Linguistics and Philosophy
https://doi.org/10.1007/s10988-024-09425-5

ORIGINAL RESEARCH

An interaction between logical vocabulary and predicate meanings

Mathieu Paillé[1]

Accepted: 2 October 2024

Abstract
Predicates within many conceptual classes are intuited as mutually exclusive. Based on these predicates' interaction with logical vocabulary like *and* or *also*, however, this paper argues that they are in fact underlyingly consistent; the strong intuited meanings arise from semantic exhaustification. In addition to demonstrating that exhaustification is more widespread than previously believed, this paper also shows that this particular exhaustification effect behaves in a hitherto undescribed manner. Indeed, a predicate's exhaustification is always computed locally at the level of the predicate, rather than the clause or sentence containing it.

Keywords Predicates · Exhaustification · Additive particles · Conjunction

1 Introduction

Is there a one-to-one correspondence between the lexical meaning of a word and its meaning as intuited in actual sentences? The standard answer to this question is that, for logical/functional vocabulary, this is often not the case. Take *most*, which is intuited as "anti-universal" (i.e., excluding a universal meaning) in many simple sentences:

(1) Arwa ate most of the apples.
$\rightsquigarrow$ Arwa did not eat all of the apples.

One might take this to show that *most* is lexically anti-universal, but this would not account for the fact that its anti-universality readily disappears in downward-entailing (DE) environments:

(2) If Arwa eats most of the apples, her dog barks.
$\not\rightsquigarrow$ Arwa's dog barks if she eats most but not all of the apples.

✉ Mathieu Paillé
mathieu.paille@mail.mcgill.ca

[1] University of Calgary, 2940 University Way NW, Calgary, AB T2N 4H5, Canada

Published online: 09 January 2025

For this reason, it is often taken that *most* is inclusive of universal meanings lexically, as observed in (2), but undergoes strengthening in sentences like (1). Asserting a sentence containing *most* comes with the negation of an alternative sentence obtained by replacing *most* with *all* (see e.g., Horn, 1972; Gazdar, 1979; Levinson, 1983; Blutner, 2002; 2004; Spector, 2003; Chierchia, 2004; 2006; van Rooij & Schulz, 2004; Sauerland, 2004; Russell, 2006; Fox, 2007; Geurts, 2010; Chierchia et al., 2012; Sauerland, 2012; Bar-Lev & Fox, 2017; Fox & Spector, 2018).

On the other hand, with a few exceptions to be discussed in Sect. 2, strengthening through the negation of alternatives has not been postulated for content vocabulary—lexical items like *teaspoon*, *tree*, *animated*, or *sophistication*. Unlike *most*, the meaning of *teaspoon*, as intuited in a sentence like (3), is assumed to match its underlying lexical–conceptual meaning.

(3) This is a teaspoon.

On the standard view, the extension of *teaspoon* in (3) is the same as the extension of the lexical item *teaspoon*: the set of teaspoons, as determined by the concept TEASPOON. No strengthening is postulated.[1] Not positing any strengthening is apparently well-motivated from the lack of observable differences between the meaning of *teaspoon* in (3) and in DE environments:

(4) If you use a teaspoon, make sure not to spill anything.

It is not as if *teaspoon* had become inclusive of tablespoons, or other artifacts like shoes or guitars, in (4)—its meaning has not changed from (3).

The present paper shows that the standard view is incorrect. I show that content vocabulary interacts with some logical vocabulary in a non-trivial way, and suggest an analysis that involves content vocabulary being strengthened in most (I will argue all) sentences. To see this, consider sentences that are intuited as contradictory due to the meaning of two predicates:

(5) #This comedy is a tragedy.

While it could be taken that *comedy* and *tragedy* are simply lexically exclusive of one another, this cannot be right. Indeed, it is possible to use logical vocabulary like *and* and *also* to remove the mutual exclusivity of such lexical items:

(6) a. This is both a comedy and a tragedy.
b. This comedy is also a tragedy.

As I argue, (6) demonstrates that predicates like *comedy* and *tragedy* are lexically weak: lexically, *comedy* means that something meets the criteria for being a comedy, without ruling out that it also meets the criteria for other genres. The intuition of a contradiction in (5) comes from the predicates being strengthened to be exclusive of other genre predicates.

[1] Some authors posit context-dependency for some predicate meanings (including, e.g., Waldon et al., 2023 for artifact predicates like *electronics*), but not strengthening through the negation of alternatives.

Having established that predicates undergo strengthening, we will see that we can learn an important lesson about strengthening by observing it in this new empirical domain. The strengthening of predicates is always computed ultra-locally, unlike previously studied effects like the strengthening of *most*. Modelling strengthening through the Exh(aust) operator of Chierchia et al. (2012), I show that the strengthening of predicates like *comedy* does not apply to clauses/propositions as standardly described: rather, with predicates, strengthening applies to individual lexical items. I call the standard exhaustification that applies to propositions (and is "free" rather than syntactically constrained) "φ-exhaustification", and the ultra-local exhaustification of predicates proposed in this paper "ω-exhaustification". Part of the motivation for this comes again from (5). If Exh could take sentential scope, its prejacent (the constituent it strengthens) would entail that the individual is both a comedy and a tragedy. Neither predicate would be strengthened to exclude the other, because Exh does not exclude alternatives that are entailed by its prejacent. Rather, to obtain a contradiction, at least one (I will claim both) must be strengthened independently of the other:

(7) #This [ω-Exh_{ALT} comedy] is a [ω-Exh_{ALT} tragedy].

The twin findings of this study are therefore both that the meaning of predicates involves strengthening, and that this strengthening behaves in a novel way, mimicking lexical meaning in being both obligatory with and ultra-local to the strengthened expression.

I begin the paper in Sect. 2 by elaborating on the few cases where predicates are strengthened on the standard view. Then, in Sect. 3, I describe in what sense predicates are intuited as "strong" in sentences. In Sect. 4, I show that this strength cannot be a fact of predicates' lexical–conceptual meanings, due to their interaction with logical vocabulary like *and* and *also*. I therefore suggest in Sect. 5 that the strong meaning is the result of grammatical exhaustification. In Sect. 6, I show that this exhaustification effect is obligatory and ultra-local, as just described. Section 7 discusses cases where weaker predicate meanings appear even in the absence of *and* or *also*; since this weakness is conditioned by discourse factors (rather than interaction with logical vocabulary), I explain them by borrowing a mechanism previously proposed for discourse-conditioned weakness in plural predication. Finally, Sect. 8 shows how the ω-exhaustification of predicates is distinct from contrastive (intonationally marked) focus on predicates; the claim in this paper is not that contrastive focus occurs more frequently than previously thought, but that some strengthening of predicates occurs without them being contrastively focused (much like *most* is usually anti-universal without being intonationally focused). Section 9 concludes.

2 Well-established cases of strengthening in predication

Let's start with a more accurate version of what I claimed to be the standard view of predicates' meanings. There are two areas in which predicates are normally taken to be intuited differently from their lexical meanings due to strengthening: when they are part of an entailment ("Horn") scale, and when they are intonationally focused.

Horn scales are sets of predicates where some asymmetrically entail others (Horn, 1972). A frequent example is {*warm*, *hot*}. Assume that lexically, these predicates lack an upper bound; informally, *warm* means 'at least hot to degree n' and *hot* means 'at least hot to degree m ($m > n$)'. In many actual sentences, however, claiming that something is *warm* implies that it does not meet the threshold for being *hot*:

(8) The soup is warm.
$\rightsquigarrow$ It is not hot.

This can be captured as a strengthening effect, whereby an alternative obtained by replacing *warm* with *hot* is negated, just as we saw in (1) for *most*/*all*. This is because *hot* asymmetrically entails *warm*; since (9b) is stronger than (9a), the assertion of (9a) comes with the negation of (9b).

(9) a. The soup is warm.
b. The soup is hot.

A second case where the intuited meaning of predicates involves strengthening is when they are contrastively focused—intonationally emphasized, with this emphasis indicating a semantic contrast with another predicate. For instance, (10) suggests that the speaker thinks there is some other salient predicate (e.g., *tea*) that could have been thought to be true of the substance at hand.

(10) This is SOUP$_F$.
$\rightsquigarrow$ This is not tea.

One context in which (10) could be used is as a correction to someone claiming the liquid is tea, for instance. One property of contrastive focus that we will return to is that when the contrastively focused predicate is within a definite DP, there is an entailment about another individual altogether:

(11) The SOUP$_F$ is hot.
$\rightsquigarrow$ The tea is not hot.

With *tea* as an alternative to *soup*, (11) means the unique soup is hot, and the unique tea is not. As with Horn scales, contrastive focus on predicates is a strengthening effect, arising from the negation of alternatives like *This is hot* or *This is tea*.

This article argues that predicates are systematically intuited as stronger than their lexical meanings, even when they are not part of a Horn scale or contrastively focused. To see this, we first have to observe that predicates are mutually exclusive of one another more often than we might expect.

3 Cotaxonyms' strong intuited meanings

This section discusses the kind of data that lets us observe predicates' strength, and briefly lays out theories of predicates/concepts that predict this strength as a fact of predicates' lexical–conceptual meaning.

3.1 Observing 'cotaxonymic exclusivity'

Taxonymy refers to the 'kind of' relation; cotaxonyms are sisters in a taxonomy, sharing a superordinate (Cruse, 2002). For instance, *poodle* is a taxonym of *dog*, and *poodle* and *labrador* are cotaxonyms. A fact about cotaxonyms that may not initially look particularly interesting is that they are mutually exclusive, as seen in (12), for example.

(12) #That sheep is a horse. (Cann, 2011, p. 459)

Call sentences like (12) "co-predications": they involve a single individual having two predicates from a particular class (here animals) assigned to it.[2]

At first glance, (12) does not appear to teach us much about language. Speakers have the world knowledge that biologically, animals cannot be two species at once. What we will now see, however, is that cotaxonyms are in fact always mutually exclusive, even when we control for world knowledge as a possible explanation. I therefore conclude that cotaxonyms are always mutually exclusive due to their being cotaxonyms; call this effect COTAXONYMIC EXCLUSIVITY.

3.1.1 Predicates whose meanings might have been inclusive

To observe cotaxonymic exclusivity, consider examples like (13) rather than (12).

(13) 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. #This train is a plane.

In contrast to (12), world knowledge is not a factor in these examples (although (13f) needs a science-fiction scenario). It is possible for objects to have white parts and green parts (13a) or for a film to be partly animated and partly live-action (13b), but (13a) and (13b) have cotaxonyms that quantify universally over the parts of the flag or film, creating a contradiction. It could have been otherwise; the predicates could have only modified *some* parts of the flag or film, resulting in the meaning that the flag has both colours and the film has animated parts and live-action parts. As for (13c) and (13d), these cannot be used about things for which English has lexicalised blends, viz. *tragicomedy* and *spork*. One could try to write this up to competition in lexical choice: it is strange to describe a utensil as a 'fork that is a spoon' when the word *spork* exists. But in fact, examples like (14) show that the existence of the

[2] Another phenomenon has also been called a "co-predication" in the literature, namely the "dot objects" of Pustejovsky (1995):

(i) The books are dusty but interesting.

In (i), *the books* refers both to the books qua information and the books qua physical object. This kind of "co-predication" is another phenomenon altogether.

words *tragicomedy* and *spork* is not a sufficiently general explanation for the semantic deviance in (13c)/(13d), because these co-predications cannot be used even to teach someone what a tragicomedy or spork is:

(14) a. #A tragicomedy is a comedy that is a tragedy.
b. #A spork is a fork that is a spoon.

Similar comments hold for (13e) and (13f): one can imagine shared jurisdictions or bimodal vehicles, but these cannot be described using the above co-predications.

The co-predications in (13) are not just semantically strange; they are specifically contradictions. They give rise to a similar intuition as $p \wedge \neg p$: *This comedy is a tragedy* entails both that the text is a comedy and that it is not a comedy, due to being a tragedy. To be sure that this contradictory effect is due to cotaxonymy, we can attempt sentences similar to those in (13) but with predicates that are not cotaxonymic. This is what the sentences in (15) do, and the contradiction effect disappears. (15a–d) mix and match predicates from different taxonomies in (13) (colour terms, genres, etc.); (15d–f) (maybe also (15c)) intentionally attempt unlikely combinations, still not reaching the contradiction effect of (13).

(15)

a.	This train is green.	cf. (13a) and (13f)
b.	Some animated films are tragedies.	cf. (13b) and (13c)
c.	??Some forks are provincial.	cf. (13d) and (13e)
d.	??Some forks are tragedies.	cf. (13c) and (13d)
e.	This spider is an accountant.	
f.	??Green ideas sleep.	(cf. Chomsky, 1957)

The acceptability of (15e) highlights how easy it is to jump into "cartoon mode", pointing out again the limitation of using world knowledge to explain the contradictory nature of (13) and maybe even (12): why can I imagine a cartoon where a spider is an accountant, but not one where a sheep is a horse (in the sense of being both, not of having been transformed from one to the other)? The meaning of some of the sentences in (15) is highly obscure, to be sure, for reasons such as presupposition failure in (15d) (*tragedy* is presumably only defined for stories/events)—but they are not contradictory like (13). They are marked with ?? instead of # to emphasize this.

In sum, cotaxonyms are interpreted as mutually exclusive, even when we control for world knowledge as a possible source of mutual exclusivity. Within particular conceptual domains like genres or colour adjectives, it appears that lexical meanings are free of overlap. In this sense, cotaxonyms are strong.

3.1.2 Pragmatic complications in the empirical picture

Sentences like those in (13) are not impossible in every discourse context or on every intended meaning. There are some aspects of language use and meaning that make co-predications of cotaxonyms appear possible, but they do not in fact go against the generalization of cotaxonymic exclusivity.

First, arguably metalinguistic uses of predicates make it possible to co-predicate cotaxonyms:

(16) a. This "comedy", as you called it, is (actually) a tragedy.

(17) a. SCENARIO: *A and B are looking at a number of flags; A is colour-blind and confuses blues and greens. A points to an entirely green flag and says:*
b. A: The blue flag has a funny shape.
B: It does. By the way, the blue flag is (actually) green.

Such uses do not disprove that cotaxonyms are mutually exclusive. In (17), for instance, *blue* means 'only blue' and *green* means 'only green'; B's statement is that 'the only-blue flag is in fact not blue at all, but only green'. We can even replicate discourses like (16)–(17) with non-cotaxonymic semantically exclusive expressions, such as *exactly n* numerals:

(18) a. SCENARIO: *A and B are looking at drawings of fantasy creatures with various numbers of legs. A suffers from pathologically poor number sense; they point to an animal with three legs and say:*
b. A: The four-legged creature has a funny face!
B: It sure does. By the way, the four-legged creature actually only has three legs.

In the same vein, co-predicated cotaxonyms are not intuited as mutually exclusive when one cotaxonym 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. (19) is an example with cotaxonyms that are true at different times.

(19) a. SCENARIO: *A formerly entirely white shirt has emerged from the wash fully green.*
b. The white shirt is green.

(19b) can be analyzed as follows, with t_1 being some time in the past, and t_0 including the present.

(20) The [white t_1] shirt is [green t_0].

Likewise, (21) involves cotaxonyms being true in different worlds (w_0 is the world of utterance and w_1 is the world of a play).

(21) a. SCENARIO: *We are setting up a play and decide to represent the character of a fox with a cat.*
b. The cat is the fox.
c. The [cat w_0] is the [fox w_1].

None of these complications should distract from the central point: if they are predicated of the same individual at the same time in the same world, cotaxonyms are mutually exclusive, even when we control for world knowledge.

For these reasons, finding instances of sentences I claim to be contradictory in a corpus does not necessarily counter my claim of contradiction. For example, a Google search of `"this comedy is a tragedy"`, which I described above as contradictory, does in fact give a few results. But they are not examples involving a

neutral claim that the work of literature is a tragicomedy. Concretely, Google (May 23, 2024) lists out twelve pages for `"this comedy is a tragedy"` (five of these are from my own work). To just focus on one example, the first page listed that is not a text written by me is an amateur film review (https://www.imdb.com/review/rw2273872/) about the 1942 film *Tish*:

(22) 4/10

> **This comedy is a tragedy**
> 4 July 2010
>
> It's nice to see the three great characters actresses, but they are given very little to work with. Marjorie Main's is the only developed character, and she seems miscast in it. Fine production values, to be sure, but this film is a mess from beginning to end. The script desperately needed many more re-writes […]

In the title of the review, the author is claiming that the film is a comedy as far as its intended genre is concerned; the author then uses the predicate *tragedy* to say something quite different, namely that the movie is bad—"a mess". Whatever the right analysis of the pragmatics of (22), the author is using the cotaxonyms as mutually exclusive: the genre of the film is that it is a comedy (and not a tragedy), while its storytelling is tragic (and not comedic).[3]

3.1.3 Cotaxonymic exclusivity outside of co-predications

So far, we have only observed cotaxonymic exclusivity in co-predications. What about simpler sentences like those in (23)? Can we still observe the strong meaning of cotaxonyms there?

(23) a. This is a comedy.
b. This is green.

I give these two examples because (23b) involves a predicate that quantifies over parts, while (23a) does not: (23b) means that all parts of the subject are green, while (23a) does not mean that all parts of the subject are comedies. Löbner (2000) calls predicates

[3] Ironically, Wikipedia (s.v. '*Tish* (film)', accessed May 23, 2024) claims that *Tish* is a tragicomedy; but that is not what the writer of (22) is attempting to say.

Another example similar to (22) is (i), part of a text on Bernard Shaw's *The Doctor's Dilemma*.

(i) At one level this comedy is a tragedy dealing with allocation of scarce medical resources. (https://www.amazon.ca/Doctors-Dilemma-Getting-Married-Shewing/dp/1417925248, accessed September 3, 2024)

As far as I can tell, the meaning here is that (i) the speaker generally takes the play to be a comedy (and no other genres), but (ii) by taking a particular perspective on the play (i.e., focusing on certain aspects of it) suddenly the speaker judges it not to be a comedy at all, but in fact a tragedy. Here too, then, we have a felicitous co-predication of cotaxonyms, but with the predicates once again holding of the individual in different ways.

like *comedy* INTEGRATIVE and predicates like *green* SUMMATIVE. It is only with the summative (23b) that we can straightforwardly observe that cotaxonymic exclusivity is maintained even in non-co-predicational environments: (23b) means that *all* parts of the subject are green, so they cannot be of other colours. As for (23a), things are not immediately so clear. The sentence certainly suggests that the text is not a tragicomedy, but from this example alone, this intuition could be a weak prototypicality effect. That is, (23a) could be taken to weakly suggest that the text is not a tragicomedy in the same way that (24) weakly suggests that the animal is not blue, for example.

(24) This is a dog.

The assumption upon hearing (24) that the animal is not blue is obviously not a semantic entailment of the sentence. One could try to say something similar for (23a), arguing that (23a) suggests that the text is not tragicomedic simply because comedies are prototypically not tragicomedic. This line of reasoning is probably wrong: the non-blueness of the animal in (24) is defeasible, but the non-tragic character of the text in (23a) is not:

(25) a. This is a dog. Strangely, it's blue.
b. This is a comedy. #Strangely, it's a tragicomedy.

(26) a. The dog is blue.
b. #The comedy is a tragedy.

Thus, the "anti-tragic" inference in (23a) must be due to cotaxonymic exclusivity being present with integrative predicates even outside of co-predications.

3.2 Approaches that derive cotaxonymic exclusivity

What, then, underlies cotaxonymic exclusivity? The structuralist linguistics of Ferdinand de Saussure (1916), a classic and formerly influential proposal, would derive this exclusivity as a fact of the lexicon; more recent approaches in cognitive science would derive it as arising from the layout of concepts.

De Saussure viewed the lexicon as forming a "structure" (a geometry) in which words' meaning is delimited by other adjacent words: "the value of each term results … from the simultaneous presence of others" (de Saussure, 1916, pp. 114-115), as in his famous diagram:

What matters for us is not the distinction between signified and signifier, but the point that lexical items exist in a partitioned space: there is no overlap between the lexical items' meanings, as indicated by the two-sided arrows. De Saussure's famous example came close to involving cotaxonyms: he pointed out that French has a single word, *mouton*, for sheep and their meat, whereas English has a word for each concept, *sheep* and *mutton*, and *sheep* cannot be used to refer to the meat. His idea was that

the concept of a sheep's meat could not be covered by *sheep* precisely because that space in the domain of animal-concepts is taken up by *mutton*. Thus, the meaning of *sheep* in English is stronger than the meaning of *mouton* in French, because *sheep* semantically excludes *mutton*. Generally, de Saussure's "structuralist" claim is that the lexical meanings of words exclude one another due to the nature of lexicalization itself.

Some caveats are needed for de Saussure, of course, to deal with content vocabulary where no exclusion is observed, between including taxonyms and their superordinates (like *sheep* and *animal*) and vocabulary from different conceptual domains (like *sheep* and *white*). Either way, de Saussure's idea would extend generally to predicates, deriving cotaxonymic exclusivity as a fact of the lexicon. There are more recent proposals that accomplish roughly the same thing. Indeed, it is a claim of both the prototype and exemplar theories of concepts that concepts are represented in geometric mental spaces which are inherently partitioned between different concepts (e.g., Smith et al., 1997; Gärdenfors, 2000; Blair & Homa, 2001; Ruts et al., 2004; Hampton, 2015; Levering et al., 2020; see also van Fraassen, 1967; Lambert & van Fraassen, 1970; Stalnaker, 1981). Prototype and exemplar theories involve competing claims on the nature of this delimitation between concepts, but not on its existence. While these theories and de Saussure's diverge in whether partitioning occurs in the conceptual space itself or in the lexicalization of this space, they agree on there being partitioning; cotaxonymic exclusivity follows from this.

4 The weak lexical meaning of predicates

I now turn to showing that the above description, while accurate for the interpretation of cotaxonyms in most sentences, is not an accurate characterization of their lexical meanings. If cotaxonyms' exclusivity is present in their lexical–conceptual meaning, we expect that it should always be observed; cotaxonyms would have an empty intersection lexically. One could only co-predicate cotaxonyms non-contradictorily using a limited set of expressions like the non-intersective adjective *wannabe* (27) or in pragmatically "marked" environments like those described in Sect. 3.1.2.

(27) This comedy is a wannabe tragedy.

Is cotaxonymic exclusivity really so persistent? Consider again the examples in (28), repeated from (13). I give these two examples to cover both summative and integrative predicates—a distinction that will occasionally be useful to make in this paper.

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

In fact, language comes with several lexical items that are both intersective in nature and capable of removing the contradictions in (28); I focus on *and* and *also*. I take them in turn in Sects. 4.1 and 4.2 respectively, since both require substantial commentary. The successful co-predication of cotaxonyms with these intersective logical

expressions will teach us that cotaxonyms are weak: *comedy* is lexically inclusive of tragicomedies and *white* lexically quantifies existentially rather than universally over parts. I elaborate briefly on these consequences for lexical meaning in Sect. 4.3. Then, in Sect. 5, I argue that the strong meanings observed in most sentences are due to exhaustification.

4.1 Cotaxonyms and conjunction

The first expression that lets us see cotaxonyms' lexically consistent meanings is *and*:[4]

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

If one assumes that *and* is intersective, (29) shows that *white*/*green* and *comedy*/*tragedy* cannot be lexically exclusive of one another.

But that is a big "if"; *and* has often been claimed to be non-intersective. In the following subsections, I claim that, when it conjoins two predicates (*white and green*, in (29a)), *and* is in fact necessarily intersective when the subject (*the flag* in (29a)) is atomic. This is in contrast to cases where the conjoined predicates are given a plural subject (e.g., *the flags*); in such cases, *and* is not necessarily intersective. All that matters for our purposes is the conclusion that *and* is intersective in (29).

4.1.1 A non-intersective 'and' with plural subjects

Let's start by observing that when conjoined predicates have a plural argument, *and* can, descriptively speaking, be given either an intersective or a non-intersective interpretation:

(30) a. The trees are coniferous and green.
⇒ *the most salient interpretation is* INTERSECTIVE*: 'All the trees are coniferous and all the trees are green.'*
b. The trees are 10 years old and 60 years old.
⇒ *the only non-contradictory interpretation is* NON-INTERSECTIVE*: 'Some of the trees are 10 years old, the rest are 60 years old.'*

To avoid needing to theorise about the best way to capture that *and* can be interpreted as either intersective or non-intersective with a plural subject (30), let's just assume incorrectly (e.g., Schmitt, 2021) that *and* is lexically ambiguous (31).[5] On this view, speakers choose with which *and* to conjoin predicates based on some notion of naturalness or contradiction-avoidance when two predicates are mutually exclusive.[6]

[4] See Harnish (1976), Levinson (1983), Krifka (1990), Lasersohn (1995), and Winter (2001) for prior discussion of conjoined colour terms like (29a).

[5] For work claiming that there is a single *and* that is underlyingly intersective, see Winter (2001), Champollion (2016), and Schein (2017). For work claiming that there is a single *and* that is underlyingly non-intersective, see Krifka (1990), Heycock and Zamparelli (2005), and Schmitt (2013, 2019).

[6] 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.

(31) 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)

The question now is whether, for the examples in (29) with an atomic subject, it is descriptively possible to claim that *and* is non-intersective. After all, there is a parallel between (30b) and (29a): (30b) predicates each conjunct of some proper part of the subject *the trees*, and likewise (29a) predicates each conjunct of some proper part of the subject *the flag*. I emphasize this parallel by repeating the examples with similar paraphrases:

(32) a. The flag is white and green.
≈ 'Some part of the flag is white, the rest is green.'
b. The trees are 10 years old and 60 years old.
≈ 'Some part of the plurality of trees is 10 years old, the rest is 60 years old.'

I now turn to arguing that this parallel is only apparent, and *and* is in fact intersective in (29).

4.1.2 'And' is intersective with summative predicates—when the subject is atomic

I give two reasons why describing (32a) as non-intersective is not tenable (*pace* Krifka, 1990).

First, when *and* conjoins predicates, adding *both* results in the non-availability of a non-intersective interpretation (Schwarzschild, 1996, p. 149). To see this, consider (33). It adds *both* to the two conjunctions with a plural argument in (30), and only the potentially intersective (30a) remains felicitous.

(33) a. The trees are (both) coniferous and green.
b. The trees are (#both) 10 years old and 60 years old.

Thus, if (32a) was only consistent due to *and* being non-intersective, it should not be possible to add *both*. But *both* is entirely acceptable:

(34) The flag is both white and green.

Therefore, an intersective interpretation of *and* must be possible in (32a). This means that the predicates *white* and *green* are lexically consistent.[7]

We can take this argumentation one step farther, to show that *and* in (32a) *only* has an intersective interpretation. If a non-intersective interpretation of *and* was available, it should be possible to make the predicates *white* and *green* mutually exclusive explicitly. To do this, we can use modifiers like *completely*. I assume the following meaning for *completely*:

(35) $[\![\text{completely}]\!] = \lambda P.\lambda x.\forall y[y \sqsubseteq x \rightarrow P(y)]$.

[7] Note that *both* was already present in the *comedies and tragedies* example in (29b), showing that this conjunction is intersective too. I return to (29b) in Sect. 4.1.3.

With a plural subject, it is possible to modify conjoined summative predicates with *completely*, as expected if such conjunctions can be non-intersective:

(36) The flags are completely white and completely green.
≈ 'Some of the flags are completely white, the rest are completely green.'

A single flag cannot be completely white as well as completely green, but the availability of a non-intersective *and* in (36) means that these two conjuncts are predicated of different subsets of flags. In contrast, with an atomic subject as in (29a), adding *completely* is not possible:

(37) #The flag is completely white and completely green.

A non-intersective *and* in (37), if available, would have given it the consistent meaning that some part of the flag is completely white, and some other part is completely green. We conclude that not only *can* the conjunction in (32a) be interpreted intersectively, in fact it *must*.

At first glance, one could take issue with this argument from *completely* by stating that (37) and (36) are not actually parallel: the right parallel for (37) with plural predication would be with floating *each* or *all*, which quantify over individuals rather than subatomic parts. By this reasoning, (37) should be compared not with (36), but with (38)—where a contradiction is in fact observed.

(38) #The flags are each completely white and each completely green.

In fact, even with a non-Boolean *and* available, (38) is only expected to be consistent under certain assumptions, which I argue are incorrect. In particular, *each* would have to compose with the conjunct first (rather than with the subject), so that *each completely white* and *each completely green* are constituents whose meaning involves universal quantification over atomic parts (39a). A non-intersective *and* (the and_2 defined in (31b)) would then predicate each universal predicate of some part of the plural subject (39b). In (39a), subscript AT stands for 'atomic'—'$Part_{\text{AT}}(x)$' is the set of atomic parts of x, and '$\sqsubseteq_{\text{AT}}$' is the 'atomic-part-of' relation.

(39) a. $[\![\text{each } P]\!] = \lambda x : |Part_{\text{AT}}(x)| > 1.\forall y[y \sqsubseteq_{\text{AT}} x \rightarrow [\![P]\!](y)]$.
b. $[\![\text{each completely white and}_2 \text{ each completely green}]\!]$
$= \lambda x.\exists x', x''[x = x' \oplus x'' \wedge [\![\text{each completely white}]\!](x')$
$\wedge [\![\text{each completely green}]\!](x'')]$.

On the assumptions in (39), we have no explanation for why (38) is contradictory. This raises the possibility that there is something wrong with my test, and whatever explains the deviance of (38) might also explain the deviance of (37).

But the assumptions just laid out are incorrect. Floating *each* has independently been argued to necessarily take as its first argument not a predicate, but a *pro* DP co-indexed with the subject (the "associate") (Doetjes, 1997; Fitzpatrick, 2006). This view is motivated by a restriction against A′-movement of the associate, agreement with the associate on the floating quantifier in languages like French, and overt clitic doubling of the associate on the floating quantifier in languages like Hebrew (Fitzpatrick, 2006).

Concretely, Fitzpatrick (2006, p. 81) gives (40a) the LF in (40b) (and assumes the same syntax for *each*).

(40) a. The students will have all had lunch.
b. [$_{DP}$ The students] λ_1 will have [$_{vP}$ [all pro_1] λ_2 [$_{vP}$ t_2 had lunch]].

In (38), then, each token of *each* has as its first argument a *pro*, each of which is co-indexed with the associate *the flags*. This means that each conjunct carries the entailment that it is every individual flag that is completely green/white. Thus, the possibility of a non-Boolean *and* with plurals will not be observed when there is a floating *each* on each conjunct. In contrast to *each*, *completely* (35) takes as its first argument a predicate rather than a *pro*, so if a non-intersective *and* was available in examples like (41), it should be possible to modify the colour terms with *completely*.

(41) The flag is white and green.

To summarise, this section has used data with *both* and *completely* to motivate that conjunctions like (41) are intersective. This shows that *white* and *green* must have a non-empty intersection lexically.

4.1.3 'And' with integrative predicates: no parallel with conjunctions with plurals

We have just seen that the non-contradictory conjunction of summative predicates (with an atomic subject) is evidence of overlap in the meanings of these predicates. What about integrative ones? The relevant examples are repeated below; (42) shows conjoined integrative predicates, and (43) shows intersective and non-intersective conjunctions with plurals.

(42) This play is both a comedy and a tragedy.

(43) a. The trees are coniferous and green.
⇒ *the most salient interpretation is* INTERSECTIVE*: 'All the trees are coniferous and all the trees are green.'*
b. The trees are 10 years old and 60 years old.
⇒ *the only non-contradictory interpretation is* NON-INTERSECTIVE*: 'Some of the trees are 10 years old, the rest are 60 years old.'*

With integrative predicates, unlike summative predicates, there is in fact no parallel whatsoever to make between (42) and (43b). While (43b) (like (32a)) can be described in terms of *and* dividing the subject into two parts, with one conjunct predicated of each part, this is not what is happening in (42), where both predicates describe the subject as a unified individual. If (42) is non-intersective, then, it would be non-intersective in a manner entirely different from (43b), lacking reference to part-structure. The non-intersective and_2 in (31b), repeated in (44), will not do.

(44) $[\![\text{and}_2]\!] = \lambda P.\lambda Q.\lambda x.\ \exists x', x''[x = x' \oplus x'' \wedge P(x') \wedge Q(x'')]$.

Can it be maintained that there is another kind of non-intersective conjunction? This putative non-intersective conjunction, and_3, would have to mean that the argument

is in a set of individuals which are in the extension of neither conjunct (i.e., neither comedies nor tragedies), but in an altogether different set (the set of tragicomedies) whose members merely resemble the members of the extension of the conjuncts. That is, and_3 would make reference not to part-quantification but to a resemblance relation. However, postulating such an and_3 immediately makes incorrect predictions. We would now expect sentences like (45) to be acceptable, for example, because platypuses resemble both ducks and beavers while being neither:

(45) #A platypus is a duck and a beaver.

In sharp contrast to (42), (45) is false. Thus, stipulating the existence of a non-intersective and_3 based on conjunct-resemblance overgenerates. The conjunction in (42) must be intersective, which means that the cotaxonyms *comedy* and *tragedy* are lexically consistent.

4.2 Cotaxonyms and additives

Conjunctions are not the only intersective expressions that make the co-predication of cotaxonyms consistent. The second such expression is *also*:

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

Note that these sentences are intended as involving clause-internal uses of *also*: *also* is not referring to previous sentences in the discourse, but rather refers anaphorically to the clause-internal material *white*/*comedies*.

Luckily, *also* is universally agreed to always be intersective (47), and we therefore do not need to investigate *also* in the same way as we did *and*.

(47) a. The trees are coniferous. They are also green.
b. The trees are 10 years old. #They are also 60 years old.

What does require commentary is the judgment. While (46b) is accepted by everyone, (46a) is not as popular to mean that the flag has a white part and a green part. But all speakers who have voiced to me that they find (46a) questionable have said that it is *better* than the sentence without *also* (28a). More to the point, (46a) can be tweaked in ways that are irrelevant to my argumentation in order to make the sentence more clearly acceptable. In (48), I both give some context to the utterance and change the definite singular *the white flag* to a partitive plural; the result is straightfowardly acceptable.

(48) a. SCENARIO: *We are at a plant that specialises in recycling cloth, and 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.

The more general point is that co-predications where one predicate is in the subject and the other predicate is in the VP generally require some pragmatic reason for the predicates to be syntactically distributed in this way; the scenario in (48) is one way to do this.

Beyond *and* and *also*, there are other expressions permitting the consistent co-predication of cotaxonyms, such as *simultaneously*:

(49) A tragicomedy is a comedy that is **simultaneously** a tragedy.

While I focus on *and* and *also* in this paper, this is clearly not an exhaustive set. All that matters for us is that there are intersective lexical items like *and* and *also* that let us observe non-exclusivity among cotaxonyms.

Summing up, the cotaxonyms *white* and *green* or *comedy* and *tragedy* cannot have empty intersections, given that they can be co-predicated with intersective expressions. As the reader can verify, the tests given in this section for these two pairs of cotaxonyms also carry over to the other cotaxonyms given in (13). What needs to be explained, then, is why cotaxonyms are intuited as inconsistent in basic sentences:

(50) #This comedy is a tragedy.

There is apparently a linguistic process that obscures cotaxonyms' lexical overlap—descriptively, a case of strengthening. Before discussing how cotaxonyms are strengthened (in Sect. 5), however, I first lay out more explicitly what we have just learned about their lexical meanings.

4.3 Cotaxonyms' lexical meanings: what it means to be "weak"

How is the fact that cotaxonyms have non-empty intersections reflected in the lexical entries of individual predicates? To answer this question, we must again distinguish between summative predicates (which quantify over parts) like *green*, and integrative predicates like *comedy*.

For integrative predicates, our finding of lexical compatibility among cotaxonyms has taught us about concepts. Focusing on *comedy*/*tragedy*, we learn that genre-concepts are not mutually exclusive. If one believes that concepts are laid out in geometric mental spaces (see Sect. 3.2), this means that the conceptual spaces they occupy are not partitioned (contra the theories cited in that section). It is not necessary to commit ourselves to any particular theory of concepts for the purposes of this article; intuitively, the idea is simply that the lexical meaning of *comedy* is that something has *at least* the properties letting it meet the requirements to count as a comedy, rather than having *only* those properties.

Since the part of lexical meaning that we have learned about for integrative predicates is the conceptual structure that underlies these lexical items, the weakness in integrative predicates' lexical meaning does not impact how the formal semanticist writes out their lexical entry:

(51) $[\![\text{comedy}]\!] = \lambda x.\textsf{comedy}(x).$

We can define *comedy* as in (51) regardless of whether *comedy* is taken to be strong or weak lexically, since the strength/weakness of the lexical entry is inherited from the strength/weakness of the concept COMEDY determining whether a particular individual can be categorised as a comedy.

In contrast, for summative predicates like colour adjectives (52), the lexical compatibility we discovered teaches us nothing about colours as concepts. (52a) is contradictory due to the colours modifying *all* parts of the flag, while the co-predications in (52b) are consistent because each cotaxonym only modifies *some* part of the flag (while together covering the entire flag). It is not the case that (52b) means that the entire flag is halfway between whiteness and greenness (i.e., pale green), as would be the case if we had moved the flag to the halfway point between the concepts GREEN and WHITE.

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

Thus, what (52b) teaches us is that colour adjectives are lexically compatible in the sense of being existential (53), not in the sense of overlapping in the conceptual space of colours.

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

In the rest of this paper, I often emphasize colours' existential meaning by writing them out as '(colour)$_\exists$'.

Thus, the lexical meanings of cotaxonyms are sometimes weak in the sense that the conceptual cores of same-domain concepts are mutually compatible, and sometimes weak in the sense that they have existential quantificational force. But there are also cotaxonyms that are apparently lexically strong. As seen in (54), *and* and *also* are powerless to co-predicate truly incompatible cotaxonyms.

(54) a. #Some triangles are **also** squares.
b. #This shape is a triangle **and** a square.

Some predicates, such as species predicates, can only be co-predicated by embedding the *also*/*and*-phrase under *like* or *as if* (55c), indicating that while world knowledge gets in the way, they are in fact underlyingly compatible:

(55) a. #A platypus is a duck that is a beaver.
b. #A platypus is a duck that is **also** a beaver.
c. (i) A platypus is **like** a duck that {is, was} **also** a beaver.
(ii) A platypus is **as if** it was a duck that was **also** a beaver.

That is, while it happens that in the real world, nothing can be both a duck and a beaver, the predicates themselves are in fact lexically–conceptually mutually inclusive. The apparent mutual exclusivity even in (55b) comes from knowledge about species, without being due to the predicates' lexical meaning itself. Ducks and beavers would remain ducks and beavers even if an unexpected shift in our knowledge of biology

meant that we were wrong all along to believe that ducks and beavers cannot have offspring that would be 'ducks that are also beavers'. It may be that the only cotaxonyms that are truly impossible to co-predicate are those whose co-predications would result in mathematical impossibilities (54). The claim about cotaxonyms' strength made in this paper is therefore this: they are all weak, modulo mathematical impossibilities. In the rest of this paper, I focus on cotaxonyms like *comedy/tragedy* and *white/green* that are lexically–conceptually compatible without requiring *like* or *as if*.[8]

Let me end this section by briefly contextualizing my claim that predicates are weak. A similar claim has been made in the field of lexical pragmatics (e.g., Blutner, 1998; Wilson & Carston, 2007; Recanati, 2010; Horn, 2017). Work in that field has taken predicates to be radically underspecified in terms of the underlying concept; if I speak of a 'dog', you will ask yourself what I must have in mind that causes me to use that word; the extension of *dog* is not in fact fixed, the theory goes. Wilson and Carston (2007) write of "ad hoc concepts"—conversation-specific concepts made on the spot for a particular set of purposes. These authors give examples they describe as "lexical narrowing" (56a) and "broadening" (56b):

(56) a. I'm not drinking tonight.
b. That bottle is empty.

(56a) can either mean that the speaker is not drinking anything, or (on the "lexically narrowed" meaning) not drinking alcohol specifically; (56b) can mean the bottle is fully empty, or (on the "lexically broadened" reading) that the bottle is close to being empty.[9]

Essentially, this tradition discusses how we link particular lexical items to concepts. Predicates, on this view, are "weak" in the sense of being underspecified for the particular underlying concept; but they undergo concept-selection upon being used in sentences, which in some sense "strengthens" their meaning. Concept-underspecification is not the same kind of weakness as I am proposing in this article. To describe my claim in terms compatible with the lexical pragmatics literature (without committing myself to this analysis), my claim has to do with the meanings that may be selected: it is *those meanings* that are weak. They may appear strong (mutually incompatible), but they are not so underlyingly. I return to the lexical-pragmatics literature and why it cannot capture the paradigm discussed in this article (it predicts fewer contradictions than are found) in Sect. 5.3.

[8] The existence of apparently strong cotaxonyms like *triangle/square* raises a question. I will soon claim that all cotaxonyms are exhaustified to exclude one another. But this cannot be observed with strong cotaxonyms like *triangle/square*, since they are already strong. The simplest hypothesis is that these are exhaustified too, but trivially so.

[9] I would analyze (56b) in terms of non-maximality rather than ad hoc concepts; see Sect. 7. For (56a), it seems to me that the 'drinking alcohol' meaning of *drink* is highly conventionalized, not ad hoc, and truly ad hoc uses are actually not attested. For example, you cannot show someone grape juice and ask them 'Do you drink?' to mean 'Do you drink grape juice?'

4.4 Section summary

In this section, I showed that cotaxonymic exclusivity cannot be due to lexicalization or the nature of concepts, because it disappears with *and*, *also*, and other expressions like *simultaneously*. These expressions are all intersective in the relevant examples, so cotaxonyms must have non-empty intersections.

5 Obtaining cotaxonymic exclusivity through exhaustification

Given that many sets of cotaxonyms are lexically consistent, why are they interpreted as mutually exclusive? I suggest in this section that this occurs through a domain-general strengthening process ("domain-general" in the sense of not being specific to cotaxonyms), namely semantic exhaustification.[10]

Section 5.1 provides initial motivation for an exhaustification-based analysis by pointing to prior analyses of obligatory additive particles in discourse; what has emerged from the literature on additives is that these are obligatory precisely to circumvent unwanted exhaustification effects. Section 5.2 walks through a first attempt at exhaustifying cotaxonyms, including discussion of the kind of alternatives that predicates trigger. Finally, Sect. 5.3 considers alternative approaches to strengthening cotaxonyms, defending the exhaustification approach against them.

In deriving cotaxonymic exclusivity through exhaustification, this section sets the stage for Sect. 6, where I show that the exhaustification of cotaxonyms has the novel characteristic of always being computed ultra-locally.

5.1 Exh and its interaction with additive particles

As initial motivation for an exhaustification-based account of cotaxonymic exclusivity, let's start by returning to the data with additive particles:

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

Descriptively, this is an "obligatory additive" effect; the sentences are semantically deviant without *also*. While the literature has not discussed this particular obligatory-additive effect, it has discussed another obligatory-additive effect substantially, in particular bisentential discourses where applying the same predicate to two different individuals results in the obligatory use of an additive:[11]

[10] In this paper, I use the term "strengthening" theory-neutrally, while using "exhaustification" to refer specifically to semantic strengthening as postulated by Chierchia et al. (2012).

[11] List intonation (e.g., Steindel Burdin & Tyler, 2018) can override the need for an additive in (58). This can be seen in examples like (ia) (Paillé, 2022b, p. 87) or (ib) (from a reviewer).

(i) a. Q: Who sang?
A: Arwa sang … Sam sang … That's it, I think.
b. Arwa sang. Sam sang. Phil danced. Mary and Sue did somersaults.

(58) Arwa sang. Sam #(also) sang.

Also is presuppositional; all it contributes to its host sentence *Sam sang* is the presupposition that someone other than Sam sang, or more specifically that Arwa sang (Kripke, 1990). Picking up on this, Amsili and Beyssade (2006), Chemla (2008), Sauerland (2008), and Singh (2008, 2011) suggest that *also* is obligatory in (58) due to Heim's (1991) principle of "Maximize Presupposition", which claims that, among a set of alternatives, speakers must utter the sentence that presupposes the most possible. On this view, sentences must have an additive presupposition whenever they can. However, Bade (2016) points out several problems with the use of Maximize Presupposition to explain (58). For instance, it incorrectly predicts that the additive should remain obligatory in (59).

(59) Arwa sang. Sam did not (also) sing.

Assuming (in light of the word order) that the additive scopes below *not* (60), the presupposition of the second sentence in (59) is that someone other than Sam sang. The presupposition projects past negation, so that the entire sentence *Sam did not also sing* presupposes that someone else *did* sing. This is entailed by the context as updated with the sentence *Arwa sang*, so the presupposition should be entirely felicitous.

(60) $[_{vP}$ not $[_{vP}$ also $[_{vP}$ Sam sing]]]

As such, *also* should be obligatory in (59) according to the Maximize Presupposition account of obligatory additives, contrary to fact.

Bade (2016) therefore takes up the second proposed theory of obligatory additives (e.g., Krifka, 1998; Sæbø, 2004; Aravind & Hackl, 2017; Paillé, 2022a), according to which the additive in (58) is obligatory because an unwanted exhaustivity effect would arise without it. Without *also*, the second sentence would be strengthened to mean that *only* Sam sang. I follow Bade in formalizing this through the Exh(aust) operator of Chierchia et al. (2012), although at this point in the argumentation, one could do this through a Gricean maxim (as done by Krifka, 1998, for example). Exh asserts that the constituent it takes (S in (61)), known as its "prejacent", is true, and all of its alternatives (S′) are false, unless they are entailed by the prejacent:

(61) $[\![\text{Exh}_{\text{ALT}}(\text{S})]\!] = 1$ iff $[\![\text{S}]\!] = 1 \wedge \forall \text{S}' \in \text{ALT}[\text{S}'$ is not entailed by S $\rightarrow \ [\![\text{S}']\!] = 0]$

In (58), the second sentence without *also* has the LF and meaning in (62):

(62) $[\![\text{Exh}_{\text{ALT}}\ [\text{Sam}_F$ sang$]]\!] = 1$ iff Sam sang $\wedge$ Arwa didn't sing ($\wedge$ Carrie didn't sing $\wedge$ …)

Sam is focused due to being a contrastive topic with *Arwa*, on this theory; the set of contrastive topics includes at least Sam and Arwa, and in certain discourse contexts could also include others (like Carrie). Thus, without *also*, a contradiction arises in the discourse.

Footnote 11 continued
We will see in this section that *also* is a "de-exhaustifier"; so, then, is list intonation.

This theory raises the question of how the additive can go about fixing the contradiction created by Exh in (62). To keep this paper focused on the strengthening of predicates, I simply adopt my proposal in Paillé (2022a) (closely related to Aravind & Hackl, 2017), according to which *also* interacts with Exh by removing the alternatives that would create a contradiction:

(63) $\text{Exh}_{\{\sout{\text{Arwa sang}}, \text{Sam sang}, \text{Carrie sang}\}}$ [also [Sam_F sang]].

Thus, with *also*, the discourse in (58) is not a contradiction: neither sentence excludes the other. But the discourse is still strengthened to mean that other individuals (here Carrie) did not sing, because Exh is only weakened, not removed entirely. To be sure, the theory is odd in some ways; it must grant *also* the power to prune alternatives in a non-compositional manner. The pruning must be done by *also* itself: otherwise, *also* would never be required because the offending alternatives could always be freely removed. And this pruning can occur even cross-sententially, as shown in discourses like (58), which presumably have the LFs in (64):

(64) $\text{Exh}_{\{\text{Arwa sang}, \sout{\text{Sam sang}}, \text{Carrie sang}\}}$ [Arwa sang]. $\text{Exh}_{\{\sout{\text{Arwa sang}}, \text{Sam sang}, \text{Carrie sang}\}}$ [also [Sam_F sang]].

Still, the theory is better than the alternatives, given Bade's criticisms of the Maximize Presupposition approach (see also the discussion in Paillé, 2022a). I therefore take it as good enough for our purposes, while emphasizing that further work should clarify *also*'s de-exhaustification abilities.

5.2 First steps in the exhaustification of cotaxonyms

Since we have an independently motivated theory stating that *also* can weaken exhaustification to avoid contradictions, the simplest hypothesis is that the interaction between *also* and cotaxonyms is of the same sort. I therefore suggest that cotaxonymic exclusivity is due to cotaxonyms being strengthened through Exh; *also* can weaken this exhaustification through the removal of alternatives. On this view, cotaxonyms form a set of alternatives and trigger exhaustification over this set, as in the following (to be modified in Sect. 6). (65) shows this for integrative cotaxonyms; I use non-co-predicational examples here and return to the contradictory co-predications in Sect. 6.

(65) a. Exh_{ALT} [*Macbeth* is a tragedy].
b. ALT = {*Macbeth* is a tragedy, *Macbeth* is a comedy, *Macbeth* is an epic, …}
c. $[\![(65a)]\!] = 1$ iff $\mathsf{tragedy}(m) \wedge \neg\mathsf{comedy}(m) \wedge \neg\mathsf{epic}(m) \wedge \neg \ldots$

This straightforwardly makes *tragedy* and other genres mutually exclusive: *Macbeth* is asserted to meet the criteria for being a tragedy, and the exhaustification results in the meaning that it does not meet the criteria for other genres.

A central property of Exh is that it only excludes non-entailed alternatives. This means it will not improperly strengthen taxonyms to exclude their superordinates. Consider (66):

(66) a. This is a dog.
b. This is a poodle.

(66a) is compatible with the dog being a poodle, and (66b) entails that it is a dog. These facts are expected on the view that predicates trigger exhaustification. (66b) might well have (66a) as an alternative (I remain agnostic about this); if it does, since Exh does not exclude entailed alternatives, (66a) would not be excluded by an Exh operator on (66b). As for (66a), the most likely possibility is that it simply does not have (66b) as an alternative (see the discussion of colour terms in Sect. 6.2.1). But consider for the sake of argument what would happen if it did. (66b) is stronger than (66a), but an Exh on (66a) would not exclude it due to the notion of Innocent Exclusion proposed by Fox (2007). Informally, Innocent Exclusion is the property of Exh whereby it only excludes alternatives that can all be consistently excluded together. If two alternatives would result in inconsistency if excluded together, Exh does not arbitrarily choose to exclude one or the other, but simply excludes neither. (66a) asserts that the individual is a dog; to be a dog, it must be some kind of dog. Alternatives about kinds of dog ({*This is a poodle*, *This is a labrador*, *This is* …}) cannot all be excluded consistently with the assertion that the individual is a dog. Thus, we do not accidentally predict that *This is a dog* would exclude the possibility of the animal being a poodle.[12]

Let's now turn to the exhaustification of summative cotaxonyms like colour terms. This is shown in (67) (cf. Harnish, 1976 and Levinson, 1983 for a similar analysis of colour terms):

(67) a. Exh_{ALT} [Mars is red].
b. ALT = {Mars is red$_{\exists}$, Mars is white$_{\exists}$, Mars is blue$_{\exists}$, …}
c. $[\![(67a)]\!] = 1$ iff $\text{red}_{\exists}(m) \wedge \neg\text{white}_{\exists}(m) \wedge \neg\text{blue}_{\exists}(m) \wedge \neg \ldots$

If Mars is at least partly red and has no other colour, it must be entirely red.[13]

The truth conditions generated by Exh in (67) do not actually give colour terms a universal quantificational force: *red* on this view is better paraphrased as 'exclusively red' than 'entirely red', and the 'entirely red' meaning only arises to the extent that we are predisposed to assuming that all parts of a surface must have a colour. But

[12] A reviewer asks a similar question about autohyponymy (Horn, 1984; Rohdenburg, 1985; Becker, 2002). Autohyponymy refers to a conceptual hyponym having the same lexicalization as the hyperonym. For instance, while a female lion is a *lioness*, a male lion is still just a *lion*. These facts do not present a challenge to the current account, not least because they involve hyponymy rather than taxonymy—a lioness is not a *kind* of lion (Cruse, 2002). But even if *lioness* and *lion* are alternatives, we do not expect any unwanted outcomes. We can start by assuming that *lioness* lexically entails both *female* and *lion*, and *lion* is unmarked for gender. Now consider these two sentences:

(i) a. This is a lion.
b. This is a lioness.

(ia) may or may not imply that the animal is male depending on whether gender is relevant. If gender is relevant, then (ib) is a stronger alternative that is excluded, resulting in (ia) entailing that the animal is male. If gender is not relevant, then (ib) is not an alternative; it is not excluded, so (ia) remains mum on gender. As for (ib), it will never exclude (ia) because it entails it, so no contradiction is expected.

[13] These truth conditions, which negate the existence of any other colour even to the smallest degree, may appear too strong—colour terms display what has been called "non-maximality", whereby certain discourse contexts can make them weaker than universal. I return to this issue in Sect. 7.

this assumption about surfaces is in fact fallible: some surfaces, in particular fully transparent ones, lack a colour. To see the difficulty for (67c), consider the scenario in (68a) and the fact that (68b) is false in this scenario.

(68) a. SCENARIO: *There is a window. Half the glass is colourless and fully transparent, while the other half is stained red (and still transparent to some extent).*
b. #The window is red.

As it stands, (67c) does not predict that (68b) should be false, since the window has a red part and no part of any other colour. We could claim that *transparent* is part of the alternative set that Exh negates, which would correctly predict (68b) to be false in that scenario. But by doing so, for the scenario (69a), we would then wrongly predict (69b) to be false, because it would negate that the window has a transparent part, when all of it is transparent.

(69) a. SCENARIO: *There is a window, all of which is stained red (and still transparent to some extent).*
b. The window is red.

On the other hand, there is the adjective *clear*, which implies a colourless transparency. If *clear* is an alternative to *red*, we predict both (68b) to be false (the window does have a clear part) and (69b) to be true (the window does not have a clear part).

While it does not in fact follow from 'red$_\exists$ and no other colour' that a surface is fully red, it does follow from 'red$_\exists$, no other colour, and not clear$_\exists$' that a surface is fully red. I therefore assume that *clear* is an alternative to colour terms, and change (67c) to (70) (pretending for ease of exposition that the only colours are red, white, and blue):

(70) $[\![(67a)]\!] = 1$ iff $\mathsf{red}_\exists(m) \wedge \neg\mathsf{white}_\exists(m) \wedge \neg\mathsf{blue}_\exists(m) \wedge \neg\mathsf{clear}_\exists(m)$.

The question this raises is on what basis a particular predicate is or is not included in the set of alternatives for this exhaustification effect. The view so far has been that cotaxonymy settles this matter; is that correct?

In fact, for other reasons entirely, I suggested in Paillé (2023) that what determines whether two predicates are alternatives for the exhaustification effect under discussion is not cotaxonymy, but rather whether two predicates contribute the same kind of information. It just happens that many cotaxonyms do in fact provide the same kind of information—*comedy* and *tragedy* both provide information about the genre of a story; *fork* and *spoon* about the form and function of a utensil; *red* and *green* about colour; etc. Like *red* or *green*, *clear* contributes information about colour (specifically the lack thereof), so if asserting *red* triggers alternatives from other predicates contributing information about colour, we now expect (70).

What, concretely, does it mean for two predicates to "contribute the same kind of information"? At this point in the article, there is nothing preventing us from framing the notion of "the kind of information a predicate provides" in terms of the question under discussion (QUD): two predications provide the same kind of information if they can both constitute answers to the same question. However, we are about to see

in Sect. 6 that the exhaustification effect discussed in this paper is always computed ultra-locally. As such, the QUD, as a phenomenon affecting sentences and not their constituents, is not quite the right notion. I hope to work more on this matter in the future.

In the remainder of this article, I continue writing as if it was cotaxonymy that determined whether two predicates feed alternatives.

5.3 Considering alternative analyses

I have begun sketching out an exhaustification analysis of cotaxonymic exclusivity. Before going on, let us take stock by asking whether there are alternative routes to strengthening cotaxonyms. I briefly discuss three general alternative approaches: (*i*) borrowing a theory from the literature on homogeneity, (*ii*) postulating a strengthening operator specific to cotaxonymic exclusivity rather than domain-general like Exh, and (*iii*) following work in lexical pragmatics.

The first alternative approach might explain summative predicates, but it would have nothing to say about integrative ones. This approach would be to lean on the observation that the part-structure quantification of summative predicates is identical to what is observed with plural predication (e.g., Löbner, 2000; Križ, 2015):

(71) a. (i) The flag is green. (⇝ it is all green)
(ii) The flag is not green. (⇝ it is not green at all)
b. (i) The children sang. (⇝ they all sang)
(ii) The children did not sing. (⇝ none of them sang)

This all-or-nothing effect in positive and negative sentences has been called "homogeneity" (see e.g., Križ & Spector, 2021; Bar-Lev, 2021 for recent discussion and Križ, 2015 for the observation of homogeneity in semantic environments other than summative/plural predication), and raises the question of whether it is right to analyze summative predicates together with integrative predicates (as done here) rather than together with plural predication. However, I argue at length in Paillé (2024) that co-predications of summative predicates make it impossible to apply theories developed for plural homogeneity (71b) to summative predicates (71a). Every theory either predicts inconsistent sentences to be consistent, or vice-versa.[14] If this is correct, there is

[14] There are many different theories of homogeneity, each facing its own particular challenge for co-predications. Let's focus on the minimal pair in (i) to see what these theories predict about it.

(i) a. This is a white and green flag.
b. #This is a white green flag.

Briefly:

1. The theory that homogeneity is due to an all-or-nothing presupposition (e.g., Löbner, 2000; Gajewski, 2005) ends up giving (ia) the meaning that all parts of the flag have a white part and also have a green part. This gives rise to an impossible flag where however small a piece one chooses, it is divisible between a partly white part and a partly green part, and so on to infinity. See Paillé (2024, §5.1).
2. The theory that homogeneity arises due to pragmatic underspecification of the quantificational force of the plural operator and of summative predicates (Krifka, 1996; Lasersohn, 1999; Winter, 2001; Malamud, 2012) predicts all co-predications to be non-contradictory, because the pragmatics would not create contradictions out of potentially non-contradictory material. See Paillé (2024, §6.1).

no theory of plural homogeneity (as applied to summative predicates) that competes with the exhaustification approach given in the present paper.

The second alternative route to deriving cotaxonymic exclusivity would not aim to strengthen cotaxonyms via a domain-general Exh operator, but by stipulating an ad hoc cotaxonym-strengthening operator (CSO). In the next section, I claim that the exhaustification leading to cotaxonymic exclusivity is always ultra-local to the cotaxonym. Rather than claiming that cotaxonyms are strengthened by an Exh operator with a special property/constraint, one could simply claim that cotaxonyms are strengthened by a local CSO. The motivation I gave in the present section to use Exh was an appeal to previous analyses of empirically different obligatory-additive effects. But there may be a stronger argument for using a domain-general Exh rather than a CSO. Indeed, in Paillé (2022b, Ch. 4, §2), I discuss similar cases of ultra-local exhaustification with phrasal adjuncts in sentences, giving reason to believe that a hypothetical CSO would not be a general enough explanation. This gives more weight to my use of Exh.

The third alternative we should consider is the lexical-pragmatic route. As described in Sect. 4.3, the literature on lexical pragmatics has discussed cases of "strengthening" in the following sense: language users have a variety of (possibly ad hoc) concepts they may associate with a given predicate, and by choosing one of these concepts, they narrow down the meaning of that predicate. This can be described as "strengthening", even though it is not strengthening via the negation of alternatives. Can such a mechanism of concept-selection be used to explain the basic minimal pairs focused on in this paper, like (72)?

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

Short of attributing conflicting opinions to a speaker, (72) is not attributable to reasoning about which ad hoc concept a speaker has in mind; presumably, the speaker would not co-predicate *comedy* and *tragedy* if they had in mind mutually exclusive ad hoc concepts for *comedy* and *tragedy*, so the sentence should be consistent even without *also*.

Footnote 14 continued

3. The theory by Križ and Spector (2021) based on the co-assertion of so-called "candidate interpretations" faces a similar problem; either the mechanism overseeing the co-assertion of these candidate interpretations is sensitive to the creation of contradiction, in which case it would never generate the contradiction in (ib), or it is not, in which case it would generate a contradiction in (ia). See Paillé (2024, §6.2).
4. The theory by Bar-Lev (2018, 2021) is based on weak lexical meaning (for the plural operator, or for summative predicates if we extend his discussion to these) being exhaustified to not *exclude* but *include* (i.e., assert the truth of) domain alternatives. To obtain the consistency of (ia), one needs to posit alternatives created via both domain-restriction and via the replacement of the conjunction ('*P* and *Q*') by disjunction ('*P* or *Q*') or single conjuncts ('*P*', '*Q*'). See Paillé (2024, §5.2) for more elaborate exposition. While this can obtain (ia) (as well as (ib), which differs from (ia) in lacking a conjunction and, therefore, lacking alternatives obtained by replacing *and* with *or*), I show there that it ends up incorrectly predicting that (ii) should be consistent (cf. Sect. 4.1 of the present article).

(ii) #The flag is completely white and completely green.

Bar-Lev's approach also cannot explain the consistency of co-predications with *also*, because *also* (unlike *and*) does not have a disjunctive alternative.

5.4 Interim conclusion

In this section, I took some first steps in modelling cotaxonymic exclusivity as the result of semantic exhaustification. We now turn to seeing that if the Exh account of cotaxonymic exclusivity is adopted, it comes with an important consequence for theories of strengthening in meaning.

6 Observing ω-exhaustification with cotaxonyms

The exhaustification of cotaxonyms has syntactic properties unlike what has been standardly described. In this section, I start by taking Chierchia et al. (2012) as a starting-point for non-cotaxonymic exhaustification phenomena (Sect. 6.1). I take these authors' description to be correct for the exhaustification of non-cotaxonymic alternative-triggering expressions (focusing on *most*), and then show that it does not carry over to cotaxonyms (Sect. 6.2). Finally, I end the section (Sect. 6.3) by returning briefly to the "de-exhaustifiers" (*and* and *also*) that let us observe cotaxonyms' weak lexical meanings.

6.1 Standard "free" exhaustification

Chierchia et al. (2012) describe Exh as being able to scope at various loci in a sentence without being constrained by the alternative-triggering expression that feeds its alternatives. Consider the exhaustification of *most* whereby it comes to exclude 'all':

(73) I read most of the books.
$\rightsquigarrow$ I did not read all of the books.

If we embed a clause containing *most*, its exhaustification can be computed either within the embedded clause or globally. This can be seen in (74):

(74) Arwa was awake while Sam hit most of the targets.

There are two prominent meanings for (74). On one meaning, (74) negates that Sam hit all the targets: for each target, Arwa was awake while he attempted to hit it, but there were some targets that he did not hit. This meaning can be brought out by the elaboration in (75a), and it arises from an Exh operator in the embedded *while*-clause (75b).

(75) a. …He never hits all of them!
b. ⟦[Arwa was awake while [Exh_{ALT} [Sam hit most of the targets]]⟧
= 1 iff Arwa was awake while Sam hit most but not all of the targets.

But (74) can also mean that Arwa was asleep while Sam hit some of the targets. On this meaning, Sam might well have hit all the targets, but Arwa was only awake during most but not all of the hits. This meaning can be brought out with the continuation in (76a), and is computed with a global Exh (76b).

(76) a. …She doesn't know whether he hit all of them, since she was asleep during some of his attempts.
b. $[\![\text{Exh}_{\text{ALT}}$ [Arwa was awake while Sam hit most of the targets]$]\!]$
= 1 iff Arwa was awake while Sam hit most of the targets $\wedge$ $\neg$[Arwa was awake while Sam hit all of the targets].

Given that Exh can occur either closer to or farther from *most*, one concludes that Exh does not have a grammatical link to such alternative-triggering expressions. It is syntactically free, with its precise placement decided by language users based on the meanings that arise. Call this FREE EXHAUSTIFICATION (rebranded shortly below as φ-exhaustification).[15] Is the syntactic freedom found for the Exh exhaustifying *most* also observed with the Exh exhaustifying cotaxonyms? One would certainly expect so. However, what we will see in this section is that in fact, Exh is necessarily in syntactic proximity to the cotaxonyms it exhaustifies. Cotaxonyms both require the presence of Exh and dictate its syntactic position. The result is that Exh's contribution resembles lexical meaning, being both obligatory and scopally inseparable from the cotaxonym.

This claim can be viewed as an unfortunate stipulation about the exhaustification of cotaxonyms, but it can also be viewed an interesting discovery. On the standard view, exhaustification is syntactically free and is a phenomenon associated with clauses/propositions. We can call this φ-exhaustification, where the letter φ is chosen for two associations it creates: a phonological association with "free" and an association with the fact that semanticists often use φ as a variable for propositions.

(77) **φ-EXHAUSTIFICATION**:
CLAUSES are strengthened without any restrictions on how close Exh is to the alternative-triggering expression(s).

My claim is that there is another kind of strengthening effect, whereby *words* are strengthened immediately upon entering the syntax, rather than having the clauses they are part of be exhaustified. Borrowing phonologists' use of ω for word-level phonological constituents, call this ω-exhaustification.[16]

(78) **ω-EXHAUSTIFICATION**:
WORDS are strengthened by themselves (i.e., ultra-locally).

There is nothing inherently more stipulative or unappealing about the claim that lexical items are strengthened immediately upon entering the syntax (i.e., that the exhaustification of cotaxonyms applies to individual predicates) than the standard view that

[15] In addition to Exh being syntactically free, Chierchia et al. (2012) also describe Exh as being optional, although this assumption has been challenged several times (e.g., Magri, 2009; Bade, 2016; Bar-Lev, 2018). With cotaxonyms, it is certainly not the case that Exh is optional; otherwise (i) would have a non-contradictory parse where the cotaxonyms are optionally non-exhaustified:

(i) #This comedy is a tragedy.

[16] I leave open for this article whether ω-exhaustification is found with any expressions other than the simplex predicates being discussed. As noted in Sect. 5.4, I argue in Paillé (2022b, Ch. 4, §2) that it is also found with some sentential adjuncts. If so, the phrase 'words' in (78) would have to be modified.

exhaustification applies to clauses, and no reason why both phenomena could not be present in natural language.[17] Work on strengthening has focused entirely on φ-exhaustification (e.g., the anti-universal meaning of *most*), but presumably this is simply because such effects are easier to find. After all, unlike φ-exhaustification, the result of ω-exhaustification is virtually indistinguishable from lexical meaning, making it difficult to identify.

The observation of locality in strengthening requires a semantic approach to the strengthening of predicates that creates cotaxonymic exclusivity. Indeed, a pragmatic Gricean approach does not predict embedded strengthening (e.g., Geurts, 2010; Sauerland, 2012), and the present section is dedicated to showing that cotaxonyms not only *can* but *must* be strengthened locally. Since the strengthening of cotaxonyms is (*i*) ultra-local and (*ii*) not due to lexical pragmatics (see Sect. 5.3), there is no option left but to view it as part of the semantic composition—using a local Exh, on my proposal.

6.2 Observing the locality constraint on Exh with cotaxonyms

Empirically, observing that Exh is necessarily local with cotaxonyms will always involve observing that cotaxonyms are interpreted as incompatible regardless of the syntactic structure or semantic environment they are in (unless there is a de-exhaustifier). Under the Exh account of cotaxonymic exclusivity, we will see that this can only be captured by strictly restricting Exh's syntax. In this section, I turn to some different ways to see this. Before doing so, however, I comment both on the semantic type of Exh's prejacent and on the kind of cotaxonyms that provide the clearest data.

I will suggest that Exh's locality requirement with cotaxonyms is such that it takes only the cotaxonym as its argument:

(79) The [Exh_{ALT} green] flag is high.

Exh takes a predicate-type argument in (79); yet, it was defined in (80), repeated from (61), as a propositional operator.

(80) $[\![\text{Exh}_{\text{ALT}}(\text{S})]\!] = 1$ iff $[\![\text{S}]\!] = 1 \wedge \forall \text{S}' \in \text{ALT}[\text{S}'$ is not entailed by $\text{S} \rightarrow [\![\text{S}']\!] = 0]$.

I therefore define a predicational Exh (81), called ω-Exh, which takes a predicate and makes it exclude other alternative predicates, based on a generalised notion of entailment (taken for granted so far in this article). A predicate P entails another predicate Q if for all x, $P(x)$ entails $Q(x)$. For instance, *dog* entails *animal*, and *scarlet* entails *red*.

(81) $[\![\omega\text{-Exh}_{\text{ALT}}(P)]\!] = \lambda x.[\![P]\!](x) \wedge \forall P' \in \text{ALT}[P'$ is not entailed by $P \rightarrow [\![P']\!](x) = 0]$.

Thus:

(82) $[\![\omega\text{-Exh}_{\text{ALT}}\ \text{green}]\!] = \lambda x.\mathsf{green}_{\exists}(x) \wedge \neg\mathsf{white}_{\exists}(x) \wedge \neg\mathsf{red}_{\exists}(x) \wedge \ldots$

[17] If both types of exhaustification exist (as I claim), some mechanism has to decide which type of exhaustification applies to which sets of alternative-triggering expressions. See Sect. 6.4.

For clarity, in addition to writing "ω-Exh" for the predicational Exh, I will also start writing "φ-Exh" for the propositional and syntactically unrestrained Exh assumed so far. When the distinction does not matter, I will go on writing "Exh".

The second point worth clarifying is which kind of cotaxonyms provide the clearest judgments for the kinds of sentences we will test in this section. Recall that we will be observing evidence that, even in precisely those environments where the exhaustification account of cotaxonymic exclusivity leads one to expect cotaxonymic exclusivity not to be observed, it is still observed. Many of these examples will not be co-predicational. But recall from Sect. 3 that cotaxonymic exclusivity with integrative predicates is more difficult to observe with certainty in non-co-predicational environments like (83), because one could always claim that the apparent mutual exclusivity between cotaxonyms is just a prototypicality effect.

(83) This is a comedy.

I argued in Sect. 3 that this would be the wrong analysis: prototypicality effects are easily defeated while cotaxonymic exclusivity is not. (84) is repeated from (26).

(84) a. The dog is blue. ⇒ *prototypicality effects are easily defeated*
b. #The comedy is a tragedy. ⇒ *contradictions from cotaxonymic exclusivity are not*

Still, since we will be toying with different syntactic/semantic environments where the exhaustification theory leads one to *expect* cotaxonymic exclusivity to disappear (at least on the standard view that there is only φ-Exh), but where I will claim it does not, a skeptical reader could blame prototypicality effects rather than exhaustification with integrative predicates. To get around this, I will rely primarily on examples involving colour terms—while also giving examples with integrative predicates for completeness. As we saw in Sect. 3, summative predicates can easily be appreciated to be mutually exclusive even in non-co-predicational environments: they are interpreted as universal.

To simplify the alternatives for examples with colour terms, I will pretend that the only colours are green, white, and red. We now turn to three types of environments in which to observe a locality condition on Exh with cotaxonyms.

6.2.1 Predicating two cotaxonyms of the same referent

I start with precisely the kinds of examples that motivated cotaxonymic exclusivity in the first place, namely co-predications. The fact that language does not allow co-predication of cotaxonyms even in a single sentence (85) is in fact unanticipated for the Exh account of cotaxonymic exclusivity.

(85) 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 (85) would in fact be non-contradictory. It would be possible for Exh to scope globally, as shown in (86) for (85a). 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, neither colour term would be excluded. Only colour terms other than green or white would be excludable. Thus, what would result is the non-contradictory meaning in (86a), rather than something like (86b), as needed to have a contradiction. For ease of exposition, (86) only shows the predicative adjective *white* as triggering alternatives, but I argue below that that *green* does too.

(86) $[\![\varphi\text{-Exh}_{\text{ALT}}$ [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
$\begin{cases} \text{the green}_\exists \text{ flag is white}_\exists \wedge \\ \neg[\text{the green}_\exists \text{ flag is green}_\exists] \wedge \\ \neg[\text{the green}_\exists \text{ flag is red}_\exists] \end{cases}$
$\Rightarrow$ contradiction

Example (85a) involves a definite subject, which triggers a uniqueness presupposition, but this is not what is at cause in creating a contradiction (at least, such a claim would be insufficiently general), since co-predicational contradictions are observed with indefinites as well:

(87) #Some green flags are white.

Besides, prior to exhaustification, the uniqueness presupposition in (85a) is merely that there is a unique 'green$_\exists$' flag, which does not result in incompatibility with *white* by itself. To blame the presupposition in (85a), one would therefore need to claim that presuppositions are independently strengthened too (as done by Geurts, 2010, for example).

Things are in fact even worse than appear in (86) if the attributive *green* triggers alternatives too. With a colour term inside a definite DP, a global Exh would create entailments about other flags altogether:

(88) $[\![\varphi\text{-Exh}_{\text{ALT}}$ [the green flag is high]$]\!]$ = 1 iff
$\begin{cases} \text{the green}_\exists \text{ flag is high} \wedge \\ \neg[\text{the white}_\exists \text{ flag is high}] \wedge \\ \neg[\text{the red}_\exists \text{ flag is high}] \end{cases}$

These are not actually intuited unless *green* is contrastively focused (see Sect. 8).

In order for the cotaxonyms in co-predications 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 colour term in its scope. One possible LF is (89); we return shortly below to whether there might be an Exh on only one colour term.

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

≈ 'The exclusively green flag is exclusively white.'
⇒ contradiction

In fact, the need for an ultra-local Exh is only motivated on the assumption that Exh can see entailment relations among colour adjectives to know what it can and cannot exclude. Reviewers from various venues have suggested the following analysis to obtain a contradiction in (86) without forcing Exh to be ultra-local. Expanding on Magri's (2009) argument that Exh does not take world knowledge into account, they suggest that Exh is so blind to the content of predicates that it does not even know that *The green*$_{\exists}$ *flag is white*$_{\exists}$ entails *The green*$_{\exists}$ *flag is green*$_{\exists}$—put another way, that one token of *green* entails another token of *green*. A global φ-Exh on *The green*$_{\exists}$ *flag is white*$_{\exists}$ would therefore negate *The green*$_{\exists}$ *flag is green*$_{\exists}$, obtaining the contradiction. This view requires Exh to be rather reckless, in being willing to negate alternatives on the mere grounds that they are not known to be entailed (while also not being known to be non-entailed).

I have two arguments against this alternative approach to deriving contradictions. They are not arguments against Magri's claim that Exh does not take world knowledge into account; world knowledge and lexical–conceptual meaning are different things. Exh can take the lexical–conceptual meaning of predicates into account (it certainly does, as I will show) without taking general world knowledge into account.

My first argument against this account of contradictory co-predications comes from colour terms specifically, and entailment relations among them. There is clear evidence that Exh does in fact know which colour predicates it can exclude. Consider the following minimal pair:

(90) a. #The red flag is green.
b. The red flag is scarlet.

In (90b), the exhaustification of *scarlet* does not exclude *red* because *scarlet* entails *red*. Crucially, while *red* is a basic-level colour term and *scarlet* is a subordinate-level term, we know that basic-level colour terms are alternatives for subordinate-level colour terms from data like (91):

(91) The flag is scarlet.
≈ 'The flag is entirely scarlet.'

If *scarlet* only had subordinate-level colour terms as alternatives, (91) would not be universal on the theory presented in this article. Only excluding subordinate-level colour adjectives would not actually exclude the entire colour wheel. While English and other languages cover the entire colour wheel with their basic colour terminology (Berlin and Kay, 1969), their subordinate colour terms are more of a patchwork. For instance, English does not have many subordinate colour terms for types of pink or orange (and most of them are structurally complex, like *hot pink*), and they certainly do not cover all hues of pink/orange. (91) would therefore be compatible with the flag having pink and orange parts if *scarlet* only excluded subordinate-level colour terms. Thus, it must be that, in (90b), *scarlet* has *red* as an alternative, but Exh does not

exclude it.[18] The important conclusion is that Exh sees entailment relations among colour predicates.

My second argument against the view that φ-Exh creates contradictions by negating alternatives like *The green*$_{\exists}$ *flag is green*$_{\exists}$ comes from more general problems this view would create. If Exh really could not see that there is an entailment relation between two tokens of the same predicate (*green* and *green*), this would be the case generally, not just in inconsistent co-predications. Now consider (92) and the exhaustification of *most*:

(92) Most of the students read books.

Putting aside the exhaustification leading to cotaxonymic exclusivity, (92) has at least the following alternatives:

(93) $\text{ALT} = \left\{ \begin{array}{c} \text{Some of the students read books,} \\ \text{Most of the students read books,} \\ \text{All of the students read books} \end{array} \right\}$

But if Exh cannot see that *student* always entails *student*, *read* always entails *read*, and *books* always entails *books*, then it will believe that the three alternatives in (93) are logically independent. In particular, Exh would negate *some of the students read books*, since it is not entailed (on this view) by *most of the students read books*. (92) would therefore be intuited as a contradiction: there are no students who read books, but most of them do. In sum, to create a contradiction in (94a) by adopting the claim that φ-Exh (due to being entirely blind to the content of predicates) does not understand that (94a) entails (94b), we end up creating contradictions in all exhaustification effects for all sentences containing any predicates.

(94) a. #The green flag is white.
b. The green flag is green.

Going back to my view that Exh creates contradictions due to a locality constraint, let's consider for a moment the fact that (95) is not the only logical possibility for creating the contradiction in (94a) via locality.

(95) The [ω-Exh$_{\text{ALT}}$ green] flag is [ω-Exh$_{\text{ALT}}$ white].

Indeed, a contradiction is predicted as long as one or the other cotaxonym is strengthened irrespective of the other. 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.

It is not clear on what basis one cotaxonym but not the other would be exhaustified. But more importantly, there is empirical evidence that both cotaxonyms must be exhaustified ultra-locally as in (95). In the set of co-predications in (96), one or both colour terms is explicitly weakened through the adverb *partly*, so that if a contradiction is observed, it must be due to the *other* colour term. The judgments for

[18] As for why *red* does not exclude *scarlet*, the simplest assumption is that *scarlet* is not an alternative to *red* because *red* is a basic-level colour term while *scarlet* is a subordinate-level colour term. Basic colour terms are alternatives to subordinate ones, but not vice-versa.

the examples can be tricky because it is easy to think of the colour term that is not modified with *partly* as referring to the background of the flag, and the one modified by *partly* as referring to some small part superimposed over this background. In this kind of situation, we do not expect the creation of a contradiction, because the colour term referring to the background can consistently be universal/exclusive: a green flag with a white circle on it can be described as 'being green' due to having an entirely green background, but this does not mean that *green* is not universal/exclusive there. For this reason, the judgments in (96) are given for a half-green, half-white flag.

(96) SCENARIO: *The flag is half white and half green.*
   a. The partly green flag is #(also) white.
   b. The green flag is ??(also) partly white.
   c. The partly green flag is ?(also) partly white.

The fact that (96a) requires *also* shows that the predicative adjective *white* is exhaustified independently of the attributive *partly green*: if *white* had no Exh or had an Exh which scoped above *partly green*, no contradiction would arise. Likewise, the fact that (96b) requires *also* shows that the attributive adjective *green* is exhaustified independently of *partly white*, for the same reason. Finally, (96c) 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 (96a–b). Thus, both attributive and predicative cotaxonymic adjectives can be shown to be exhaustified independently of the other. If this is the case in (96a–b), it is presumably also the case in (94a). This is evidence that the LF in (95), with a local ω-Exh operator on each colour term, is the only LF associated with the sentence #*The green flag is white* (94a).

6.2.2 Cotaxonyms with a scalar expression

In addition to contradictory co-predications, we can also observe the locality requirement on Exh by considering cases where a cotaxonym co-exists with another weak alternative-triggering expression, such as an existential quantifier:

(97) 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.'

Weak expressions like *some* are usually exhaustified. As we will see, if Exh could scope above both the cotaxonym and the other alternative-triggering expression *some*, we would obtain wrong results.

When a single Exh operator has a prejacent with more than one alternative-triggering expression, I assume that the set of alternatives it takes includes all the sentences obtained by replacing one or more alternative-triggering expression with one of its alternatives (Sauerland, 2004). Thus, if (97a) is exhaustified with a single global φ-Exh, the alternatives are in (98).

(98) $\text{ALT} = \left\{ \begin{array}{ll} \text{Some flags are green,} & \text{All flags are green,} \\ \text{Some flags are white,} & \text{All flags are white,} \\ \text{Some flags are red,} & \text{All flags are red} \end{array} \right\}$

The question is which of these alternatives Exh actually excludes. It cannot exclude all the ones not entailed by the prejacent. If it did, we would obtain the truth conditions in (99).

(99) $[\![\varphi\text{-Exh}_{\text{ALT}}\ [\text{some flags are green}]]\!] = 1 \text{ iff } \left\{ \begin{array}{l} \text{some flags are green}_{\exists} \wedge \\ \neg[\text{some flags are white}_{\exists}] \wedge \\ \neg[\text{some flags are red}_{\exists}] \wedge \\ \neg[\text{all flags are green}_{\exists}] \wedge \\ \neg[\text{all flags are white}_{\exists}] \wedge \\ \neg[\text{all flags are red}_{\exists}] \end{array} \right.$

This meaning goes against the world knowledge that surfaces must have a colour (or be clear; recall from Sect. 5 that I am assuming that *clear* is an alternative to colour terms): (99) 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 (99), there must be at least one colourless flag.

We can try to salvage (99) by appealing to Innocent Exclusion, at least if Innocent Exclusion takes world knowledge about colours into account. In (99), the alternatives whose exclusion lead to (non-logical) inconsistency are the following:

(100) {some flags are white$_{\exists}$, some flags are red$_{\exists}$, all flags are green$_{\exists}$}

We could try to claim that Innocent Exclusion ensures that these alternatives are simply not excluded. This leaves us with the truth conditions in (101) instead of (99):

(101) $[\![\varphi\text{-Exh}_{\text{ALT}}\ [\text{some flags are green}]]\!] = 1 \text{ iff } \left\{ \begin{array}{l} \text{some flags are green}_{\exists} \wedge \\ \neg[\text{all flags are white}_{\exists}] \wedge \\ \neg[\text{all flags are red}_{\exists}] \end{array} \right.$

While this no longer entails that some flags lack a colour entirely, this is not a good result either. First, (101) only means that some flags are *partly* green, not *exclusively* green. What is more, it does not strengthen *some* to mean 'not all'. The result is that some (maybe all) flags are partly (maybe entirely) green, rather than some but not all flags being exclusively green. Our attempt at modifying the unwanted result of (99) via Innocent Exclusion has failed.

In contrast, we can obtain the right truth conditions by having two Exh operators in the sentence, as shown in (102a). The first is an ω-Exh that is immediately above the cotaxonym and has cotaxonymic alternatives. The other is a global φ-Exh that takes alternatives created by replacing *some* with its scalemates (just shown as *all* in (102)), but not replacing the cotaxonym with anything, as if the cotaxonym has been rendered inactive by having already been used by the lower ω-Exh. The alternatives for the two Exh operators in (102a) are shown in (102b). This creates the truth conditions in (102c).

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

b. (i) ALT-1 = {green, white, red}
(ii) ALT-2 = {some flags are [ω-Exh$_{\text{ALT-2}}$ green], all flags are [ω-Exh$_{\text{ALT-1}}$ green]}

c. $[\![(102a)]\!] = 1$ 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: some but not all flags are exclusively green.

To reiterate, the Exh leading to cotaxonymic exclusivity has to be local to the cotaxonym. The strengthening of *some* is an instance of standard, free φ-exhaustification; it is not subject to the notion of ω-exhaustivity being developed for cotaxonyms in this section.

6.2.3 Cotaxonyms in downward-entailing (DE) environments

The last case in which we observe a locality requirement on Exh with cotaxonyms is when the cotaxonym is under a DE operator. Exh under such operators is normally dispreferred, because it leads to global weakening rather than strengthening. Under *if*, for example, *most* is not strengthened to being anti-universal (at least without focus intonation on *most*):

(103) If you read most of the books, you'll be really knowledgeable.

Expressions in DE environments *can* be strengthened (leading to global weakening), as in (104) (inspired by a similar example with *or* from Chierchia et al. (2012, p. 2306)).

(104) If you read most of the books, you'll be really knowledgeable; but if you read all of them you'll lose your mind.

The meaning in (104) comes about from a φ-Exh operator below *if*:

(105) If [φ-Exh$_{\text{ALT}}$ [you read most of the books]], you'll be really knowledgeable; but …

Cotaxonyms, however, do not let their weak lexical meanings surface in DE environments. They remain mutually exclusive:[19]

(106) a. If the flag is white, you can do a jumping jack.
≈ 'If the flag is entirely white, you can do a jumping jack.'

[19] As a note on the judgment, colours terms have been described as sometimes being weaker than universal due to discourse factors; see Sect. 7. It is not impossible to interpret *white* weakly in (106a), but crucially, unlike other exhaustification effects, putting *white* under *if* does not *preferably* lead to a weak meaning. I therefore maintain the claim in the main text: *white* is exhaustified below *if*, and to the extent that we can intuit (106a) as meaning that the addressee may do a jumping jack as long as the flag has any white, this is due to post-exhaustification discourse-based weakening, as discussed in Sect. 7.

b. If the play is a comedy, I am willing to pay a lot of money.
≈ 'If the play is a non-tragic comedy, I am willing to pay a lot of money.'

Of course, for (106a), simply not exhaustifying *white* will not create this meaning. But even if an Exh is present, it cannot create the strong meaning of *white* in (106a) if it is non-local:

(107) $[\![\varphi\text{-Exh}_{\text{ALT}}$ [If the flag is white, you can do a jumping jack]$]\!]$
= 1 iff you can do a jumping jack if the flag is white$_{\exists}$ $\wedge$
$\neg$[you can do a jumping jack if the flag is red$_{\exists}$] $\wedge$
$\neg$[you can do a jumping jack if the flag is blue$_{\exists}$].

On the other hand, an Exh below *if* does obtain the right truth conditions. (108) is not the only logical possibility, as long as Exh is below *if*.

(108) $[\![$If the flag is [$\omega\text{-Exh}_{\text{ALT}}$ white], you can do a jumping jack$]\!]$
= 1 iff you can do a jumping jack if the flag is [white$_{\exists}$ & not blue$_{\exists}$ & not red$_{\exists}$].
≈ 'You can do a jumping jack if the flag is exclusively white.'

The empirical picture for DE environments is more complex than this, however. There are a handful of exceptional DE environments where cotaxonyms are interpreted as non-exclusive. These are all contexts that have a negative flavour (cf. Chierchia, 2004)—including most prominently sentential negation, but also possibly *doubt* and *no* (see Paillé, 2022b, Chs. 2 and 4 for brief discussion). Consider *shirt* or *white* under *not*:[20]

(109) a. This is not a shirt.
b. This is not white.

Without focus intonation on the cotaxonyms, the sentences in (109) negate the weak meanings of the predicates *shirt* and *white*, rather than their meanings displaying cotaxonymic exclusivity. (109a) denies that the garment has the necessary properties to count as a shirt, rather than denying that it is a 'shirt and nothing else'; if the garment was a 'shirt that was also a hat' or the like, (109a) would only be felicitous with focus intonation on *shirt*:

(110) This is not a SHIRT$_F$. It's a [SHIRT THAT'S ALSO A HAT]$_F$.

Likewise, (109b) denies that the subject has any white (e.g., Löbner, 2000)—the existential lexical meaning of *white*. The sentence would not be true if the subject was both white and some other colour.

To deal with this, there are two possibilities. The first option is to weaken the claim that Exh is obligatory/necessarily local with cotaxonyms; perhaps there are a small handful of exceptional environments where Exh is not required. This is not a trivial explanandum: work on exhaustification usually treats all DE environments as a single

[20] For integrative cotaxonyms (those not referring to part structure), I have usually been using *comedy*/*tragedy*. I change *comedy* to *shirt* in (109) because I find the existence of the lexical item *tragicomedy* gets in the way for the judgment for *comedy* under negation.

class, but the exceptional environments we are concerned with are not just DE, but specifically negative-flavoured DE environments. The second option, which I adopt in this article, is to modify Exh's meaning slightly so that it can be claimed to be local and obligatory in all environments, including under negation. Bassi et al. (2021) suggest that Exh excludes alternatives in its truth conditions, but not its falsity conditions (it is trivalent):

$$(111)\qquad [\![\varphi\text{-Exh}^{triv\cdot}_{\text{ALT}}(\text{S})]\!] = \begin{cases} 1, \text{ iff } [\![\text{S}]\!] = 1 \wedge \forall \text{S}' \in \text{ALT}[\text{S}' \text{ is not entailed by S} \rightarrow [\![\text{S}']\!] = 0]; \\ 0, \text{ iff } [\![\text{S}]\!] = 0; \\ \#, \text{ otherwise} \end{cases}$$

Naturally, sentential negation simply switches the truth and falsity conditions of the negated constituent:

$$(112)\qquad [\![\text{not(S)}]\!] = \begin{cases} 1, \text{ iff } [\![\text{S}]\!] = 0; \\ 0, \text{ iff } [\![\text{S}]\!] = 1; \\ \#, \text{ otherwise} \end{cases}$$

Since the trivalent Exh in (111) does not affect falsity conditions, one expects its contribution to the truth conditions of negative sentences to be vacuous. This is shown in (113) for (109a), where I hypothesise the ultra-local presence of a trivalent ω-Exh with a cotaxonym under *not*.

$$(113)\qquad [\![\text{not}\,[\text{this}\,(a)\,\text{is}\,[\omega\text{-Exh}^{triv\cdot}_{\text{ALT}}\,\text{white}]]]\!] = \begin{cases} 1, \text{ iff } \neg\mathsf{white}_{\exists}(a); \\ 0, \text{ iff } \mathsf{white}_{\exists}(a) \wedge \neg\mathsf{blue}_{\exists}(a) \wedge \neg\mathsf{red}_{\exists}(a) \wedge \neg \ldots; \\ \#, \text{ otherwise} \end{cases}$$

On this view, the sentences in (109) are not actually counterexamples to the claim that Exh is necessarily local with cotaxonyms. With the trivalent Exh, we expect negative-flavoured DE environments to behave differently from non-negative DE environments, as observed with cotaxonyms.

6.2.4 Interim summary: Exh immediately above cotaxonyms

We have gone through three classes of data where cotaxonymic exclusivity is observed despite non-local Exh operators failing to create it. On the Exh account of cotaxonymic exclusivity, then, it must be that Exh is subject to a locality constraint; the strengthening of cotaxonyms does not behave as a run-of-the-mill φ-exhaustivity effect applying freely to clauses at any syntactic distance from the alternative-triggering expression(s). The following examples are repeated from the previous subsections; they show how cotaxonymic exclusivity can be generated through local ω-Exh operators.

(114) [The [ω-Exh$_{\text{ALT}}$ green] flag is [ω-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 exclusively green flag is exclusively white.'
⇒ contradiction

(115) $[\![\varphi\text{-Exh}_{\text{ALT}}\ [\text{some flags are}\ [\omega\text{-Exh}_{\text{ALT}}\ \text{green}]]]\!]$

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

$\approx$ 'Some but not all flags are exclusively green.'

(116) $[\![$If the flag is $[\omega\text{-Exh}_{\text{ALT}}$ white], you can do a jumping jack$]\!]$
= 1 iff you can do a jumping jack if the flag is [white$_{\exists}$ & not green$_{\exists}$ & not red$_{\exists}$].
$\approx$ 'If the flag is exclusively white, you can do a jumping jack.'

I emphasize that this locality requirement is even more stringent than Exh just needing to be in the same clause as the cotaxonym: (114)–(115) are monoclausal, but a global Exh must still be ruled out. In fact, Exh with cotaxonyms is so local that it effectively mimics lexical meaning, due to being obligatory and local to the predicate; it is only with *and*, *also*, and the like that we can appreciate that the meaning it contributes is not lexical.

There are ways to claim that Exh is not so ultra-local in some of these examples; for instance, there is nothing in (116) preventing Exh from being farther away from *white*, as long as it is under *if*. But these would be ad hoc attempts at moving Exh farther away from the cotaxonym in a sentence-by-sentence manner; in truth, we can never actually observe any syntactic distance between Exh and the cotaxonym (but see my discussion of *and* immediately below), so the simplest hypothesis is that cotaxonyms are exhaustified by themselves as soon as they merge into the syntax via ω-exhaustification.

6.3 De-exhaustifiers and local Exh

Having established locality in the exhaustification of cotaxonyms, we now turn briefly to how *also* and *and* can let us observe cotaxonyms' weak meanings. We briefly discussed *also*, but not *and*, in Sect. 5.1. The goal for this article is not to develop a proper theory of these "de-exhaustifers", so I merely to suggest a way forward—and most importantly, show that there is no need to view de-exhaustifiers as going against my claim that there is always a local ω-Exh with cotaxonyms.

Let me first repeat from Sect. 5.1 that de-exhaustifiers do not really de-exhaustify: while *comedy* and *tragedy* (for example) become mutually compatible in the presence of *and* and *also* (117), they are still exclusive of other genres—the play referred to in (117) cannot be an epic, for example.

(117) a. This play is both a comedy and a tragedy.
b. This comedy is also a tragedy.

Some Exh operator(s) must still be present in (117). Therefore, still following Sect. 5.1, I assume that *also* makes *comedy*/*tragedy* in (117b) mutually compatible by pruning alternatives from Exh, rather than removing Exh operators entirely:

(118) a. This [ω-$\text{Exh}_{\text{ALT-1}}$ comedy] is also a [ω-$\text{Exh}_{\text{ALT-2}}$ tragedy].
b. ALT-1 = ALT-2 = {~~comedy~~, ~~tragedy~~, epic …}
c. $[\![(118a)]\!] = 1$ iff $\text{tragedy}(\iota x[\text{comedy}(x) \wedge \neg\text{epic}(x)]) \wedge \neg\text{epic}(\iota x[\text{comedy}(x) \wedge \neg\text{epic}(x)])$

As for the data with *and*, one option would be to claim that ω-Exh is always as local as possible on cotaxonyms, but undergoes pruning as with *also*:

(119) This play is both a [ω-$\text{Exh}_{\text{ALT-1}}$ comedy] and a [ω-$\text{Exh}_{\text{ALT-2}}$ tragedy].

Another option (cf. Bade, 2016) would be to define the locality requirement on Exh in such a way as to allow a single ω-Exh to appear above the entire conjunction:

(120) This play is both [ω-Exh_{ALT} a comedy and a tragedy].

This is the option taken in Paillé (2022b, Ch. 6), where it is suggested that if Exh's requirement is that it must be in the minimal XP of the cotaxonym, and if conjunction phrases inherit the label of the conjuncts, (120) is predicted to be possible. ω-Exh's predicate argument entails both *comedy* and *tragedy*, so Exh will exclude alternatives like *epic* without creating a contradiction—no pruning necessary.

6.4 Interim conclusion

In Sect. 5, I suggested that cotaxonymic exclusivity is due to Exh. What I showed in the present section is that this is a different kind of exhaustification effect. In the standard φ-exhaustification observed with sentences containing *most* (for example), Exh is free; it can occur at a distance from the alternative-triggering expression. But with cotaxonyms, Exh is both obligatory and necessarily very local. I called this ω-exhaustification, the insight being that strengthening in natural language is sometimes a process that occurs to predicates ("words") rather than clauses/propositions.

In this article, I do not attempt to formalize ω-exhaustification beyond observing its existence. In Paillé (2022b, Ch. 6), I suggest modelling its twin properties (obligatoriness and locality) through a syntactic Agree relation between derivational morphemes and Exh. Another approach could take inspiration from recent work by Sauerland et al. (2023) also using very local exhaustification on *et*-type nodes, and claim that it is generally the case that every *et*-type node must be exhaustified (expanding on Magri's (2009) claim that every *t*-type node must be exhaustified). I leave this for future work.

The formalization of ω-exhaustification is intrinsically tied with exactly how local to the cotaxonym ω-Exh must be. On the hypothesis that all *et*-type nodes need to be exhaustified, the expectation is that all cotaxonyms (and more expressions too) should be given an ω-Exh operator immediately upon merging into the syntax. The Agree-based approach, I show in Paillé (2022b, Ch. 6), provides a bit more flexibility, due to

the possible percolation of agreement features up the cotaxonym's XP. As mentioned above, this allows a single ω-Exh to scope over an entire conjunction.

7 Discourse-conditioned weakness in predication

This paper so far has attempted to derive strong meanings for predicates through exhaustification. But empirically, weak meanings are often observed too, even in the absence of de-exhaustifiers like *and* or *also*. (121) gives an example for both summative and integrative predicates.[21]

(121) a. Arwa sold every spoon she had.
b. The car is red.

(121a) might suggest that Arwa also sold her sporks (I comment on the judgment below), and (121b) is certainly compatible with the car not being entirely red (presumably the tires and seats are not red, for example). In the following subsections, I take summative and integrative predicates one at a time, and suggest explaining weak meanings in examples like (121) by borrowing an account of weakness in plural predication from Križ (2015).

7.1 Discourse-conditioned weakness with summative predicates

In the literature on part-structure in predication (in particular plurals and summative predicates), the observation of less-than-universal quantification has been called "non-maximality" (e.g., Dowty, 1987; Brisson, 1988; 2003; Lasersohn, 1999; Malamud, 2012; Schwarz, 2013; Križ, 2015; Bar-Lev, 2021; Križ & Spector, 2021); I call it DISCOURSE-CONDITIONED WEAKNESS instead, in anticipation of Sect. 7.2, where I discuss the same phenomenon with integrative predicates. These lack part-quantification, so they are not well framed in terms of "(non-)maximality".

Discourse-conditioned weakness has received considerable attention in the literature on plural predication. Križ (2015) points out that (122) could be true and felicitous in many discourse contexts even if only most (e.g., eight out of ten) of the professors smiled, for example.

(122) The professors smiled.

What leads to this? Let us first appreciate a long-standing (but not uncontroversial; see Bar-Lev 2021) view that plural predication and summative predicates give rise to truth-value gaps (Löbner, 2000; Spector, 2013; Križ, 2015; Križ and Spector, 2021). Indeed, Löbner (2000) assumes that the falsity conditions of a sentence p are the same

[21] (121a) is from a reviewer, who asks if the example might motivate that cotaxonymic exclusivity is a property of co-predications rather than being consistently present in intregrative predicates' meanings. I gave some evidence in Sect. 3 in favour of viewing cotaxonymic exclusivity as present even in non-co-predicational sentences, and make the case in this section that (121a) should be understood in terms of discourse pragmatics, since *spoon* is not in fact necessarily inclusive of sporks in the example.

as the truth conditions of its negation $\neg p$. The falsity conditions of (123a) therefore correspond to the truth conditions of (123b).

(123) a. The flag is green.
≈ 'The flag is all green.'
b. The flag is not green.
≈ 'The flag is not green at all.'

If (123a) is true if the flag is all green, and false if it is not green at all, it must be neither true nor false if the flag is only partly green. The same goes for plural predication, given the truth conditions of positive and negative sentences:

(124) a. The professors smiled.
≈ 'All of the professors smiled.'
b. The professors didn't smile.
≈ 'None of the professors smiled.'

On the exhaustification-based theory of summative predicates presented in this paper, truth-value gaps arise immediately given the claim from Sect. 6.2.3 that Exh is trivalent:

(125) $$[\![\text{This } (a) \text{ is } [\omega\text{-Exh}^{triv}_{\text{ALT}} \cdot \text{ red}]]\!] = \begin{cases} 1, \text{ iff } \mathsf{red}_{\exists}(a) \wedge \neg\mathsf{blue}_{\exists}(a) \wedge \neg\mathsf{white}_{\exists}(a) \wedge \neg \ldots; \\ 0, \text{ iff } \neg\mathsf{red}_{\exists}(a); \\ \#, \text{ otherwise} \end{cases}$$

(125) is neither true nor false if 'this' is only partly red—it has some red, contrary to the falsity conditions, but it is not exclusively red, contrary to the truth conditions. The possibility of adopting the trivalent Exh operator was raised in Sect. 6.2.3 to deal with sentential negation; now, if we accept the existence of truth-value gaps in summative predication, the need for a trivalent Exh becomes even clearer.

Križ (2015) suggests that discourse-conditioned weakness arises from the existence of truth-value gaps. The connection between discourse-conditioned weakness and truth-value gaps can be appreciated from the fact that both of these phenomena disappear with *all* (Križ, 2015):

(126) a. All the professors smiled. (false if eight out of ten smiled)
b. All of the shirt is red. (false if it is partly red)

The basic insight of Križ's (2015, pp. 76ff) theory is that sentences that are neither true nor false in the world of utterance w_0 can be used felicitously if w_0 is, for the purposes of the conversation, equivalent to a world in which the sentence was true. The first ingredient of this theory is the standard assumption (e.g., van Rooij, 2003) that QUDs partition worlds by how they resolve it. Let's consider the sentence (127) with a toy model of three worlds (128) corresponding to different amounts of red on the shirt:

(127) The shirt is red.

(128) $\left\{\begin{array}{ll} w_1: & \text{the shirt is all red,} \\ w_2: & \text{the shirt is half red,} \\ w_3: & \text{the shirt is not red at all} \end{array}\right\}$

If the QUD is 'How much red does the shirt have?' or 'What does the shirt look like?', all of these worlds are in their own cell—they all correspond to different answers to the QUD. On the other hand, if the QUD is 'Does the shirt have any red on it?', w_1 and w_2 both correspond to the answer 'yes', so they are in the same cell.

From here, Križ proposes modifying Grice's (1975, p. 75) Maxim of Quality so that speakers are only prevented from saying things that are false, rather than necessarily needing to say things that are true. Križ suggests that speakers can utter sentences that are neither true nor false as long as they are true in some of the worlds in the cell of the partition containing the world of utterance. That is, a sentence must correctly identify the cell containing the real world, but does not need to identify the real world as such. Thus, QUD permitting, speakers may say things that are neither true nor false in the real world.

This makes it possible to capture weak meanings with summative predicates even while maintaining the claims in this article that they are ω-exhaustified in all sentences. In all discourse contexts, (129a) has the LF in (129b); this is true if the flag is only green, false if it is not green at all, and neither true nor false otherwise.

(129) a. The flag is green.
b. The flag is [$\omega\text{-Exh}^{triv.}_{\text{ALT}}$ green].

The observation of weak meanings is due to the sentence sometimes being usable even if neither true nor false.

This approach to discourse-based weakness is better than apparently obvious alternatives, which would be to derive non-universal meanings for summative predicates by not exhaustifying them at all or removing certain colour terms from their alternatives (cf. Bar-Lev, 2021). These simpler approaches would only manage to produce existential meanings for weak summative predication. But in fact, in many discourse contexts, summative predicates are weaker than universal while still stronger than existential. For example:

(130) a. SCENARIO: *For a temporary art installation, you are making a large mosaic using leaves. There's a part of the drawing that should all be solid orange, but this part is still missing a lot of leaves. People will be looking at the mosaic from a distance to appreciate it as a drawing, so it's okay if the leaves you find are not actually fully orange.*
b. This leaf is orange.
⇒ FELICITOUS *for a leaf that is mostly orange, with some green/brown*
⇒ INFELICITOUS *for a leaf that is mostly green/brown, with some orange*

In (130), not exhaustifying *orange*, or exhaustifying without all basic colour terms as alternatives, would produce an existential meaning. Yet, the sentence is stronger than existential; it means something like 'the leaf has as much orange as I need it to have'.

7.2 Discourse-conditioned weakness with integrative predicates

Integrative predicates have not been described as involving truth-value gaps or anything akin to the discourse-based weakness found in plural predication. I suggest, however, that this is exactly what is at play in the weak meanings observed in sentences like (131), repeated from (121a).

(131) Arwa sold every spoon she had.

Let me first make two points about the empirical picture. First, it happens that *spoon* is in a DE environment in (131), but this is incidental; similar meanings as in (131) can also be observed outside of DE environments. For example:

(132) Arwa only sold spoons.

To my ear, it is possible to interpret (132) as compatible with Arwa having sold spoons and sporks, as long as she did not sell forks or knives. Besides, as the reviewer who provided the example points out about (131), it cannot be claimed that there is generally no ω-exhaustification in the restrictor of *every*, because co-predications maintain the pattern generally reported in this article even when part of *every*'s restrictor:

(133) Arwa sold every spoon that is #(also) a fork that she had.

The second empirical point about (131) is that, while it can be interpreted as meaning that Arwa sold her sporks too, this is certainly not necessary. Whether *spoon* in (131) is interpreted as inclusive or exclusive of sporks seems entirely discourse-dependent. To bring out an interpretation where *spoon* is exclusive of sporks, consider the longer sentence in (134):

(134) Arwa sold every spoon she had, but she still has all her forks, sporks, and knives.

Thus, even in examples where an integrative predicate seems weak, this is conditioned by discourse, including speaker preferences about what to group together.

I therefore analyze the possibly weaker meaning of the integrative *spoon* in (131) as a case of discourse-based weakness. Integrative predicates are ω-exhaustified with a trivalent Exh, leading to truth-value gaps. Such gaps arise when a predicate P holds of an individual but there is also another predicate Q (a cotaxonym of P) that holds of the same individual, e.g.:

(135) $[\![\textit{Macbeth}\text{ is a }[\omega\text{-Exh}^{triv}_{\text{ALT}}\text{ tragedy}]]\!] = \begin{cases} 1, \text{ iff } \mathsf{tragedy}(m) \wedge \neg\mathsf{comedy}(m) \wedge \neg \ldots; \\ 0, \text{ iff } \neg\mathsf{tragedy}(m); \\ \#, \text{ otherwise} \end{cases}$

That is, *Macbeth is a tragedy* is true if it is only a tragedy, false if it is not a tragedy at all (regardless of the status of other genre-predicates), and neither true nor false if *Macbeth* is both a tragedy and some other genre. In Paillé (to appear), I defend empirically the existence of a truth-value gap in examples like (135) based in part on data independent of the discourse-conditioned weakness I focus on here.

From the existence of truth-value gaps with integrative predicates (135), Križ's mechanism kicks in as for summative predicates, predicting weakness in discourse environments where being 'only *P*' is equivalent to being '*P* and maybe also *Q*'. The optionality I just described in whether the predicate *spoon* excludes or includes sporks in examples like (131) can be understood in terms of whether the speaker considers sporks to be equivalent to spoons for certain purposes.

As already discussed for summative predicates, this solution is more complicated than a conceivable alternative, which would be to claim that examples like (131), repeated in (136), show that ω-Exh is in fact optional at least in certain conditions.

(136) Arwa sold every spoon she had.

But as pointed out above, postulating optional exhaustification would make it entirely mysterious why *also* is still required if one were to co-predicate *spoon* and *fork* in (136):

(137) Arwa sold every spoon that is #(also) a fork that she had.

Križ's pragmatic mechanism is therefore the preferable option.

7.3 Section conclusion

In this section, I discussed cases where, contrary to the claims so far in the article, cotaxonyms are interpreted as weak even in the absence of de-exhaustification. I explained these via the pragmatic weakening effect independently postulated by Križ (2015) for plural predication. To do this, I relied on a trivalent Exh, as I did in Sect. 6 to explain the lack of observable strengthening under negative-flavoured DE operators.

The mechanism to create weakness that I just adopted might appear to create a problem for co-predications like (138) (or (137), for that matter). Indeed, if both cotaxonyms can be interpreted as weak, then *also* is no longer expected to be required.

(138) a. This fork is #(also) a spoon.
b. The white flag is #(also) green.

Why are the cotaxonyms apparently necessarily strong? Recall that Križ's mechanism is a global, pragmatic mechanism, and it relies crucially on the existence of worlds in which the sentence is true (it allows a sentence to be uttered in a world in which it is neither true nor false, as long as it is in the same cell as the worlds where it is true). But the sentences in (138) are logical contradictions due to the presence of ω-Exh operators. There are no worlds in which they are true. Thus, it is impossible for the real world to be in the same cell as a world in which they would be true. A de-exhaustifier is thus predicted to always be required with cotaxonymic co-predications.

This concludes my discussion of the strength of predicates; in the next section, I turn to clarifying what my claims about exhaustification in language are, and what they are not.

8 Cotaxonymic exclusivity is not contrastive focus

In Sect. 2, I claimed that the standard view of the strengthening of predicates is that it occurs in two circumstances: when predicates are part of a Horn scale, and when they are contrastively (intonationally) focused. In this section, I drive home the point that the proposed ω-exhaustification of predicates leading to cotaxonymic exclusivity constitutes a new kind of strengthening effect with predicates, by showing that it differs from either of the processes described in Sect. 2. It is straightforward to appreciate that the creation of cotaxonymic exclusivity does not involve Horn scales: *comedy* and *tragedy* are not on an entailment scale, and in fact they are logically independent (given the claim in this article that they are weak), other than both being kinds of genres. But what about contrastive focus?

The next subsections point out two differences between the exhaustification leading to cotaxonymic exclusivity (henceforth CE-exhaustivity) and the exhaustification leading to contrastive focus on predicates (CF-exhaustivity). To be clear, I am not assuming or positing two different kinds of exhaustification in addition to the distinction between φ-exhaustification and ω-exhaustification; my claim is that the exhaustification leading to the intuition of cotaxonymic exclusivity arises from ω-Exh, while the exhaustification leading to the intuition of contrastive focus is an instance of φ-Exh. The first way to observe that CE-exhaustivity is different from CF-exhaustivity is the descriptively obvious point that only contrastive focus necessarily involves focus intonation. The second is the fact that CE-exhaustivity and CF-exhaustivity are computed at different syntactic scopes; CE-exhaustivity is an ω-exhaustification phenomenon, while CF-exhaustivity is a φ-exhaustification phenomenon computed over and above CE-exhaustivity. A contrastively focused cotaxonym is therefore exhaustified twice.

8.1 The first difference: focus intonation

The apparently obvious way to argue that CE-exhaustivity is different from CF-exhaustivity is to show that cotaxonymic exclusivity persists even when a predicate is clearly not intonationally focused:

(139) a. Do you like the LONG_F comedy or the SHORT_F comedy?
b. Do you want the white FLAG_F or the white TABLE_F?

(139a) refers to real comedies and (139b) refers to an entirely white flag or table, so the relevant cotaxonyms are exhaustified in (139) even without being focused.

Another way to differentiate CE-exhaustivity from CF-exhaustivity is to construct a minimal pair where one sentence has contrastive focus on a cotaxonym and the other does not:

(140) a. The comedy was good.
b. The COMEDY_F was good.

Both (140a) and (140b) at least strongly suggest (on my theory, entail) that the play is a true comedy, i.e., not a tragicomedy. But (140b) clearly has additional meaning, namely that some other thing (most saliently a play of another genre) was *not* good.

Straightforwardly, if contrastive focus and cotaxonymic exclusivity were the same effect, the difference between (140a) and (140b) could not be captured.

Thus, cotaxonymic exclusivity is more like the "anti-universal" meaning of *most* than the meaning arising from contrastive intonational focus on predicates. Indeed, *most* and cotaxonyms pattern together in not needing to be intonationally focused to signal the presence of alternatives:

(141) a. Arwa saw most of the children.
⇒ *excludes a universal alternative without intonational focus on 'most'*
b. This is a comedy.
⇒ *excludes cotaxonymic alternatives without needing intonational focus on 'comedy'*

8.2 The second difference: syntactic scope

The second difference between CE-exhaustivity and CF-exhaustivity is that CE-exhaustivity is an obligatory effect that is necessarily computed locally (giving it a lexical-like flavour), while CF-exhaustivity is optional and not necessarily local. In other words, cotaxonymic exclusivity is due to ω-exhaustivity, while contrastive focus on predicates is due to φ-exhaustivity.

Consider (142), where *white* is only contrastively focused in (142b) (a similar point could be made for (140)).

(142) a. The white flag is high.
b. The WHITE$_F$ flag is high.

In both cases, *white* is exclusive of other colour terms (the flag under discussion is entirely white), so it must be exhaustified. Let's start by focusing on (142a). The Exh strengthening *white* must be local, i.e., below *the*:

(143) a. $[\![$The [ω-Exh$_{\text{ALT}}$ white] flag is high$]\!]$
$= 1$ iff the $\left(\begin{array}{c}\text{white}_\exists\ \&\\ \text{not green}_\exists\ \&\\ \text{not red}_\exists\end{array}\right)$ flag is high.
b. *$[\![$φ-Exh$_{\text{ALT}}$ [the white flag is high]$]\!]$
$= 1$ iff $\left\{\begin{array}{l}\text{the white}_\exists \text{ flag is high } \wedge\\ \neg\text{the green}_\exists \text{ flag is high } \wedge\\ \neg\text{the red}_\exists \text{ flag is high}\end{array}\right.$

(143a) corresponds to the intuited meaning, while (143b) is problematic in two ways: it carries non-intuited inferences about other flags, and it does not make *white* universal. Thus, as already established in Sect. 6, (142a) involves an ultra-local ω-Exh—the LF in (143a) is the only one available.

But now consider (142b), with contrastive focus on *white*. Here too, *white* is universal, so an ω-Exh below *the* is still needed. But unlike (142a), in (142b) there *is* an inference about some other flag: the contribution of the contrastive focus is to mean that there is some other, non-white colour such that it is *not* the case that the unique flag

that is entirely of that colour is high. Which colour this happens to be depends on what the speaker has in mind; for concreteness, imagine the alternative colour is green. To capture that *white* is universal and that there is also an inference about another flag, we need two Exh operators: the ultra-local ω-Exh below *the*, and a free φ-Exh computing contrastive focus above *the*:

(144) $[\![\varphi\text{-Exh}_{\text{ALT}}$ [the [$\omega\text{-Exh}_{\text{ALT}}$ white] flag is high]$]\!]$
= 1 iff the only-white flag is high $\wedge\neg$[the only-green flag is high].

The ω-Exh below *the* takes the set of colour terms as its alternatives, deriving the universal/exclusive meaning for *white*; the global φ-Exh takes only sentences obtained by replacing *white* with the contextually salient alternative *green*, but not the other colours.

Thus, CE-exhaustivity and CF-exhaustivity differ in two ways. Descriptively, there is a difference in intonation and meaning between these effects. But they also have a different syntax–semantics; CF-exhaustivity (a case of φ-exhaustification) is computed over and above CE-exhaustivity (a case of ω-exhaustification), and these effects involve different sets of alternatives: cotaxonymic for the latter but discourse-dependent for the former (but see Sect. 5.2, where I suggest that the kind of information provided by a predicate, rather than cotaxonymy, determines alternativehood for CE-exhaustivity/ω-exhaustification).

9 Conclusion

In this article, I have shown that there is a systematic mismatch between the meaning of content vocabulary items as present in sentences, and their underlying lexical–conceptual meaning. Many (perhaps virtually all) cotaxonyms have overlap in meaning lexically, but due to an ultra-local exhaustification effect, they come to delimit each other's meaning in the semantic composition. This is a universal phenomenon, as far as I am aware:

(145) Cette comédie est #(aussi) une tragédie. (French)
this comedy is also a tragedy

This exhaustification effect, due to being obligatory and necessarily local, is easily mistaken as part of predicates' lexical meaning—which is presumably why it has not been noticed until now. Semantically, the mutual exclusion of two given cotaxonyms disappears in two situations. This article has focused on when it disappears with *and* or *also* in positive sentences (146), but we also saw in Sect. 6.2.3 that it disappears under *not* (147).

(146) a. The flag is white and green.
b. This comedy is also a tragedy.

(147) a. The flag is not white.
b. The play is not a comedy.

It is only in (147) that we directly observe cotaxonyms' lexical meanings, however. The conjunctive elements *and* and *also* (146) can make two given cotaxonyms lexically consistent, but they keep them inconsistent with other cotaxonyms like *red* for (146a) or *epic* for (146b). A white and green flag is *only* white and green, and uttering that a comedy is also a tragedy entails that it is not an epic (unless a subsequent additive prunes that alternative—that is, unless one goes on to say it is 'also an epic'). Thus, in positive sentences, cotaxonyms' weak lexical meanings can only be inferred, never observed.

This paper is in many ways programmatic, and opens up a number of areas for future work. The three important domains of research are the following. First, what exactly is the lexical meaning of integrative predicates? For summative predicates, we have learned in this paper that they are existential; for integrative predicates, while we have learned that cotaxonyms can be lexically–conceptually consistent, this paper has not modelled what the conceptual underpinning of their lexical meaning actually is. I have shown that the concept COMEDY associated with the lexical meaning of *comedy* must be compatible with the concept TRAGEDY. This is clearly a step forward in understanding concepts (and goes against certain approaches to concepts, as mentioned in Sect. 3.2), but leaves open how best to model the meaning of such concepts.[22]

The second important domain pertains to the formal properties of exhaustification, including its locality, trivalency, and obligatoriness. As mentioned in Sect. 6, while I suggest that *in the basic case* cotaxonyms are strengthened so locally that ω-Exh has nothing but the cotaxonym as its predicative argument, there are open questions and possibilities about how consistent this is (see in particular the discussion of conjunction in Sect. 6.3) and how to model why such a locality constraint would hold. Likewise, the suggestion to adopt a trivalent Exh to deal with negative sentences entails a general rethinking of exhaustification effects (Bassi et al., 2021).

The third domain of research affected by this paper pertains to de-exhaustifiers—at least *and* and *also*. For both, I have only sketched out possibilities for how they de-exhaustify. For *and* specifically, there is another topic of inquiry: my discussion relied on the suggestion that it is always intersective when it conjoins predicates that have an atomic argument (as in *the flag is white and green*); future work should integrate this claim with already-existing theories of *and*'s lexical meaning and the apparent presence of intersective and non-intersective meanings for conjunctions.

The finding that cotaxonyms have a different meaning in sentences than in the lexicon or the conceptual module has foundational consequences for the relationship between language and mind. A tacit consensus in cognitive science, linguistics, psychology, and philosophy takes the nature of concepts to be observable directly from the meaning of predicates in simple sentences. In fact, language directly affects our intuitions of concepts, making them appear narrower in meaning than they really are.

Acknowledgements First, I thank my doctoral supervisors, Bernhard Schwarz and Luis Alonso-Ovalle—this article would not exist without years of supervision and mentorship by them, and a steady stream of input on almost everything in it. I also thank Aron Hirsch for many helpful meetings, and Junko Shimoyama for the class on copular sentences at McGill from which this project first emerged. For helpful meetings, discussion, and encouragement, I thank Viola Schmitt, Uli Sauerland, Nina Haslinger, Alexander Göbel, Nadine Bade, Michael Wagner, Roni Katzir, Vera Hohaus, Justin Royer, and Jacob Hoover. Finally, I thank

[22] For an overview of various theories of concepts, see Margolis and Laurence (2021).

the reviewers and audiences at the various conferences where parts of this work have been presented, as well as three anonymous *Linguistics and Philosophy* reviewers. All shortcomings remain my own.

Declarations

Conflict of interest The author has no conflicts of interest to declare.

References

Amsili, P. & Beyssade, C. (2006). Compulsory presupposition in discourse. In Sidner, C., Harpur, J., Benz, A., & Kühnlein, P., (Eds.), *Proceedings of the Workshop on Constraints in Discourse*, (pp. 5–10), Maynooth. National University of Ireland.

Aravind, A. & Hackl, M. (2017). Against a unified treatment of obligatory presupposition trigger effects. In Burgdorf, D., Collard, J., Maspong, S., & Stefánsdóttir, B., (Eds.), *Proceedings of SALT 27*, (pp. 173–190), Washington, DC. Linguistic Society of America.

Bade, N. (2016). *Obligatory presupposition triggers in discourse: Empirical foundations of the theories Maximize Presupposition and Obligatory Implicatures*. PhD thesis, University of Tübingen, Tübingen.

Bar-Lev, M. E. (2018). *Free Choice, Homogeneity, and Innocent Inclusion*. PhD thesis, The Hebrew University of Jerusalem, Jerusalem.

Bar-Lev, M. E. (2021). An implicature account of Homogeneity and Non-maximality. *Linguistics and Philosophy, 44*(10), 1045–1097.

Bar-Lev, M. E. & Fox, D. (2017). Universal Free Choice and Innocent Inclusion. In Burgdorf, D., Collard, J., Maspong, S., & Stefánsdóttir, B., (Eds.), *Proceedings of SALT 27*, (pp. 95–115).

Bassi, I., Del Pinal, G., & Sauerland, U. (2021). Presuppositional exhaustification. *Semantics and Pragmatics, 14*(11), 1–48.

Becker, T. (2002). Autohyponymy: Implicature in lexical semantics, word formation, and grammar. *Journal of Germanic Linguistics, 14*(2), 105–136.

Berlin, B., & Kay, P. (1969). *Basic colour terms: Their universality and evolution*. Berkeley: University of California Press.

Blair, M., & Homa, D. (2001). Expanding the search for a linear separability constraint on category learning. *Memory & Cognition, 29*(8), 1153–1164.

Blutner, R. (1998). Lexical pragmatics. *Journal of Semantics, 15*(2), 115–162.

Blutner, R. (2002). Lexical semantics and pragmatics. In Hamm, F. & Zimmermann, T. E., (Eds.), *Linguistische Berichte Sonderheft 10. Semantics*, (pp. 27–58). Helmut Buske Verlag, Hamburg.

Blutner, R. (2004). Pragmatics and the lexicon. In L. R. Horn & G. Ward (Eds.), *The handbook of pragmatics* (pp. 488–514). Malden, MA: Blackwell.

Brisson, C. (1998). *Distributivity, maximality, and floating quantifiers*. PhD thesis, Rutgers University, New Brunswick, NJ.

Brisson, C. (2003). Plurals, *all*, and the nonuniformity of collective predication. *Linguistics and Philosophy, 26*(2), 129–184.

Cann, R. (2011). Sense relations. In C. Maienborn, K. V. Heusinger, & P. Portner (Eds.), *Semantics: An international handbook of natural language meaning* (Vol. 1, pp. 456–479). Berlin: de Gruyter Mouton.

Champollion, L. (2016). Ten men and women got married today: Noun coordination and the intersective theory of conjunction. *Journal of Semantics, 33*(3), 561–622.

Chemla, E. (2008). An epistemic step for anti-presuppositions. *Journal of Semantics, 25*(2), 141–173.

Chierchia, G. (2004). Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In A. Belletti (Ed.), *Structures and beyond: The cartography of syntactic structures* (Vol. 3, pp. 39–103). Oxford: Oxford University Press.

Chierchia, G. (2006). Broaden your views: Implicatures of domain widening and the "logicality" of language. *Linguistic Inquiry, 37*, 535–590.

Chierchia, G., Fox, D., & Spector, B. (2012). Scalar implicatures as a grammatical phenomenon. In C. Maienborn, K. von Heusinger, & P. Portner (Eds.), *Semantics: An international handbook of natural language meaning* (Vol. 3, pp. 2297–2331). Berlin: De Gruyter Mouton.

Chomsky, N. (1957). *Syntactic structures*. The Hague: Mouton.

Cruse, A. (2002). Hyponymy and its varieties. In R. Green, C. A. Bean, & S. H. Myaeng (Eds.), *The semantics of relationships: An interdisciplinary perspective* (pp. 3–21). Dordrecht: Springer.
de Saussure, F. (2011[1916]). *Course in general linguistics*. Columbia University Press, New York.
Doetjes, J. (1997). *Quantifiers and selection: On the distribution of quantifying expressions in French, Dutch and English*. PhD thesis, Leiden University.
Dowty, D. R. (1987). Collective predicates, distributive predicates, and *all*. In Miller, A. & Zhang, Z.-S., (Eds.), *Proceedings of the 3rd Eastern States Conference on Linguistics (ESCOL)*, (pp. 97–115), Columbus, OH. Ohio State University.
Fitzpatrick, J. M. (2006). *The syntactic and semantic roots of floating quantification*. PhD thesis, MIT, Cambridge, MA.
Fox, D. (2007). Free choice and the theory of scalar implicatures. In U. Sauerland & P. Stateva (Eds.), *Presupposition and implicature in compositional semantics* (pp. 71–120). Basingstoke, Hampshire (UK): Palgrave Macmillan.
Fox, D., & Spector, B. (2018). Economy and embedded exhaustification. *Natural Language Semantics, 26*, 1–50.
Gajewski, J. (2005). *Neg-raising: Polarity and presupposition*. PhD thesis, MIT, Cambridge, MA.
Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought*. Cambridge, MA: MIT Press.
Gazdar, G. (1979). *Pragmatics: Implicature, presupposition, and logical form*. New York: Academic Press.
Geurts, B. (2010). *Quantity implicatures*. Cambridge: Cambridge University Press.
Grice, H. P. (1975). Logic and conversation. In C. Peter & J. L. Morgan (Eds.), *Syntax and semantics, volume 3: Speech acts* (pp. 41–58). New York: Academic Press.
Hampton, J. A. (2015). Categories, prototypes and exemplars. In N. Riemer (Ed.), *The Routledge handbook of semantics* (pp. 125–141). Abingdon, UK: Routledge.
Harnish, R. (1976). Logical form and implicature. In T. G. Bever, J. J. Katz, & D. T. Langendoen (Eds.), *An integrated theory of linguistic ability* (pp. 313–391). New York: Thomas Cromwell.
Heim, I. (1991). Artikel und Definitheit [Articles and definiteness]. In A. von Stechow & D. Wunderlich (Eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung* (pp. 487–535). Berlin: de Gruyter.
Heycock, C., & Zamparelli, R. (2005). Friends and colleagues: Plurality, coordination, and the structure of DP. *Natural Language Semantics, 13*, 201–270.
Horn, L. (1972). *On the semantic properties of logical operators in English*. PhD thesis, UCLA, Los Angeles.
Horn, L. (1984). Ambiguity, negation, and the London school of parsimony. In Jones, C. & Sells, P., (Eds.), *Proceedings of NELS 14*, (pp. 108–131), Amherst, MA. GLSA (Graduate Linguistics Student Association), University of Massachusetts, Amherst.
Horn, L. (2017). Pragmatics and the lexicon. In Y. Huang (Ed.), *The Oxford handbook of pragmatics* (pp. 511–531). Oxford: Oxford University Press.
Krifka, M. (1990). Boolean and non-Boolean 'and'. In L. Kálmán & L. Pólos (Eds.), *Papers from the Second Symposium on Logic and Language* (pp. 161–188). Budapest. Akadémiai Kiadó.
Krifka, M. (1996). Pragmatic strengthening in plural predications and donkey sentences. In Galloway, T. & Spence, J., (Eds.), *Proceedings of SALT 6*, (pp. 136–153), Ithaca, NY. Cornell University.
Krifka, M. (1998). Additive particles under stress. In: Strolovitch, D. & Lawson, A., (Eds.), *Proceedings of SALT 8*, (pp. 111–129), Ithaca, NY. Cornell University.
Kripke, S. A. (2009 [1990]). Presupposition and anaphora: Remarks on the formulation of the projection problem. *Linguistic Inquiry*, *40*(3), 367–386.
Križ, M. (2015). *Aspects of homogeneity in the semantics of natural language*. PhD thesis, University of Vienna, Vienna.
Križ, M., & Spector, B. (2021). Interpreting plural prediction: homogeneity and non-maximality. *Linguistics and Philosophy, 44*, 1131–1178.
Lambert, K., & van Fraassen, B. C. (1970). Meaning telations, possible objects, and possible worlds. In K. Lambert (Ed.), *Philosophical problems in logic* (pp. 1–19). Dordrecht: Springer.
Lasersohn, P. (1995). *Plurality, conjunction and events*. Dordrecht: Kluwer Academic Publishers.
Lasersohn, P. (1999). Pragmatic halos. *Language, 75*(3), 522–551.
Levering, K. R., Conaway, N., & Kurtz, K. J. (2020). Revisiting the linear separability constraint: New implications for theories of human category learning. *Memory & Cognition, 48*, 335–347.
Levinson, S. C. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
Löbner, S. (2000). Polarity in natural language: Predication, quantification and negation in particular and characterizing sentences. *Linguistics and Philosophy, 23*, 213–308.

Magri, G. (2009). *A theory of individual-level predicates based on blind mandatory implicatures. Constraint promotion for Optimality Theory*. PhD thesis, MIT, Cambridge, MA.
Malamud, S. A. (2012). The meaning of plural definites: A decision-theoretic approach. *Semantics & Pragmatics*, 5:Article 3.
Margolis, E. & Laurence, S. (2021). Concepts. In Zalta, E. N., (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University, Spring 2021 edition.
Paillé, M. (2022a). On additives' interaction with exhaustivity: the view from negative continuations. In Bakay, Ö., Pratley, B., Neu, E., & Deal, P., (Eds.), *Proceedings of the fifty-second annual meeting of the North East Linguistic Society*, volume 3, pages 1–10, Amherst, MA. GLSA (Graduate Linguistics Student Association), University of Massachusetts, Amherst.
Paillé, M. (2022b). *Strengthening predicates*. PhD thesis, McGill University, Montréal.
Paillé, M. (2023). Alternatives and jurisdiction in predication. In Onoeva, M., Staňková, A., & Šimík, R., (Eds.), *Proceedings of Sinn und Bedeutung 27*, pages 483–499, Prague. Charles University.
Paillé, M. (to appear). Truth-value gaps with integrative predicates. In Kiper, Y., Lyu, L., Wang, R., & Webster, N., (Eds.), *Proceedings of the 41st West Coast Conference on Formal Linguistics*, Somerville, MA. Cascadilla Press.
Paillé, M. (2024). Co-predications and the quantificational force of summative predicates. *Journal of Semantics*.
Poortman, E. B. (2017). Concept typicality and the interpretation of plural predicate conjunction. In J. A. Hampton & Y. Winter (Eds.), *Compositionality and concepts in linguistics and psychology* (pp. 139–162). Springer.
Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: MIT Press.
Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Oxford University Press.
Rohdenburg, G. (1985). Dogs, bitches and other creatures. *Journal of Semantics, 4*(2), 117–135.
Russell, B. (2006). Against grammatical computation of scalar implicatures. *Journal of Semantics, 23*, 361–382.
Ruts, W., Storms, G., & Hampton, J. A. (2004). Linear separability in superordinate natural language concepts. *Memory & Cognition, 32*(1), 83–95.
Sæbø, K. J. (2004). Conversational contrast and conventional parallel: Topic implicatures and additive presuppositions. *Journal of Semantics, 21*, 199–217.
Sauerland, U. (2004). Scalar implicatures in complex sentences. *Linguistics and Philosophy, 27*, 367–391.
Sauerland, U. (2008). Implicated presuppositions. In A. Steube (Ed.), *Language, context and cognition, volume 8: The discourse potential of underspecified structures* (pp. 581–599). Berlin: Mouton de Gruyter.
Sauerland, U. (2012). The computation of scalar implicatures: Pragmatic, lexical or grammatical? *Language and Linguistics Compass, 6*(1), 36–49.
Sauerland, U., Bassi, I., Bill, C., Bimpeh, A., Hirsch, A., Jeretič, P., Meyer, M.-C., Nicolae, A., Yatsushiro, K., & Alexiadou, A. (2023). An Algebra of Thought that predicts key aspects of language structure. In Hrdinková, K., Onoeva, M., Staňková, A., & Šimík, R., (eds), *Proceedings of Sinn und Bedeutung 27*, pages 589–608, Prague. Charles University.
Schein, B. (2017). *'And': Conjunction reduction redux*. Cambridge, MA: MIT Press.
Schmitt, V. (2013). *More pluralities*. PhD thesis, University of Vienna, Vienna.
Schmitt, V. (2019). Pluralities across categories and plural projection. *Semantics & Pragmatics, 12*, 1–55.
Schmitt, V. (2021). Boolean and non-Boolean conjunction. In D. Gutzmann, L. Matthewson, C. Meier, H. Rullmann, & T. E. Zimmermann (Eds.), *The Wiley Blackwell companion to semantics* (pp. 1–32). Hoboken, NJ: John Wiley & Sons Inc.
Schwarz, F. (2013). Maximality and definite plurals: Experimental evidence. In Chemla, E., Homer, V., & Winterstein, G., (eds), *Proceedings of Sinn und Bedeutung 17*, pages 509–526.
Schwarzschild, R. (1996). *Pluralities*. Boston: Kluwer Academic.
Singh, R. (2008). *Modularity and locality in interpretation*. PhD thesis, MIT, Cambridge, MA.
Singh, R. (2011). *Maximize Presupposition!* and local contexts. *Natural Language Semantics, 19*, 149–168.
Smith, J. D., Murray, M. J., & Minda, J. P. (1997). Straight talk about linear separability. *Journal of Experimental Psychology, 23*(3), 659–680.
Spector, B. (2003). Scalar implicatures: Exhaustivity and Gricean reasoning. In M. Aloni, A. Butler, & P. Dekker (Eds.), *Questions in dynamic semantics* (pp. 225–249). Bingley, UK: Emerald Group Publishing.

Spector, B. (2013). Homogeneity and plurals: From the Strongest Meaning Hypothesis to supervaluations. Talk given at *Sinn und Bedeutung 18*. University of the Basque Country, Vitoria-Gasteiz.
Stalnaker, R. (1981). Anti-essentialism. *Midwest Studies of Philosophy, 4*, 343–355.
Steindel Burdin, R., & Tyler, J. (2018). Rises inform, and plateaus remind: Exploring the epistemic meanings of "list intonation" in American English. *Journal of Pragmatics, 136*, 97–114.
van Fraassen, B. C. (1967). Meaning relations among predicates. *Noûs, 1*(2), 161–179.
van Rooij, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy, 26*, 727–763.
van Rooij, R., & Schulz, K. (2004). Exhaustive interpretation of complex sentences. *Journal of Logic, Language and Information, 13*, 491–519.
Waldon, B., Condoravdi, C., Levin, B., & Degen, J. (2023). On the context dependence of artifact noun interpretation. In Hrdinková, K., Onoeva, M., Staňková, A., & Šimík, R., (eds), *Proceedings of Sinn und Bedeutung 27*, pages 674–692, Prague. Charles University.
Wilson, D., & Carston, R. (2007). A unitary approach to lexical pragmatics: Relevance, inference and ad hoc concepts. In N. Burton-Roberts (Ed.), *Pragmatics* (pp. 230–259). New York: Palgrave-Macmillan.
Winter, Y. (2001). *Flexibility principles in Boolean semantics*. Cambridge, MA: MIT Press.