoral dissertation, University of Massachusetts Amherst, Amherst.

Kobele, Gregory M. 2005. Features moving madly: A formal perspective on feature percolation in the minimalist program. *Research on Language and Computation* 3:391–410.

Koopman, Hilda J., & Dominique Sportiche. 1991. The position of subjects. *Lingua* 85:211–258.

Kratzer, Angelika. 1996. Severing the external argument from its verb. In *Phrase structure and the lexicon*, ed. Johan Rooryck & Laurie Zaring, 109–137. Dordrecht: Kluwer Academic Publishers.

Krifka, Manfred. 1990. Boolean and non-Boolean 'and'. In *Papers from the Second Symposium on Logic and Language*, ed. László Kálmán & László Pólos, 161–188. Budapest: Akadémiai Kiadó.

Krifka, Manfred. 1992. A compositional semantics for multiple focus constructions. In *Informationsstruktur und Grammatik*, ed. Joachim Jacobs, 17–53. Wiesbaden: VS Verlag für Sozialwissenschaften.

Krifka, Manfred. 1996. Pragmatic strengthening in plural predications and donkey sentences. In *Proceedings of SALT 6*, ed. Teresa Galloway & Justin Spence, 136–153. Ithaca, NY: Cornell University.

Krifka, Manfred. 1998. Additive particles under stress. In *Proceedings of SALT 8*, ed. Devon Strolovitch & Aaron Lawson, 111–129. Ithaca, NY: Cornell University.

Kripke, Saul A. 2009[1990]. Presupposition and anaphora: Remarks on the formulation of the projection problem. *Linguistic Inquiry* 40:367–386.

Križ, Manuel. 2015. Aspects of homogeneity in the semantics of natural language. Doctoral dissertation, University of Vienna, Vienna.

Križ, Manuel. 2019. Homogeneity effects in natural language semantics. *Language and Linguistics Compass* 13:1–17.

Križ, Manuel, & Benjamin Spector. 2021. Interpreting plural prediction: homogeneity and non-maximality. *Linguistics and Philosophy* 44:1131–1178.

Kuroda, S.-Y. 1988. Whether we agree or not: a comparative syntax of English and Japanese. *Lingvisticae Investigationes* 12:1–47.

Lambert, Karel, & Bas C. van Fraassen. 1970. Meaning relations, possible objects, and possible worlds. In *Philosophical problems in logic*, ed. Karel Lambert, 1–19. Dordrecht: Springer.

Lasersohn, Peter. 1995. *Plurality, conjunction and events*. Dordrecht: Kluwer Academic Publishers.

Lasersohn, Peter. 1999. Pragmatic halos. *Language* 75:522–551.

Lee, Youngjoo. 2004. The syntax and semantics of focus particles. Doctoral dissertation, MIT, Cambridge, MA.

Legate, Julie Anne. 2002. Phases in "Beyond Explanatory Adequacy". MS: MIT.

Legate, Julie Anne. 2003. Some interface properties of the phase. *Linguistic Inquiry* 34:506–515.

Legate, Julie Anne. 2014. *Voice and v: Lessons from Acehnese*. Cambridge, MA: MIT Press.

Levinson, Stephen C. 1983. *Pragmatics*. Cambridge: Cambridge University Press.

Levinson, Stephen C. 2000. *Presumptive meanings: The theory of generalized conversational implicature*. Cambridge, MA: MIT Press.

Lieber, Rochelle. 1989. On percolation. In *Yearbook of Morphology*, ed. Geert Booij & Jaap van Marle, volume 2, 95–138. Berlin: De Gruyter Mouton.

Link, Godehard. 1983. The logical analysis of plurals and mass terms: a lattice theoretical approach. In *Meaning, use, and the interpretation of language*, ed. R. Bäuerle, C. Schwarze, & A. von Stechow, 302–323. Berlin: de Gruyter.

Löbner, Sebastian. 2000. Polarity in natural language: predication, quantification and negation in particular and characterizing sentences. *Linguistics and Philosophy* 23:213–308.

Magri, Giorgio. 2009. A theory of individual-level predicates based on blind mandatory implicatures. Constraint promotion for Optimality Theory. Doctoral dissertation, MIT, Cambridge, MA.

Magri, Giorgio. 2014. An account for the homogeneity effects triggered by plural definites and conjunction based on double strengthening. In *Pragmatics, semantics and the case of scalar implicatures*, ed. Salvatore Pistoia Reda, 99–145. Basingstoke: Palgrave Macmillan.

Malamud, Sophia A. 2012. The meaning of plural definites: A decision-theoretic approach. *Semantics & Pragmatics* 5:1–58.

Marantz, Alec. 1997. No escape from syntax: Don't try morphological analysis in the privacy of your own lexicon. *University of Pennsylvania working papers in linguistics* 4:201–225.

Marantz, Alec. 2001. Words. University of Southern California: Handout from WCCFL 20.

Marantz, Alec. 2007. Phases and words. In *Phases in the theory of grammar*, ed. Sook-Hee Chol, 191–222. Seoul: Dong In.

Margolis, Eric, & Stephen Laurence. 2021. Concepts. In *The Stanford encyclopedia of philosophy*, ed. Edward N. Zalta. Metaphysics Research Lab, Stanford University, Spring 2021 edition.

Matsumoto, Yo. 1995. The conversational condition of Horn scales. *Linguistics and Philosophy* 18:21–60.

McCloskey, James. 1986. Inflection and conjunction in Modern Irish. *Natural Language & Linguistic Theory* 4:245–281.

McCloskey, Jim. 1997. Subjecthood and subject positions. In *Elements of grammar: Handbook in generative syntax*, ed. Liliane Haegeman, 197–235. Dordrecht: Kluwer.

McNally, Louise. 1992. VP coordination and the VP-internal subject hypothesis. *Linguistic Inquiry* 23:336–341.

Meyer, Marie-Christine. 2016. Generalized Free Choice and missing alternatives. *Journal of Semantics* 33:703–754.

Montague, Richard. 1973. The proper treatment of quantification in ordinary English. In *Approaches to natural language*, ed. K. J. J. Hintikka, J. M. E. Moravcsik, & P. Suppes, 221–242. Dordrecht: Springer.

Munn, Alan Boag. 1993. Topics in the syntax and semantics of coordinate structures. Doctoral dissertation, University of Maryland, College Park, Maryland.

Nie, Yining. 2020. Licensing arguments. Doctoral dissertation, New York University, New York.

Oikonomou, Despina. 2016a. Covert modals in root contexts. Doctoral dissertation, MIT, Cambridge, MA.

Oikonomou, Despina. 2016b. Imperatives are existential modals: Deriving the strong reading as an implicature. In *Proceedings of SALT 26*, ed. Mary Moroney, Carol Rose Little, Jacob Collard, & Dan Burgdorf, 1043–1062. Linguistic Society of America and Cornell Linguistics Circle.

Österle, Daniel. 2015. The insertion of 'auch' under negation. B.A. thesis, University of Tübingen, Tübingen.

Ott, Dennis. 2009. The evolution of I-language: Lexicalization as the key evolutionary novelty. *Biolinguistics* 3:255–269.

Paillé, Mathieu. 2021. Exhaustivity and the meaning of colour terms. In *Proceedings of the 38th West Coast Conference on Formal Linguistics*, ed. Rachel Soo, Una Y. Chow, & Sander Nederveen, 334–344. Somerville, MA: Cascadilla Press.

Paillé, Mathieu, & Bernhard Schwarz. 2018. Knowing *whether* and ignorance inferences. In *Proceedings of the 36th West Coast Conference on Formal Linguistics*, ed. Richard Stockwell, Maura O'Leary, Zhongshi Xu, & Z.L. Zhou, 306–315. Somerville, MA: Cascadilla Press.

Parsons, Terence. 1995. Thematic relations and arguments. *Linguistic Inquiry* 26:635–662.

Payne, John. 1993. The headedness of noun phrases: slaying the nominal hydra. In *Heads in grammatical theory*, ed. Greville G. Corbett, Norman M. Fraser, & Scott McGlashan, 114–139. Cambridge: Cambridge University Press.

Perlmutter, David M., & Paul M. Postal. 1977. Toward a universal characterization of passivization. In *Proceedings of the 3rd Annual Meeting of the Berkeley Linguistics Society*, 394–417.

Pietroski, Paul M. 2018. *Conjoining meanings: Semantics without truth values*. Oxford: Oxford University Press.

Pollard, Carl, & Ivan Sag. 1994. *Head-driven phrase structure grammar*. Chicago: University of Chicago Press.

Poortman, Eva B. 2017. Concept typicality and the interpretation of plural predicate conjunction. In *Compositionality and concepts in linguistics and psychology*, ed. James A. Hampton & Yoad Winter, Language, Cognition, and Mind, 139–162. Springer.

Preminger, Omer. 2011. Agreement as a fallible operation. Doctoral dissertation, MIT, Cambridge, MA.

Pustejovsky, James. 1995. *The generative lexicon*. Cambridge, MA: MIT Press.

Quang, Phuc Dong. 1971. A note on conjoined noun phrases. In *Studies out in left field: defamatory essays presented to James D. McCawley on his 33rd or 34th birthday*, ed. Arnold M. Zwicky, Peter H. Salus, Robert I. Binnick, & Anthony L. Vanek, 11–18. Edmonton: Linguistics Research.

Rescher, Nicholas. 1959. The distinction between predicate intension and extension. *Revue philosophique de Louvain* 57:623–636.

Roberts, Craige. 2012[1998]. Information structure in discourse: towards an integrated formal theory of pragmatics. *Semantics & Pragmatics* 5:1–69.

van Rooij, Robert. 2003. Questioning to resolve decision problems. *Linguistics and Philosophy* 26:727–763.

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

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

Rosch, Eleanor. 1978. Principles of categorization. In *Cognition and categorization*, ed. Eleanor Rosch & Barbara Lloyd, 27–48. Hillsdale, NJ: Lawrence Erlbaum Associates.

Rosch, Eleanor, & Carolyn B. Mervis. 1975. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology* 7:573–605.

Rotstein, Carmen, & Yoad Winter. 2004. Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics* 12:259–288.

Russell, Benjamin. 2006. Against grammatical computation of scalar implicatures. *Journal of Semantics* 23:361–382.

Ruys, E. G. 2015. On the anaphoricity of *too*. *Linguistic Inquiry* 46:343–361.

Sæbø, Kjell Johan. 2004. Conversational contrast and conventional parallel: Topic implicatures and additive presuppositions. *Journal of Semantics* 21:199–217.

Sag, Ivan A., Gerald Gazdar, Thomas Wasow, & Steven Weisler. 1985. Coordination and how to distinguish categories. *Natural Language & Linguistic Theory* 3:117–171.

van der Sandt, Rob, & Bart Geurts. 2001. Too. In *Proceedings of the 13th Amsterdam Colloquium*, 180–185. University of Amsterdam.

Sassoon, Galit, & Natalia Zevakhina. 2012. Granularity shifting: Experimental evidence from degree modifiers. In *Proceedings of SALT 22*, 226–246.

Sauerland, Uli. 2002. The present tense is vacuous. *Snippets* 6:12–13.

Sauerland, Uli. 2004. Scalar implicatures in complex sentences. *Linguistics and Philosophy* 27:367–391.

Sauerland, Uli. 2008. Implicated presuppositions. In *Language, context and cognition*, ed. Anita Steube, volume 8: The Discourse Potential of Underspecified Structures, 581–599. Berlin: Mouton de Gruyter.

Sauerland, Uli. 2012. The computation of scalar implicatures: pragmatic, lexical or grammatical? *Language and Linguistics Compass* 6:36–49.

Sauerland, Uli, Jan Anderssen, & Kazuko Yatsushiro. 2005. The plural is semantically unmarked. In *Linguistic evidence: Empirical, theoretical, and computational perspectives*, ed. Stephan Kepser & Marga Reis. Berlin: De Gruyter Mouton.

de Saussure, Ferdinand. 2011[1916]. *Course in general linguistics*. New York: Columbia University Press.

Schein, Barry. 2017. *'And': Conjunction reduction redux*. Cambridge, MA: MIT Press.

Schlenker, Philippe. 2012. *Maximize Presupposition* and Gricean reasoning. *Natural Language Semantics* 20:391–429.

Schmitt, Viola. 2013. More pluralities. Doctoral dissertation, University of Vienna, Vienna.

Schmitt, Viola. 2019. Pluralities across categories and plural projection. *Semantics & Pragmatics* 12:1–49.

Schmitt, Viola. 2021. Boolean and non-Boolean conjunction. In *The Wiley Blackwell companion to semantics*, ed. Daniel Gutzmann, Lisa Matthewson, Cécile Meier, Hotze Rullmann, & Thomas Ede Zimmermann, 1–32. Hoboken, NJ: John Wiley & Sons, Inc.

Schwarzschild, Roger. 1994. Plurals, presuppositions and the sources of distributivity. *Natural Language Semantics* 2:201–248.

Schwarzschild, Roger. 1996. *Pluralities*. Boston: Kluwer Academic.

Sells, Peter. 1985. *Lectures on contemporary syntactic theories*. CSLI lecture notes 3. Stanford, CA: CSLI Publications.

Singh, Raj. 2008a. Modularity and locality in interpretation. Doctoral dissertation, MIT, Cambridge, MA.

Singh, Raj. 2008b. On the interpretation of disjunction: asymmetric, incremental, and eager for inconsistency. *Linguistics and Philosophy* 31:245–260.

Singh, Raj. 2011. *Maximize Presupposition!* and local contexts. *Natural Language Semantics* 19:149–168.

Sivik, Lars, & Charles Taft. 1994. Color naming: A mapping in the NCS of common color terms. *Scandinavian Journal of Psychology* 35:144–164.

Soames, Scott. 1982. How presuppositions are inherited: a solution to the projection problem. *Linguistic Inquiry* 13:483–545.

Soames, Scott. 1989. Presupposition. In *Handbook of philosophical logic*, ed. Dov Gabbay & Franz Guenthner, volume 4, 553–616. D. Reidel.

Speas, Margaret Jean. 1986. Adjunctions and projections in syntax. Doctoral dissertation, MIT, Cambridge, MA.

Spector, Benjamin. 2003. Scalar implicatures: exhaustivity and Gricean reasoning. In *Questions in dynamic semantics*, ed. Maria Aloni, Alastair Butler, & Paul Dekker, volume 17 of *Current Search in the Semantics/Pragmatics Interface*, 225–249. Bingley, UK: Emerald Group Publishing.

Spector, Benjamin. 2013. Homogeneity and plurals: From the Strongest Meaning Hypothesis to supervaluations. Sinn und Bedeutung 18: University of the Basque Country, Vitoria-Gasteiz. Available at: `https://ehutb.ehu.es/uploads/material/Video/3289/Sinn18_01.pdf`.

Spelke, Elizabeth S. 2000. Core knowledge. *American Psychologist* 55:1233–1243.

Spelke, Elizabeth S. 2003. What makes us smart? Core knowledge and natural language. In *Language in mind: Advances in the study of language and thought*, ed. Dedre Gentner & Susan Goldin-Meadow, 277–311. Cambridge, MA: MIT Press.

Spelke, Elizabeth S., & Katherine D. Kinzler. 2006. Core knowledge. *Developmental Science* 10:89–96.

Spelke, Elizabeth S., & Katherine D. Kinzler. 2009. Innateness, learning, and rationality. *Child Development Perspectives* 3:96–98.

Sportiche, Dominique. 1988. A theory of floating quantifiers and its corollaries for constituent structure. *Linguistic Inquiry* 19:425–449.

Stalnaker, Robert. 1981. Anti-essentialism. *Midwest Studies of Philosophy* 4:343–355.

Staniszewski, Frank. 2020. *Until*-phrases as obligatory free choice items. In *Proceedings of Sinn und Bedeutung 24*, ed. Michael Franke, Nikola Kompa, Mingya Liu, Jutta L. Mueller, & Juliane Schwab, volume 2, 271–288.

Steindel Burdin, Rachel, & Joseph Tyler. 2018. Rises inform, and plateaus remind: Exploring the epistemic meanings of "list intonation" in American English. *Journal of Pragmatics* 136:97–114.

Sugawara, Ayaka. 2017. The morpheme *-ish* is a degree head. In *A pesky set: Papers for David Pesetsky*, ed. Claire Halpert, Hadas Kotek, & Coppe van Urk. Cambridge, MA: MIT Working Papers in Linguistics.

Szabolcsi, Anna. 2015. What do quantifier particles do? *Linguistics and Philosophy* 38:159–204.

Szabolcsi, Anna, & Bill Haddican. 2004. Conjunction meets negation: A study in cross-linguistic variation. *Journal of Semantics* 21:219–249.

Talmy, Leonard. 1978. Figure and ground in complex sentences. In *Universals of human language*, ed. Joseph Greenberg, volume 4: Syntax, 625–649. Stanford: Stanford University Press.

Tsarfaty, Reut, & Khalil Sima'an. 2007. Three-dimensional parametrization for parsing morphologically rich languages. In *Proceedings of the 10th Conference on Parsing Technologies*, ed. Harry Bunt & Paola Merlo, 156–167. Prague: Association for Computational Linguistics.

Wagner, Michael. 2005. Prosody and recursion. Doctoral dissertation, MIT, Cambridge, MA.

Wagner, Michael. 2006. Givenness and locality. In *Proceedings of SALT 16*, ed. Jonathan Howell & Masayuki Gibson, 295–312. Ithaca, NY: CLC Publications.

Wagner, Michael. 2012. Focus and givenness: a unified approach. In *Contrasts and positions in information structure*, ed. Ivona Kučerová & Ad Neeleman, 102–147. Cambridge: Cambridge University Press.

Williams, Alexander. 2015. *Arguments in syntax and semantics*. Cambridge: Cambridge University Press.

Williams, Edwin. 1978. Across-the-board rule application. *Linguistic Inquiry* 9:31–43.

Winter, Yoad. 1998. Flexible Boolean semantics: Coordination, plurality and scope in natural language. Doctoral dissertation, Utrecht University, Utrecht.

Winter, Yoad. 2001. *Flexibility principles in Boolean semantics*. Cambridge, MA: MIT Press.

Winter, Yoad. 2019. On presupposition projection with trivalent connectives. In *Proceedings of SALT 29*, 582–608.

Winter, Yoad, & Remko Scha. 2015. Plurals. In *The handbook of contemporary semantic theory*, ed. Shalom Lappin & Chris Fox, 77–113. Hoboken, NJ: John Wiley & Sons, 2nd edition.

Wągiel, Marcin. 2021. *Subatomic quantification*. Berlin: Language Science Press.

Yoon, Youngeun. 1996. Total and partial predicates and the weak and strong interpretations. *Natural Language Semantics* 4:217–236.

Zagona, Karen T. 1982. Government and proper government of verbal projections. Doctoral dissertation, University of Washington, Seattle.

Zeevat, Henk. 1992. Presupposition and accommodation in update semantics. *Journal of Semantics* 9:379–412.

Zeevat, Henk. 2002. Explaining presupposition triggers. In *Information sharing: reference and presupposition in language generation and interpretation*, ed. Kees van Deemter & Rodger Kibble, 61–87. Stanford: CSLI Publications.

Zhang, Niina Ning. 2010. *Coordination in syntax*. Cambridge: Cambridge University Press.