Lexical Illusions, (Non-)Maximality, and Invisible Gaps

by

Diego Feinmann

A thesis submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

The University of Sheffield
Faculty of Arts and Humanities
Department of Philosophy

March 2020

ABSTRACT

The first part of this work (Chapters I, II and III) is aimed at testing the lexical modulation hypothesis, which has it that word meanings can be adjusted, either narrowed or broadened, during the process of semantic composition in response to pragmatic pressures. I examine many of the examples that have been given in support of this hypothesis and, through the application of linguistic tests, try to determine whether these examples can, as a matter of fact, be taken as evidence that something like modulation is operative in natural language. The examination of the data will lead me to the conclusion that the lexical modulation hypothesis is likely to be false. The second part of this work (Chapters IV, V, and VI) is concerned with the phenomenon of (im)precision, which many have analysed as an instance of lexical modulation. I develop a formal account of (im)precise interpretation, which builds upon Križ's (2015) seminal work on (non-)maximality. Furthermore, I show how homogeneity, (arguably) a semantic pre-requisite for (im)precise interpretation, is implicated in constructions that are not obviously homogenous, such as absolute adjectives.

ACKNOWLEDGEMENTS

I shall be brief, despite being indebted to so many people. Thanks to Jennifer Saul and Rosanna Keefe, for their patience and dedication, for reading (and insightfully commenting on) countless drafts of mine, and for giving me complete freedom to pursue my research interests. Thanks to Benjamin Spector, for investing so much time in me (many hours per week over almost three years), for helping me to develop my ideas, for teaching me formal semantics and probability theory, and for being so incredibly kind and supportive. This thesis would not have been possible without Benjamin's guidance. Thanks to Daniel Rothschild, for giving me a sense of direction that I altogether lacked when I started to work on this project, for reading my early drafts (some of them very bad), for putting me in touch with the Paris crowd, for his thought-provoking feedback (which greatly enhanced the quality of this work), and for arranging for me to give talks at UCL and Oxford. Thanks to Manuel Križ, for writing such a brilliant thesis (which the present work is hugely indebted to), for responding to all my e-mails and WhatsApp messages, for always having something profound to say (no matter the topic), and for his friendship. Thanks to Philippe Schlenker, for his willingness to discuss my work, for his insightful comments and suggestions, and for being kind and generous. Thanks to Salvador Mascarenas, for his great semantics courses, for all his encouragement and support, for his challenging feedback, and for always being there to lend a helping hand.

Many other people helped me along the way; in no particular order, thanks to: Andrés Soria Ruiz, Mora Maldonado, Amir Anvari, Milica Denić, Daniel Hoek, Leon Bergen, Matt Mandelkern, Diego Tajer, Angelika Kratzer, Emmanuel Chemla, Tadeusz Ciecierski, Robyn Carston, François Recanati, Daniel Kricheff, Tristan Thommen, Kajsa Djärv, Jack Byron, Michael Murez, Ramiro Caso, Andrés Saab, Moshe Elyashiv Bar-Lev, Stephen Laurence, Eleonora Orlando, Heather Burnett, Nathan Klinedinst, James Lewis, Pawel Grabarczyk, Ewan Dunbar, Jeremy Kuhn, Brian Buccola, Émile Enguehard, Janek Guerrini, Cécile Crimon, Nadine Bade, WooJin Chung, Louis Rouillé, David Blunier, Alexandre Cremers, and Luca Gasparri.

This research was made possible by a University of Sheffield's *University Scholarship*, which I am honoured to have been awarded.

All my life I have been clumsily following the steps of a shadow I can hardly discern. This is the shadow of my grandfather, Alberto Claudio Blasetti. He would sit next to me, and talk for hours, about Rilke, about Saint-John Perse, about Dante… How could one ever be grateful enough for such a gift?

FOR MARY. FOR EVERYTHING.

TABLE OF CONTENTS

DECLARATION

*I, Diego Feinmann, confirm that the Thesis is my own work. I am aware of the University's Guidance on the Use of Unfair Means (*www.sheffield.ac.uk/ssid/unfair-means*). This work has not been previously been presented for an award at this, or any other, university.*

Chapter I
INTRODUCTION

1 THE LEXICAL MODULATION HYPOTHESIS

When words are said to be context-sensitive, what is it meant exactly? Typically, two things: that a word is ambiguous, and that the resolution of this ambiguity can depend on contextual factors, or that a word has indexical features, and that the determination of its actual denotation requires access to contextual information. The former class includes homonyms, such as 'bank' (*financial intuition*) and 'bank' (*river bank*), but also more subtle instances of ambiguity;[1] the latter class includes *dedicated* indexical expressions such as 'here', 'I' and 'this', but also lexical items such as relative adjectives (e.g. 'tall', 'hot'), relational nouns (e.g. 'enemy', 'foreigner') as well as relational adjectives (e.g. 'foreign', 'local'), all of which require the provision of a contextual standard or parameter (cf. Mitchell 1986, Partee 1989, Kennedy and McNally 2005, among others).

Advocates of what has come to be known as *truth-conditional pragmatics*,[2] such as Sperber and Wilson (1986), Carston (2002), and Recanati (2004, 2010), among others, believe that words are context-sensitive in yet another sense. According to these theorists, whom I will refer to with the term *contextualists*, the meaning of (disambiguated) content words can be adjusted during the process of semantic composition to maximise contextual fitness. This proposal goes by the name of *lexical modulation* (LM, henceforth).

To illustrate the (alleged) workings of this (alleged) mechanism, consider (1)-(3):

(1) I would like to live in a world without *banks*.

(2) I need to get some money… Do you know of any *bank* around here?

(3) Context: In a meeting of the heads of the G20's Central Banks, one of the heads utters…
Every *bank* should raise the interest rates, there's no other solution.

The term 'bank' would typically be interpreted as *bank institution* in (1), as *retail bank* in (2), and as *central bank from a G20 country* in (3). Since 'bank' is not semantically ambiguous between *bank institution*, *retail bank,* and *central bank from a G20 country*, it is tempting to conclude that the meaning of 'bank', once disambiguated in favour of the *bank institution* sense (and against the *river bank* sense),

[1] For example, 'file' (*computer file*) and 'file' (*document in a cabinet box*).
[2] The term originates in Recanati (1993).

can be either interpreted as meaning *bank institution*—as in (1)—, or subjected to further pragmatic refinements—as in (2) and (3); that is, it seems reasonable to suppose that considerations about what the speaker means, in addition to disambiguating 'bank', can modulate the denotation of its disambiguated senses.[3]

For contextualists, therefore, LM is a *compositional* mechanism, just like lexical disambiguation. In plain terms, what this means is that LM is believed to interact with the recursive semantics and, correspondingly, play a role in determining the truth-conditional profile of utterances (i.e. what Grice (1975) termed 'what is said' and I will simply refer to as 'the literal meaning'[4]). That this may be the case should not call into question the principle of compositionality; if the traditional version of the principle has it that the meaning of a complex expression is 'a function of the disambiguated meanings of its parts and the way they are put together', the LM-friendly version would have it that the meaning of a complex expression is 'a function of the modulated disambiguated meanings of its parts and the way they are put together'. See Recanati (2009) for discussion on this point.

Two types of LM are often distinguished in the literature: *lexical narrowing*, on the one hand, and *lexical broadening or loosening*, on the other. For example, according to Wilson and Carston (2007: 6),

> lexical narrowing involves the use of a word to convey a more specific sense than the encoded one, with a more restricted denotation (picking out a subset of the items that fall under the encoded concept). Lexical broadening involves the use of a word to convey a more general sense than the encoded one, with a consequent expansion of the linguistically specified denotation.

Similar characterisations can be found in Carston and Powell (2008), Recanati (2010), and Ludlow (2014). (2) and (3), under the assumption that lexical modulation is at work in these examples (and assumption that I will dispute in the next chapter), should be analysed as instances of lexical narrowing, as each of the alleged modulations (i.e. *retail bank* and *central bank from a G20 country*) asymmetrically entails (or is logically stronger than) the unmodulated or root meaning, namely, *bank institution*. Conversely, hyperbole/exaggeration and loose talk are, in this framework, analysed as instances of lexical broadening; this is because, in these cases, the root meaning appears to undergo lexical weakening (consider, for example, an utterance of 'the water in the bath is *boiling*', which would typically be interpreted as meaning 'the water in the bath is *very hot*').[5]

---

[3] In Chapter II, § 6, it will be argued that (2) and (3)'s perceived interpretations come about via (phrasal) domain restriction, and not via lexical modulation.

[4] A sentence may have more than one literal meaning as a result of containing an ambiguous word or because more than one parsing is available (among other possibilities). As is customary in linguistics, I will refer to the possible literal meanings that a sentence can have as *readings*.

[5] I will reserve the term modulation for cases in which the (allegedly) modulated sense is either logically stronger or weaker than the root/unmodulated meaning. It should be noticed, however, that contextualists also talk about modulation in cases in which the resulting meaning does not stand in an entailment relation with the root/unmodulated meaning (for example, cases

What sort of pragmatic considerations are expected to drive the LM mechanism? In principle, the same pragmatic considerations that drive the processes of word disambiguation and indexical saturation.[6] There are *at least* two pragmatic considerations that are going to matter; in broadly Gricean terms, there is, on the one hand, *Quality* (the presumption that the speaker does not say things which s/he believes to be false) and, on the other hand, *Relevance* (the presumption that the speaker says things which are related to what is being talked about). The term 'bank', for example, can be disambiguated either through Quality, as illustrated in (4), or through Relevance, as in (5) and (6).

(4) Clark went to the *bank* to get some money.

(5) Context: There have been claims that a dangerous criminal is wandering about by the river bank. Inspector Clark is asked to go to the river bank and check.

Clark went to the *bank*.

(6) Context: There have been claims that a dangerous criminal is hidden somewhere inside the HSBC branch of Regent Street, London. Inspector Clark is asked to go to the said HSBC branch and check.

Clark went to the *bank*.

In (4), in the absence of any context, 'bank' will clearly be disambiguated in favour of its *bank institution* sense. In principle, 'bank' could have been interpreted in its *river bank* sense; such a move, however, would have generated a proposition which, given world knowledge, is implausible. Indeed, the presumption that the speaker does not say things that s/he believes to be false is much easier to maintain (i.e. it requires less 'blind' belief) if the proposition expressed is likely to be true, as opposed to unlikely to be true. Quality-related considerations, therefore, are expected to conspire so that 'bank' is interpreted in its *bank institution* sense.

In (5) and (6), on the contrary, Quality considerations do not seem to play any prominent role in disambiguating 'bank': the two possible readings of 'Clark went to the bank' are, *a priori*, equally plausible, and it is not clear why an interpreter should favour one over the other. However, (5) and (6), as a result of being uttered in different contexts, have different expectations of Relevance: these expectations alone seem enough to resolve the ambiguity of 'bank' in each case. It is worth noting that, in real-world interactions, Quality and Relevance are typically intertwined; for example, (4), if uttered

involving metaphor or metonymy)—see, for example, Carston (2002), Wilson and Carston (2007), and Wilson (2011). In this thesis, I will have nothing to say about such cases.

[6] Recanati (2004) calls the pragmatic processes that contribute to determining the literal meaning of sentences *primary pragmatic processes*, and those that are responsible for the derivation of conversational implicatures *secondary pragmatic processes*. LM, lexical disambiguation and indexical saturation all belong to the first category.

in a conversation, is likely not to be uttered out of the blue but in relation to some salient issue that is being discussed, one that, in one way or another, makes banks of the financial sort relevant.

The LM proposal has had a great deal of influence on recent approaches to lexical interpretation. In philosophy, Ludlow's *Living Words: Meaning Underdetermination and the Dynamic Lexicon* (2014) echoes many of the ideas pioneered in Sperber and Wilson (1986). In semantics, Pagin and Pelletier (2007), and more recently, Del Pinal (2018) incorporate a modulation function in the compositional semantics, which permits lexical denotations to modulate (either narrow or broaden) in response to contextual demands. Finally, the LM thesis has influenced probabilistic models of communication. For example, the RSA (Rational Speech Act) model (Bergen 2016; Bergen, Levy, and Goodman 2016) implements what is known as *lexical uncertainty*, a technique which has the net effect of allowing pragmatic considerations to strengthen lexical meanings during the process of semantic composition.

Bergen's proposal has some distinctive features that should not pass unnoticed. In this framework, and simplifying for the sake of exposition, there is no fixed lexicon but rather a set of lexica, containing every possible lexicon (i.e. every possible mapping from form to meaning). In lexicon α, 'bank' may mean *central bank*, in lexicon β 'bank' may mean *investment bank*, in lexicon γ 'bank' may mean *retail bank*, etc. A listener has uncertainty about which lexicon the speaker is using. To interpret an utterance, a listener considers how likely the speaker would have been to choose the uttered sentence given each lexicon and, by so doing, accounts for her uncertainty (for details, see Bergen, Levy, and Goodman 2016: 25). One consequence of this account is that, once a lexicon is chosen to interpret a given utterance or discourse fragment, every occurrence of each word is bound to mean what it means in the chosen lexicon; for example, on this account, if an occurrence of 'bank' in discourse fragment F means *retail bank* because lexicon γ was chosen, then any other occurrence of 'bank' in F is bound to also be interpreted as meaning *retail bank*.

Understandably, LM and associated proposals have been met with scepticism; as Rothschild (2015) argues, if lexical items are indeed that context-sensitive, then why is it that 'believe' in (7)a cannot be strengthened to *be certain/sure*? Indeed, such a modulation would turn (7)a, a contradiction, into a pragmatically acceptable statement, as illustrated in (7)b.

(7) a. # Tim think it's raining, but he doesn't believe it.
b. Tim thinks it's raining, but he isn't sure about it.

Rothschild (2015) concludes that there must be restrictions on lexical modulation, otherwise a whole range of non-existent readings would be suddenly predicted to be possible.

2 SUMMARY OF CORE CHAPTERS

2.1 Chapter II: On the Nature of Lexical Narrowing Effects

In Chapter II, I will be concerned with testing whether there is such a thing as lexical narrowing: that is, whether pragmatic considerations can in fact strengthen the denotation of content words during the process of literal meaning composition. To do this, I will examine many of the linguistic examples that have been used to motivate the lexical narrowing hypothesis and will try to determine, through the application of linguistic tests, whether these examples can in fact be taken to suggest that lexical narrowing is operative in natural language. The conclusion that I will reach goes against what many might expect: I have found no convincing data indicating that content words support lexically narrowing; if such data exists, it is yet to be found.

2.2 Chapter III: The Character of Imprecision

Chapter III has two main components. The first part will be concerned with testing whether there is such a thing as lexical broadening (the logical counterpart of lexical narrowing). Here, too, the verdict will be negative: I have found no evidence indicating that the meaning of individual content words can be logically weakened during the process of literal meaning composition. The second part of this chapter will be concerned with providing a characterisation of the phenomenon of (im)precise interpretation (which, in later chapters, will be referred to with the more technical term (*non-*)*maximality*); for example, a sentence such as ‘the bookcase is empty’ can, in certain contexts, be interpreted precisely, as meaning ‘the bookcase is completely empty’ (false if there is one book in the bookcase) but, in other contexts, can be interpreted imprecisely (true enough if there is one or two books in the bookcase). Within the lexical modulation framework, imprecision has been analysed as an instance of lexical broadening (e.g. the word ‘empty’, in certain contexts, can be weakened to ‘more or less empty’). I will argue that this analysis is incorrect and argue that imprecise interpretation is a global, post-compositional mechanism that operates on whole sentences.

2.3 Chapter IV: (Non-)Maximality Revisited

In Chapter IV, I will do three things. First, I will discuss the semantic condition known as homogeneity; a homogenous sentence is a sentence that has an extension gap of a very specific sort (a gap that, as will be argued in Chapter VI, does not have to be visible). Second, I will introduce Križ’s (2015) theory of (non-)maximality (or (im)precision); according to Križ’s (2015) account, (non-)maximality has two legs: a semantic leg (homogeneity) and a pragmatic leg (an important modification of the maxim of

Quality). It will be shown how Križ's account makes sense of the complex data associated with imprecision phenomena. To conclude, I will point out some limitations of his account.

2.4 Chapter V: Probabilistic (Non-)Maximality

Chapter V has three components: (i) a probabilistic account of relevance; (ii) a theory of (non-)maximality that builds upon Križ's (2015) main theoretical insights; (iii) a comparison between Križ's (2015) account and the proposed account (which I will refer to as probabilistic (non-)maximality or PNM). In (i), I will briefly summarise the classical conception of relevance associated with the QUD framework (Roberts [1996] 2012), and propose a revision of this conception along the lines suggested in Büring (2003). In (ii), I will put forward an account of (non-)maximality that, like Križ's (2015), has two components: a semantic one (homogeneity) and a non-semantic one (an account of how the information carried by trivalent sentences is assimilated into the common ground). In (iii), I will compare the two accounts and show that PNM can make sense of data that Križ's (2015) account leaves unexplained.

2.5 Chapter VI: Invisible Gaps

In Chapter VI, the case is made that homogeneity is not always visible. It should be noted that Križ (2015) never claimed that homogeneity had to be visible and, in fact, his attempt to derive homogeneity in numerals (see Križ 2015: 105) indicates that, at the time of writing his dissertation, he was very much aware that homogeneity did not need to entail gap visibility.

The chapter has the following structure. First, I will put forward a homogenous semantics for absolute adjectives, a semantics that, if coupled with the appropriate alignment condition, predicts gap invisibility. Second, I will call into question the standard explanation of the distribution facts of *slightly* and *completely*, an explanation which relies on what I believe to be incorrect assumptions about the structure of adjectival scales. Last but not least, I will present an analysis in which *completely* functions as the homogeneity remover for total adjectives, whereas *slightly* functions as the homogeneity remover for partial adjectives. This analysis is able to predict the distribution of these modifiers with greater accuracy than the scale structure approach (or so I will argue).

Chapter II
ON THE NATURE OF LEXICAL NARROWING EFFECTS

1 OVERVIEW

In the present chapter, I will be concerned with testing whether there is such a thing as *lexical narrowing* (LN henceforth): that is, whether pragmatic considerations can have the effect of strengthening word meanings during the process of literal meaning composition (see Chapter I, § 1). Scrutinising every alleged case of LN is an impossible task; what is possible to do, however, is to select the most compelling cases, and then proceed to examine whether they live up to the hype. This is the strategy that I will pursue here.

2 BACKGROUND ASSUMPTIONS

Searle (1980, 1992) famously claimed that the word 'cut' makes different contributions to the truth-conditions of sentences such as (1)a and (1)b.

(1) a. Bill cut the grass.
b. Sally cut the cake.

In *The Rediscovery of Mind* (1992), for example, Searle writes: 'the utterances contain the literal occurrence of the verb "cut", but that word, on a normal interpretation, is interpreted differently in each sentence' (p. 179). Each sentence, according to Searle, appeals to different 'background assumptions' (his terminology), and these assumptions (roughly, bundles of world knowledge that dictate that cakes are normally cut with a knife, and that the grass is normally cut with a lawn mower) end up having an effect on how the word 'cut' is interpreted (1992: 179-180). In LN terms, this is equivalent to saying that, in response to pragmatic pressures (in this case, Quality-related considerations), 'cut' lexically narrows to *cut with a lawn mower* in (1)a and to *cut with a knife* in (1)b. Recanati (2002, 2009) endorses Searle's analysis, whereas Cohen (1986) and Carston (2002) make almost identical claims to those of Searle in connection with verbs such as 'drop' and 'open'.[7]

It is true that the 'intuitive' truth-conditions of (1)a and (1)b appear to encode something like a manner of cutting: indeed, it is not immediately clear whether an utterance of (1)a would be deemed to be true if Bill had cut the grass using scissors, or whether an utterance of (1)b would be accepted as true if Sally had cut the cake with a chainsaw. But the intuitive truth-conditions of an utterance are not necessarily

[7] It is worth noting that these claims been met with scepticism by many (e.g. Pelletier 2013; Unnsteinsson 2014). However, a data-focused examination of these claims has been lacking in the literature. This section is aimed at filling this gap.

a faithful guide to its literal meaning. If it is unclear what the literal meaning of an utterance is, then linguistic tests need to be applied to shed light on the matter.

One way to determine whether ‘cut’ supports LN is to come up with a sentence which, unless ‘cut’ were to modulate to a stronger meaning, would give rise to a contradiction (thereby violating both Quality and Relevance). This technique, to my knowledge, originates in Rothschild (2015), and is implemented in examples below (2) and (3).

(2) a. # John didn’t cut the cake, but he cut it with a chainsaw.
b. John didn’t cut (with a knife) the cake, but he cut it with a chainsaw

(3) a. # John didn’t cut the grass, but he cut it with scissors.
b. John didn’t cut (with a lawn mower) the grass, but he cut it with scissors.

The data speaks for itself. If the first occurrence of ‘cut’ was able to modulate to *cut with a knife* in (2)a or to *cut with a lawn mower* in (3)a, then one would expect this mechanism to take place (i.e. such modulations would turn these examples into non-contradictory sentences). The fact that (2)a and (3)a cannot be interpreted as meaning (2)b and (3)b, respectively, suggests that ‘cut’, *contra* Searle and others, does not support LN.

To account for the contradictoriness of (2)a and (3)a, contextualists could put forward the following principle: once a word ‘blah’ in some fragment of discourse F is interpreted as meaning $blah_1$, then all other occurrences of ‘blah’ in F should be interpreted as meaning $blah_1$ (Recanati p.c.). This constraint (which I will refer to as *uniformity*), is built in Bergen, Levy, and Goodman’s (2016) *lexical uncertainty*, as discussed in Chapter I, § 1. If such a constraint were to be assumed, and the first occurrence of the word ‘cut’ in (2)a were to be interpreted as meaning *cut with a knife*, then the second occurrence of ‘cut’ would too need to be interpreted in this way, which would create a contradiction. The same outcome would hold if ‘cut’ were to be interpreted uniformly as meaning *cut simpliciter*.

However, on closer examination, *uniformity* falls short of making sense of the data. Notice that, according to this constraint, ‘cut’ should always be interpreted uniformly, including in (4) below. But, if so, why is it that the first conjunct of (4) has the same intuitive truth-conditions as (1)a, and that the second conjunct of (4) has the same intuitive truth-conditions as (1)b?

(4) Bill cut the grass and Sally cut the cake.

Indeed, if the intuitive truth-conditions of (1)a and (1)b are the consequence of the word ‘cut’ being lexically narrowed and, further, there is such a thing as *uniformity*, then the intuitive truth-conditions of

(4) cannot be explained by an appeal to LN—and it would be completely ad-hoc to stipulate that the intuitive truth-conditions of (4) and those of (1)a and (1)b come about via different mechanisms.

In short, the LN hypothesis, with or without *uniformity*, fails to make sense of the data. If *uniformity* is assumed, the contradictoriness of (1)a and (3)a is accounted for, but the intuitive truth-conditions of (4) are unexpected; alternatively, if *uniformity* is not assumed, the intuitive truth-conditions of (4) can be explained by an appeal to LN, but the contradictoriness of (1)a and (3)a is unexpected. Thus, it seems reasonable to conclude that, *pace* Searle and other contextualists, that LN is not responsible for delivering the intuitive truth-conditions of (1)a and (1)b.

Another way to put Searle's claims to test is to ask whether the alleged readings of (1)a and (1)b can be openly challenged. Indeed, if (1)b, for example, were to mean something close to *Sally cut the cake with a knife* by virtue of the word 'cut' being interpreted as meaning something close to *cut with a knife*, then the audience should, at least in principle, be able to challenge (1)b by producing an utterance of the form 'That's not true! She cut it with a chainsaw'.[8] This is tested in (5) and (6) below.

(5) A: Sally cut the cake.
B: # That's not true! She cut it with a chainsaw.

(6) A: Sally cut the cake with a knife.
B: That's not true! She cut it with a chainsaw.

The data, once again, are clear: (5)B is completely infelicitous, whereas (6)B is perfectly felicitous. This indicates that the word 'cut' in (5)A is not interpreted as meaning *cut with a knife*: if it was, then the exchange in (5) should be as natural (or near as natural) as the one in (6).

It could be argued that LN, by virtue of being an implicit mechanism, does not quite have the same effect on (5)A as actually adding 'with a knife' explicitly, as in (6)A. In this regard, a case could be made that implicit content has a somewhat different status from explicitly provided material, and that this difference manifests itself in the extent to which a direct challenge is admissible (e.g. it may be that implicit content cannot be directly challenged).[9] The minimal pairs reported in (7) control for this possible confounding factor.

[8] This sort of test, known in the literature as the 'That's not true!' test (e.g. Tonhauser 2012) is typically used to diagnose for 'at-issueness', a technical notion which is close (but not identical) to Grice's (1975) notion of 'what is said'.
[9] Thanks to Benjamin Spector for pressing on this point.

(7) Context: Prof. Smith is talking about the students newly admitted to the ENS…

a_1. Prof. Smith: Not every student is excellent.
a_2. Prof. Black: That's not true—all the students newly admitted to the ENS are excellent.

b_1. Prof. Smith: Not every student newly admitted to the ENS is excellent.
b_2. Prof. Black: That's not true—all the students newly admitted to the ENS are excellent.

c_1. Prof. Smith: Not every student (in the world) is excellent.
c_2. Prof. Black: # That's not true—all the students newly admitted to the ENS are excellent.

(7)a_1 is a case of implicit domain restriction, a mechanism of local pragmatic enrichment that has the net effect of strengthening the denotation of the NP restrictor (cf. § 6 for an extensive discussion of this phenomenon). Due to implicit domain restriction, the asserted content of (7)a_1 ends up being more or less equivalent to that of (7)b_1: the fact that (7)a_1 and (7)b_1 can both be challenged by exactly the same sentence is an indication of this; furthermore, notice that (7)c_1, in which domain restriction is explicitly blocked by the addition of '(in the world)', cannot be challenged by this sentence.

If domain restriction, which is an implicit mechanism, has the net effect of making (7)a_1 and (7)b_1 informationally equivalent (at least to the extent that they can both be challenged by the same sentence), then why is it that the (alleged) strengthening from *cut* to *cut with a knife*, which is as implicit as domain restriction, fails to make (5)A and (6)A equivalent in the same sort of way? Indeed, if one thinks that the implicitness of LN is responsible for the contrast between (5) and (6), then one is going to have a hard time explaining why it is that both (7)$a_{1\text{-}2}$ and (7)$b_{1\text{-}2}$ are perfectly natural exchanges.

The results obtained in Rothschild's contradiction test and in the 'That's not true!' test show that (1)a entails nothing about the manner in which the grass was cut and that (1)b entails nothing about the manner in which the cake was cut. The thesis, held by Searle and others, that 'cut' is interpreted as meaning something like *cut with a lawn mower* in (1)a and something like *cut with a knife* in (1)b predicts the existence of these entailments: Searle's thesis, therefore, must be wrong.

The literal meaning of (1)a and that of (1)b, undoubtedly, interact with what interpreters know about the world (i.e. their 'background assumptions'), but this interaction is much subtler than the one envisaged by contextualists. Let's take (1)b, for example. In light of world knowledge, the prior probability of a cake being cut with a knife is much higher than the prior probability of a cake being cut with something other than a knife; upon learning that some cake was cut, one is bound to infer that it is much more likely than not that the cake was cut with a knife.[10] Thus, though 'Sally cut the cake' is true

[10] In more technical terms:
P(the cake was cut with a knife | the cake was cut) >> P(the cake was cut with something other than a knife | the cake was cut)

if Sally cut the cake with a chainsaw, 'Sally cut the cake', if seen through the lens of world knowledge, is much less likely to be true *as a result of* Sally cutting the cake with a chainsaw than *as a result of* Sally cutting the cake with a knife.

This sort of interaction between world knowledge and assertion, it is worth noting, can be observed *everywhere* (and not just in constructions involving the verb 'cut'). Imagine, for example, that someone from Austin (Texas, US) utters 'I had breakfast'; further, imagine that someone from Paris (France) utters the same sentence. Given world knowledge, one can infer that the person from Austin and the person from Paris are likely to have had different things for breakfast; however, it is a mistake to think that world knowledge considerations modulate the meaning of 'breakfast' (or strengthen the truth-conditions of 'I had breakfast' through some other compositional procedure), in the same way that it is a mistake to think that world knowledge considerations modulate the meaning of 'cut' in (1)a and (1)b.

3 LOCAL EXHAUSTIFICATION

In (8), the verb *eat* appears to be strengthened to *eat but not devour*: if no strengthening was at work, (8) would be a contradiction (however, it does not feel like one).

(8) Anyone who eats their food will get a prize, but those who devour it won't.[11]
⇒ Anyone who eats (but doesn't devour) their food will get a prize, but those who devour it won't.

Is this an instance of LN? I do not think it is; the strengthening procedure at work in (8) is of a very specific type: it appears to involve combining *eats* with the negation of *devour*. If LN were to be responsible for such a procedure, it should be possible to find a context in which *eats* gets to mean something else other than *eats but doesn't devour*; however, in (8), *eats* can only be interpreted as *eats but doesn't devour*, irrespective of context.[12]

The strengthening observed in (8) is, as far as I can tell, the result of a scalar implicature being computed in the first conjunct. Scalar implicatures are instances of silent exhaustification,[13] a (plausibly grammaticalised) mechanism whereby a given syntactic constituent is combined via conjunction with the negation of each of its logically stronger alternatives. In (8), the scalar implicature appears to be triggered due to lexical competition between *eat* and *devour*, competition which is 'forced' by the contrastive construction (notice that strong prosodic focus on 'devour' is required).[14] The computation

[11] Thanks to Leon Bergen for the example.
[12] Thanks to Dan Hoek for pressing on this point.
[13] As opposed to *overt* exhaustification, executed by means of an overt particle such as 'only'.
[14] The competition between 'eat' and 'devour' is clearly not a feature of the English lexicon (as opposed to the competition between 'some' and 'all', for example): indeed, an utterance of 'John ate the cake' does not give rise to the implicature 'John

of this implicature, it is worth noting, must be taking place in an embedded position: (8), if read logically, is a contradiction,[15] and no scalar reasoning can be applied to a contradiction.[16]

The fact that scalar implicatures can be local calls into question the received, 'pragmatic' view of implicature calculation, which has it that implicatures are the output of global, domain general processes of rational thought (e.g. Grice 1975, Horn 1972, and many others).[17] It remains controversial how to reformulate the pragmatic account to make room for local (or embedded) implicatures; for the purposes of this discussion, I shall subscribe to the grammatical approach, as laid out in (Chierchia, Fox, and Spector 2008).[18] On this account, scalar implicatures are the consequence of a grammatical mechanism, implemented through a syntactically realised operator *exh*, which can be thought of as a silent variant of 'only'.

To generate the attested reading of (8), *exh* can be inserted, for example, right above the embedded TP 't_1 eats their food', as illustrated in (9).

(9) [TP [anyone [CP who$_1$ ***exh***[TP t_1 eats their food]]] [will get a prize]], but those who devour it won't.

The *exh* operator, as defined in Chierchia, Fox, and Spector (2008), takes the literal meaning of the relevant constituent and adds to it the negations of all the alternatives which are logically stronger than the literal meaning. Under the assumption that *eats* and *devour* are in competition with each other, the alternative 't_1 <u>devours</u> their food' can be generated (by replacing 'eats' with 'devour' in the embedded TP). This alternative is stronger than the literal meaning of the TP; as a result, local exhaustification is predicted to take place: 't_1 eats their food' should be strengthened to 't_1 eats and doesn't devour their food' (see Chierchia, Fox, and Spector (2008) for a detailed presentation of the algorithm).

Here I am of course stipulating that the exhaustification procedure is carried out at the level of the TP; it could also be stipulated that *exh* directly applies to *eats*: indeed, if the latter analysis were to be pursued, (8)'s attested reading would also be generated. I know of no good reason to prefer one analysis over the other; however, there are good reasons to think that the mechanism of local exhaustification

didn't devour the cake'. If anything, 'John ate the cake', in combination with world knowledge, affords one the inference that John is likely to have eaten the cake in the manner in which people stereotypically eat cakes.

[15] Indeed, 'anyone who eats their food will get a prize' strictly entails 'anyone who devours their food will get a prize'.

[16] As mentioned, exhaustification is a mechanism that involves negating those alternatives that are logically stronger than the literal meaning (and, subsequently, adding them to the literal meaning via conjunction); however, if the literal meaning is a contradiction, none of its alternatives will asymmetrically entail it and, as a result, there will be nothing to exhaustify it with.

[17] Cohen (1971) with examples such as (i), was, to my knowledge, the first to show that Gricean reasoning could be embedded.

(i) Bob believes that Anna ate some of the cookies.
$\Rightarrow$ Bob believes that Anna ate some *but not all* of the cookies.

[18] See Sauerland (2014) for a comparison between the so-called 'pragmatic' approach vs. Chierchia, Fox, and Spector (2008).

cannot be operationally subsumed into lexical exhaustification; consider, for example, the following example, due to Benjamin Spector (p.c.):

(10) a. John [has to solve some of the problems to pass], while other students have to solve all.
⇒ b. John has to solve some of the problems to pass but he does not have to solve all the problems to pass, while other students have to solve all.

The fact that (10)a can have (10)b as a reading shows that *exh* can target phrases: indeed, to generate (10)b, the whole VP of the first clause must be exhaustified. Hence, even if *exh* were to target lexical constituents in some cases, it would be erroneous to characterise local exhaustification as a lexical mechanism. This is another reason to reject (8) as evidence of LN: not only the strengthening observed in (8) is blind to contextual pressures: the mechanism responsible for performing such strengthening is not lexical in nature.

4 ACCEPTABLE CONTRADICTIONS

By 'acceptable contradictions', I mean natural language sentences which, although they look as formal contradictions, appear to have non-contradictory readings or interpretations (cf. Alxatib, Pagin, and Sauerland 2013). I have identified three 'types' of these contradictions:

(11) Acceptable contradictions:

a. *but*-contradictions	John is tall, but he isn't TALL.	[Capitalisation signals focal stress.]
b. *and*-contradictions	John is tall and John isn't tall.	[Or 'John is and isn't tall.']
c. *neither*-contradiction	Neither is John tall, nor is he not tall.	[Or 'John neither is nor isn't tall.']

To my knowledge, only *but*-contradictions have been used explicitly to argue for LN (see Carston 2002: 324); however, all the sentences in (11) could in principle be brought forward to motivate the LN hypothesis: at least intuitively, in all these cases, it looks as if the word *tall* was being subjected to some sort of adjustment procedure, a procedure whose net effect could be argued to be that of 'rescuing' the sentences from contradictoriness. It is worth pointing out that LN, if constrained by *uniformity* (cf. § 2), could not be invoked as a possible account of acceptable contradictions. LN, of course, does not need to incorporate this constraint; but, if it is not constrained by this principle, one cannot appeal to it to account for the contradictoriness of (1)a-b, or similar examples.

In the next section, I will examine the three types of contradictions listed in (11) and investigate whether LN can be held responsible for their non-contradictory interpretations. The anticipated conclusion is this: none of the sentences in (11) can be taken to be evidence of LN; furthermore, each of these sentences is 'rescued' from contradictoriness by a different mechanism.

4.1 But-contradictions: intensification of gradable adjectives via stress placement[19]

The phenomenon observed in (11)a is stress-dependant: in particular, stress on 'tall' appears to strengthen the meaning of the gradable adjective: from *tall* to *very tall*. The upshot of this is that (11)a is interpreted as in (12).

(12) John is tall, but he isn't *very* tall.

This phenomenon, discussed in Bergen (2016) in quite some detail, is independent of whether the sentence in which the stressed constituent appears is a contradiction; indeed, in (13) below, the same effect as in (11)a is observed and (13) is not a contradiction:

(13) John is TALL.
⇒ John is very tall.

These data could hardly be taken as evidence of LN. For one thing, context is playing no role in triggering the observed strengthening: the intensification of the adjective, quite clearly, is being triggered by the placement of stress. In addition, this kind of intensification-by-stress mechanism only works with threshold-sensitive gradable adjectives (cf. Chapter VI, § 2), as the non-rescuable contradictoriness of (14)a and (14)b exposes.

(14) a. # He is dead, but he isn't DEAD.
b. # He is vegan, but he isn't VEGAN.

It should also be noticed that the phenomenon observed in (11)a is of a very different nature from (11)b-c. First, no constituent in (11)b-c needs to be stressed for the sentences to be rescued from contradictoriness. Second, whereas (11)b-c admit right-node-raised paraphrases (as indicated between square brackets next to the examples), (11)a does not—as shown in (15). Finally, (11)a, unlike (11)b-c, requires *but*-conjunction (to mark contrast), as revealed by the oddness of (16).

(15) # John is, but he isn't TALL.

(16) # John is tall and he isn't TALL.

[19] I thank Manuel Križ for an especially helpful conversation on the topic of *but*-contradictions. I also thank Philippe Schlenker for an e-mail exchange on the matter.

To conclude, on the basis of *but*-contradictions, it cannot be argued that the meaning of content words can be strengthened more or less freely to meet contextual demands. If anything, what these data show is that a sub-group of lexical items, irrespective of context, can be subjected to a very specific type of strengthening (i.e. strengthening along a degree scale: from *x* to *very x*), if (and only if) stress is placed on them.

4.2 And- and neither-contradictions: conjunctions versus negated disjunctions

And- and *neither*-sentences, in their non-contradictory interpretations, appear to mean different things. *And*-sentences, at least intuitively, give rise to what I will call a two-angle interpretation, i.e., **there's a sense in which** X is *F* & **there's a sense in which** X isn't *F* (cf. Egré 2019). *Neither*-sentences, by contrast, seem to be interpreted as meaning *x is neither* **definitely** *F, nor* **definitely** *not F.*

(17) *and*-contradiction
a. x is *F* and x isn't *F*.
b. John is tall and John isn't tall. [Or 'John is and isn't tall.']
c. ⇒ There's a sense in which John is tall and there's a sense in which he isn't.

(18) *neither*-contradiction
a. Neither is x *F* nor is x not *F*.
b. Neither is John tall, nor is he not tall. [Or 'John neither is nor isn't tall.']
c. ⇒ Neither is John definitely tall, nor is he definitely not tall.

Let's first focus on (18): on the assumption that (18)b is indeed interpreted as in (18)c, then it is hard resist the following account: there is something like a vagueness-trimming operation that applies to both 'tall' and 'not tall'; this operation has effects comparable to those of an overt *definitely*: it sharpens, at least to some extent, the meaning these constituents (via removal of borderline cases; cf. Chapter III, § 3.3); as a result of this operation, (18)b is rescued from contradictoriness: (18)b ends up meaning that John is a borderline case of tall.

If the picture just sketched is more or less accurate, one expects *neither*-sentences (i.e. neither is x *F* nor is x not *F*) to give rise to a non-contradictory interpretation only if *F* is a vague term (i.e. non-vague terms are not expected to be sensitive to a vagueness-trimming operation). By contrast, *and*-sentences (i.e. x is *F* and x isn't *F*) should, at least in principle, give rise to a two-angle (non-contradictory) interpretation irrespective of whether *F* is vague. Though it is not clear what *there is a sense in which* does to meaning that it applies to, one thing is clear: there is no reason to expect its application to have effects comparable to those of an overt *definitely*.

To test these predictions, I shall use the non-vague comparative predicate 'taller than x'. These predicates are notorious for not giving rise to the sorties paradox (a landmark test for vagueness)—cf.

Kennedy (2007, 2011); for example, in (19), one is not at all compelled to accept the inductive premise and, as a result, no paradox can be generated.

(19) Rodman is taller than Jordan.

Inductive premise:
For any value *n* (Rodman's height), if Rodman is taller than Jordan at *n*, then he is taller than Jordan at *n* – 1cm.

Let's now imagine the following context. The couch of the Chicago Bulls is vacillating as to whether give the power forward position to Rodman or Jordan. He finally makes his mind and chooses Rodman for the position, because Rodman is taller than Jordan. One of his assistants objects to this decision, and he does so by saying…

(20) Think twice boss. Rodman is and isn't taller than Jordan.

Let's assume that it is common ground that Rodman is taller than Jordan, but it is also common ground that Jordan can jump much higher than Rodman and that the power forward is expected to be able to jump very high. In this context, (20) appears to mean *there is a sense in which Rodman is taller than Jordan* (i.e. Rodman is as a matter of fact taller than Jordan) *and there is a sense in which he isn't* (i.e. Rodman cannot jump as high as Jordan).

(21), by contrast, is a contradiction, and cannot be rescued neither in the stipulated context nor in any other context.

(21) # Neither is Rodman taller than Jordan nor is he not taller than Jordan.

The sharp contrast between (20) and (21) confirms that *and*- and *neither*-sentences, when assigned a non-contradictory interpretation, give rise to different interpretations. *Neither*-sentences are interpreted as descriptions of borderline cases (when *F* is vague); when *F* is not vague, these sentences are contradictions. *And*-sentences, on the other hand, give rise to a much more complex and nuanced interpretation, a 'two-angle' interpretation, which is not correlated with the presence of a vague predicate.

4.2.1 LN and and-contradictions

LN cannot be the mechanism responsible for rescuing *and*-sentences, the reason for this is simple: LN cannot generate two-angle interpretations; as illustrated in (22) and (23) below, *there's a sense in which x is F* does not entail *x is F*.

(22) a. John was upset and John wasn't upset.
b. There's a sense in which John was upset and there's a sense in which John wasn't upset.
> **There's a sense in which John was upset ⇏ John was upset**
> **John was mildly upset ⇒ John was upset** [control]

(23) a. John is and socialist and John isn't a socialist.
b. There's a sense in which he is a socialist and there's a sense in which he is not.
> **There's a sense in which John is a socialist ⇏ John is a socialist**
> **John is a devoted socialist ⇒ John is a socialist** [control]

Another argument against invoking LN to account for two-angle interpretations is the following: if one could LN-adjust the meaning of individual words to avoid contradictoriness (and hence keep afloat the presumption that the speaker follows both Quality and Relevance), then (24), for example, should be able to be interpreted in a non-contradictory manner:

(24) # When I'm upset, I cry, but when I'm upset, I don't cry.

However, I cannot think of *any* context in which (24) would not be perceived as a contradiction: however, if the first occurrence of *upset* could be LN-adjusted to *terribly upset* and the second to *only mildly upset*, then (24) should have a non-contradictory interpretation, as illustrated in (25).

(25) When I'm terribly upset, I cry, but when I'm only mildly upset, I don't cry.

Of course, (24) could be rescued via stress placement (see § 4.1.), as shown in (26).

(26) When I'm upset, I cry, but when I'm UPSET, I don't cry.
⇒ When I'm upset, I cry, but when I'm *very* upset, I don't cry.

The fact that (24) can be rescued via the placement of stress, however, is orthogonal to the point that I am making. *And*-sentences do not require having stressed constituents in order to escape contradictoriness; if LN were to be responsible for rescuing these kind of sentences, then (24) should also be salvageable by the application of this mechanism: the fact that (24) can *only* be rescued via 'intensification by stress' suggests that, whatever mechanism is involved in rescuing *and*-sentences, it is not LN.

In § 4.3.2, I will discuss some minimal pairs that show that *and*-sentences, when embedded, fail to trigger two-angle interpretations: such data, as far as I can tell, is fatal for any account that has it that that two-angle interpretations are derived via LN or any other 'local' mechanism (see, for example, Egré and Zehr 2018).

4.2.2 LN and *neither*-contradictions

Can *neither*-sentences be put forward as evidence of LN (or lexical modulation more broadly)? I do not think so. First, these sentences are interpreted as meaning *x is neither* **definitely** *F, nor is* **definitely** *not F* (if *F* is a vague predicate), *irrespective* of context: in fact, this is the only possible reading of these sentences as far as I can tell. What these data suggest is that there must me something like a vagueness-trimming operation that interacts with the recursive semantics: however, I see no evidence here for the thesis that the meaning of words can be adjusted more or less freely at the compositional stage in response to contextual pressures.[20] In addition, note that the sharpening operation involved in the interpretation of *neither*-sentences appears to target *F* (which may be a lexical constituent) but also *not F* (which is not a lexical constituent). Hence, the mechanism that rescues *neither*-sentences, on the assumption that (18)b is in fact interpreted as (18)c, cannot be adequately characterised as being lexical in nature.

4.3 How do acceptable contradictions come about?

4.3.1 Neither-contradictions

Neither-sentences are compatible with at least two (closely related) accounts. The most intuitive one consist of having, as shown in (27), a silent *definitely* operator that more or less mirrors (syntactically and semantically) an overt *definitely* (perhaps with some distributional constraints). A solution in this spirit is provided in Egré and Zehr (2018). An alternative proposal, as illustrated in (28), is to posit the existence of a clause-level operator ***S*** that has the effect of collapsing the extension gap (the worlds that neither verify nor falsify the clause) into falsity; ***S***, on this account, would perform a task analogous to that of a local accommodation operator (e.g. Heim 1983; Beaver and Krahmer 2001).[21] This solution is hinted at in Spector (2012, 2016).

[20] Notice, furthermore, that *definitely p* may or may not be logically stronger than *p*: this will depend on how vagueness is formalised and the precise notion of logical entailment that one adopts.

[21] Take, for example, (ii):

(ii) Mary hasn't met the king of France, because there is no king of France. (Spector 2016b)

The definite description 'the king of France' triggers the presupposition that there is a king of France; if this presupposition were to be accommodated globally, that is, if one were to update the common ground with this presupposition, then (ii) would yield a contradiction. (ii), however, does not feel contradictory.

(27) Bakewell neither is (*definitely*) a city, nor it is (*definitely*) not a city.

(28) NOT (***S***(Bakewell is a city) or ***S***(Bakewell is not a city)).

Questions remain, as always. Do these accounts make different predictions or are they, at bottom, the same account? Is the gap-removal operation some sort of repair mechanism made available by the grammar in exceptional circumstances, or is it applied in a more systematic fashion, and not just to rescue otherwise contradictory sentences? Further research is needed to make progress on these matters.

4.3.2 And-contradictions

As already discussed, *neither*-sentences (of the form 'neither is x *F* nor is x not *F*') require *F* to be a vague constituent in order to be interpreted as descriptions of borderline cases; by contrast, *and*-sentences (of the form 'x is *F* and x isn't *F*') do not require *F* to be vague in order to be interpreted in a 'two-angle' manner. There are further (and subtler) differences between these sentences. First, *and*-sentences, even if assigned a two-angle interpretation, always remain, to some extent, contradictory. Contrast (29) with (30) below, for example:

(29) John neither is nor isn't rich.

(30) John is and isn't rich.

(29) does not feel *at all* like a contradiction, whereas (30), at least according to my intuitions, has more of a mixed status: if the pun is tolerable, one may say that (30) is and isn't a contradiction.

Second, *neither*-sentences, insofar as they trigger non-contradictory interpretations, they do so irrespective of context; by contrast, *and*-sentences, in order to be assigned a non-contradictory interpretation, often need some 'help' from context. Take (20), for example, repeated below as (31).

(31) Rodman is and isn't taller than Jordan.

To handle sentences such as (ii), the standard solution is to define an accommodation operator *A* as follows:

$\boldsymbol{\mathcal{A}}(\varphi) = 1$ if $[\![\varphi]\!] = 1$
$\boldsymbol{\mathcal{A}}\ (\varphi) = 0$ otherwise (i.e. if $[\![\varphi]\!] = 0$ or $[\![\varphi]\!] = \#$)

Embedding $\boldsymbol{\mathcal{A}}$ below negation in the first conjunct of (ii) has the net effect of locally accommodating the presupposition (signaled in **bold** font):
NOT [$\boldsymbol{\mathcal{A}}$ [Mary has met the king of France]] = NOT [**there's a king of France** and Mary met him]

After the insertion of $\boldsymbol{\mathcal{A}}$, (ii) is predicted to mean: *It is not the case that there's a king of France and Mary met him, because there is no king of France.*

If uttered in an out-of-the-blue context, (31) feels like a contradiction; however, if uttered in the context stipulated in (20), (31) lends itself to be interpreted in a two-angle manner.

Third, two-angle interpretations, as far as I am aware, can only be assigned to *matrix* sentences; indeed, as soon as *and*-sentences are embedded, two-angle interpretations do not get off the ground. Compare (32) with (33), for example:

(32) a. Mary neither is nor isn't rich.
b. If **Mary neither is nor isn't rich**, then she must be borderline rich.
c. If Mary neither is definitely rich nor is she definitely not rich, then she must be borderline rich.

(33) a. Mary is and isn't rich.
b. # If **Mary is and isn't rich**, those who are rich across the board will look down on her.
c. If there's a sense in which Mary is rich, and a sense in which she isn't, those who are rich across the board will look down on her.

(32)a, as expected, is interpreted as meaning *Mary neither is definitely rich nor is she definitely not rich*; when embedded in the antecedent of the conditional in (32)b, (32)a continues to mean this, as revealed by the equivalence in meaning between (32)b and (32)c. (33)a, notably, fails to reproduce this pattern. (32)a may be able to be interpreted as meaning *there's a sense in which Mary is rich, and a sense in which she isn't*; however, as soon as it is embedded in (33)b, the two-angle interpretation is not triggered, as revealed by the contrast between (33)b and (33)c.[22] The fact that two-angle interpretations are not embeddable strongly suggest that the mechanism that generates these interpretations does not feed composition.

There is a final (an important) difference: *and*-contradictions, as opposed to *neither*-contradictions, have a strong poetic flavour and, to some extent, betray intellectual sophistication. It is therefore unsurprising that these kind of sentences are found in abundance in literary discourse; ancient Majorca storytellers, for example, began their stories with the line *Aixo era y no era* ('It was and it was not…') (Jakobson 1960): the construction highlights, in a concise and elegant manner, the double-sidedness of narrative, the fact that fiction and reality both overlap and diverge.

In the light of these considerations, it is clear to me that *and*-contradictions, unlike *neither*-contradictions, are not 'rescued' at the compositional stage (cf. the embeddability data); if *and*-contradictions are interpreted the way they are, it is due to (or partly due to) their contradictory nature. I suspect that the *and*-contradiction is a species of *enantiosis* (also known as *discordia concors*), a figure

[22] Salvador Mascarenhas pointed out to me that, in order for (33)b to work, 'Mary is and isn't rich' would need to appear between quotation marks. I agree with this assessment, which provides further evidence that *and*-sentences, once embedded, can no longer be interpreted in a two-angle manner.

of speech by which opposites are juxtaposed in an attempt to emphasise contrast and double-sidedness.[23] Why *and*-contradictions, as opposed to other contradictions, are associated with this trope is a hard question; clearly, there are other features of *and*-contradictions that are responsible for triggering enantiosis (contradictoriness is a necessary but not a sufficient condition). Much remains to be done to have a better grasp of this phenomenon.

4.3.3 The strict-tolerant approach: a note

According to the Strict-Tolerant (ST) approach to vague predication (Cobreros, Egré, Ripley, and van Rooij 2012), the semantic evaluation of natural language sentences rests on the interaction of three notions of truth: the classical notion of truth, a notion of tolerant truth, as well as notion of strict truth. On this account, both 'x is neither *F* nor not *F*' and 'x is *F* and not *F*' are contradictions if evaluated *strictly*; however, if evaluated *tolerantly* (and if *F* is a vague predicate), these sentences are not predicted to be contradictions: their meaning is predicted to be 'x is a borderline case of *F*'.

I cannot discuss this interesting framework at the length it deserves; I just want to point out that the arguments given in the previous sections call into question ST's treatment of acceptable contradictions. First, according to the data discussed, *and*-sentences, when interpreted in a non-contradictory manner, are not interpreted as meaning 'x is a borderline case of *F*'; in fact, *and*-sentences, unlike *neither*-sentences, can be assigned a non-contradictory interpretation independently of whether *F* is a vague constituent.[24] In addition, whereas *neither*-sentences can be embedded (and still be interpreted in a non-contradictory manner), *and*-sentences cannot, which calls into question the idea that these sentences are all 'rescued' via the same mechanism.

5 THE GREEN LEAVES

So-called 'Travis cases' are context-shifting thought experiments due to the philosopher Charles Travis. They tend to involve non-ambiguous utterances that turn out to have different readings when evaluated

[23] The following passage from Paul's second letter to the Corinthians is a classic example of enantiosis: 'By honor and dishonor, by evil report and good report: as deceivers, and yet true; as unknown, and yet well known; as dying, and, behold, we live; as chastened, and not killed; as sorrowful, yet always rejoicing; as poor, yet making many rich; as having nothing, and yet possessing all things.' (2 Corinthians 6:10)

[24] Empirical work has shown, it is true, that participants tend to accept both *and*- and *neither*-sentences as descriptions of borderline cases (e.g. Ripley 2011; Alxatib and Pelletier 2011; Serchuk, Hargreaves, and Zach 2011; Egré, de Gardelle, and Ripley 2013). This, however, is not evidence that *and*- and *neither*-sentences, when interpreted in a non-contradictory manner, mean the same thing. It should also be noticed that, although both *and*- and *neither*-sentences are accepted as description of borderline sentences, *neither*-sentences are to a large extent preferred over *and*-sentences (see, for example, Egré and Zehr 2018). The analysis that I have proposed, unlike the ST analysis, immediately makes sense of this empirical finding.

in different contexts. In this section, I will discuss the 'green leaves' case, which is often cited as a landmark example of lexical context-sensitivity.

Travis imagines someone called Pia uttering the sentence 'the leaves are green' in two different scenarios. In the first scenario, Pia is talking to a photographer who needs some green leaves for a photo shooting. Pia offers her some leaves that are green by virtue of having been painted in green (they are in fact naturally red). In this scenario, the intuition is that Pia said something true. In the second scenario, a botanist comes along searching for green leaves for an experiment; referring to the very same leaves, Pia utters 'the leaves are green'. This time, the intuition is that Pia said something false.[25]

To account for the green leaves case, several lexico-pragmatic theories have been put forward; for example: that *green*, in the botanist situation, is strengthened to *naturally green* via LN (e.g. Del Pinal 2018); or that colour terms such as *green* are ambiguous between (at least) two distinct senses, one paraphrasable as *naturally green* (non-gradable) and the other as *green in looks* (gradable) (Kennedy and McNally 2010). Whether colour terms are context-sensitive in some way or another is a question that I will not attempt to answer here. What is clear to me, however, is that there is absolutely no need to stipulate that colour terms are context-sensitive to explain what is going on in the green leaves case.

I will assume that 'green' is monosemic (as there is no evidence to the contrary; see § 5.1 for discussion on this point). Now, what does *green* mean? Well, for one thing, it does not mean *naturally green*: as shown in (34) and (35), 'x is green' entails nothing about whether x is naturally green or not, and 'x is not green' entails that x is not naturally green.

(34) The leaves are green, but they are not naturally green.

(35) # The leaves are not green, but they are naturally green.

Colour terms, at least in English, behave like individual-level predicates; these are terms that denote a (more or less) stable property of an individual—as opposed to stage-level predicates, which denote a transient one; see, for example, Milsark (1974), Carlson (1977) and Chierchia (1995). The standard test used to diagnose whether a predicate is individual- or stage-level comes from Milsark (1974): stage-level predicates (e.g. *available*, *sick*, *concerned*) can occur in *there-insertion* sentences; individual-level predicates (*smart*, *green*, *altruisitc*) cannot. Consider, for example, the contrast between (36)a and (36)b.

[25] Taken from Travis (1997).

(36) a. There were several seats *available*.
b. # There were several seats *green*.

Thus, insofar as *green* is an individual-level predicate, it must mean something like *dispositionally green-looking* (as opposed to *in a transient state of greenness*, the meaning that one would presumably have to posit if *green* was a stage-level predicate).[26] This meaning is compatible with the fact that something may be green (something may be dispositionally green-looking) yet not look green (because of an intruding shade, for example); it is also compatible with the fact that something may not be green (something may not be dispositionally green-looking), yet nonetheless look green (because of a light effect, for example). I do not know whether there is more to the meaning of *green*, but this is all that needs to be made explicit in order to account for the green leaves case.

In the botanist scenario, there is a botanist in need of some green leaves (for the purposes of conducting an experiment). Pia knows that this is the case. If Pia is a cooperative speaker (in the Gricean sense), she should provide the botanist with as much information as she (the botanist) needs (that is, Pia should abide by the maxim of Quantity). The fact that the leaves are not naturally green is information that Pia is expected to communicate: indeed, for the botanist, given that she is running an experiment, it makes all the difference whether she gets leaves which are green by virtue of biological constitution or green by virtue of having been painted in green. Thus, Pia not revealing that the leaves are not naturally green is a blatant violation of the maxim Quantity: she should have said more. The botanist, upon hearing 'the leaves are green' is going to assume that the leaves are naturally green, for the same reason that, if one hears 'Sally cut the cake', one is going to assume that Sally cut the cake with a knife: give world knowledge, it is much (much) more likely than not that the leaves are green as a result of being naturally green than as a result of having been painted in green (cf. § 2.). Furthermore, the botanist has no reason to think that Pia conceals important information from her: if the leaves were not naturally green, Pia surely would have said so.

According to the analysis that I am presenting, Pia said something true: the green-painted leaves are, as a matter of fact, dispositionally green-looking; however, by concealing crucial information (the leaves are not naturally green), she misled the botanist into thinking that the leaves are naturally green. Misleadingness and falsehood, by making the speaker look dishonest, lend themselves to conflation; these notions, however, are importantly distinct: one can mislead someone by telling the truth (though maybe no by telling the whole truth), but one cannot say something false by saying something true.

[26] In case that it helps to spell out this further, 'x is dispositionally green-looking' is true if x looks green in the absence of any interference in the worlds in which there is light (or something along these lines). Of course, something can be green (i.e., dispositionally green-looking) and not be *naturally* green (i.e., dispositionally green-looking by virtue of biological constitution). I thank Benjamin Spector and Manuel Križ for extensive discussions on the meaning of 'green'.

Things are different in the photographer scenario: when the photographer hears the utterance 'the leaves are green', she, like the botanist, is bound to assume that the leaves in question are naturally green (as per world knowledge); the crucial difference, however, is that the photographer does not care whether the leaves are naturally green or whether they are green by artificial means: indeed, she needs the leaves for a photo shooting (and not for a botanical experiment). The photographer may end up with the (false) belief that the leaves are naturally green, but one has the strong intuition that Pia did not mislead the photographer: she said all that it was required from her in that context and did not conceal relevant information.

To sum up, according to the proposed account, 'the leaves are green' is a true sentence in both scenarios: what changes from one scenario to the other is how much information Pia is expected to give away under the assumption that she is a cooperative speaker. In one scenario, Pia conceals crucial information and, as a result, one's intuition about whether Pia's utterance is true is tainted by one's perception that Pia acted uncooperatively (i.e. she concealed crucial information, possibly with the intention to deceive the botanist); in the photographer scenario, she gives as much information as the context requires and, as a result, one's intuition about whether Pia's utterance is true is not confounded by considerations about cooperativity.

Why should one prefer this account over a 'lexicalist' account, which posits that, in the botanist scenario, *green* lexically narrows to *naturally green* or, alternatively, that *green* is disambiguated to *naturally green* (on the assumption that *green* is ambiguous)? There are several reasons:

(i) Occam's razor: the presumption that speakers are cooperative (e.g. do not conceal relevant information) is enough to account for the data: if the data can be explained by appealing to such a basic principle, why would anyone want to stipulate something exorbitant about the meaning of *green*?

(ii) If the green leaves case is evidence that *green* is context-sensitive, then more or less every word is context-sensitive.[27] Take, for example, the adjective *serrated*, which means, roughly, *sawlike*. Now, reproduce the green leaves case using *serrated* instead of *green*: Pia has a bunch of naturally serrated leaves and another bunch of leaves which she has cut in a sawlike fashion using scissors. Imagine that the botanist, on this occasion, is looking for serrated (instead of green) leaves, and that there is a photographer who is also looking for serrated leaves. The same effect is obtained. Are we going to conclude that *serrated* is context-sensitive? LN advocates may; after all, these theorists think that every content word can be modulated more or less freely. I would be surprised, however, if Kennedy and McNally wanted to claim that 'serrated' is ambiguous between *naturally serrated* and *serrated in looks*.

[27] Thanks to Émile Enguehard for pressing on this point.

(iii) Imagine that, in the botanist scenario, Pia had uttered 'the leaves are yellow': Pia's utterance would have clearly judged to be false. The lexicalist account has it that, in the botanist scenario, 'the leaves are green' means *the leaves are naturally green* (because 'green' is interpreted as meaning *naturally green*); according to this account, therefore, both 'the leaves are yellow' and 'the leaves are green' express false propositions (in the botanist scenario). Quite clearly, however, once one considers the matter carefully, there is a contrast (in terms of truth-valueness) between these two utterances. The reader is now invited to check their own intuitions. In the proposed account, this contrast is predicated: 'the leaves are green' is true but highly misleading (and hence may be mistaken by a false utterance), while 'the leaves a yellow' is false. The lexicalist approach (either LN or the ambiguity thesis), by contrast, fails to account for this basic contrast.

I have considered three reasons to prefer the misleading account (over the lexicalist approach). Is there any consideration that could tilt the scales back in the other direction? Perhaps. If one could find independent evidence that 'green' is ambiguous between *naturally green* and *green in looks*, for example, that would get the lexicalist account back into the race. Kennedy and McNally (2010) provide two arguments, based on consideration about gradability, that there is in fact a systematic ambiguity at the heart of colour terms. I will review their arguments in the next section.

5.1 On the purported ambiguity of colour terms

Kennedy and McNally (K&M) (2010) put forward the view that colour terms are ambiguous between a 'quality' gradable meaning (*green in looks*) and what they call a 'classificatory' meaning.[28] The latter is stipulated to denote a non-gradable property of individuals; in the case of leaves, this property would be something like *green by virtue of biological constitution* or the shorter *naturally green*.[29] If 'green', in addition to a quality-gradable meaning, were to be associated with a classificatory meaning in K&M's

[28] K&M (2010), in fact, stipulate that colour terms are associated to *two* distinct gradable meanings, a 'quality' meaning and a 'quantity' one (in addition to the 'classificatory' meaning). The 'quality' meaning, a function from individuals to degrees, measures how closely the object's manifestation of the colour approximates the appropriate prototype; the 'quantity' meaning, also a function from individuals to degrees, measures how much of the object manifests the colour. This quality/quantity ambiguity is not meant to play any role in explaining the green leaves puzzle and is motivated by independent considerations. I am sceptical that colour terms exhibit a quality/quantity ambiguity; an examination of this proposal, however, falls outside the scope of the present work.

[29] K&M (2010) provide the following lexical entry for the alleged 'classificatory' (non-gradable) meaning of colour terms (e.g. 'green'):

$[\![\text{green}^{\text{nongr}}]\!] = \lambda x.P_i(x) \wedge \mathbf{cor}(Pi, \mathbf{green})$

This adjective, of type $\langle e,t \rangle$, applies truthfully to its argument iff that argument has some other property which is correlated with the colour green. The correlation relation is indicated as **cor**, and the correlated property is the value of a free variable P_i. Though K&M treat *Pi* as a free variable, they assume that its value is conventionally determined by features of the object denoted by the colour adjective's argument (cf. K&M 2010: ftn. 13). Thus, when $[\![\text{green}^{\text{nongr}}]\!]$ is applied to leaves, the value for *Pi* will (presumably) be something like *having chlorophyll content* (which is correlated to the colour green).

sense, then the green leaves case would have a straight-forward solution, i.e.: in the botanist scenario 'green' would be disambiguated in favour of *naturally green* (the classificatory meaning) and, in the photographer scenario, in favour of *green in looks* (the quality-gradable meaning). However, the data that these theorists present in support of the claim that 'green' has a classificatory meaning among its associated senses is not at all persuasive.

K&M consider a modified version Travis's original experiment. Pia now has a pile of leaves that consists of both naturally green leaves of different shades and naturally red leaves painted in green, also of varying shades. Once again, there are two scenarios. First, Pia's photographer friend walks in and asks if she can have some green leaves to include in a mixed-media piece. Pia invites her to sort through the leaves and take any leaves she wants. Second, Pia's botanist friend walks in and asks if she can have some green leaves for a research project. Pia invites her to choose some leaves as well. The photographer and the botanist know beforehand—or find out while sorting through the leaves—that one pile has naturally green leaves, while the other doesn't. K&M point out that the photographer could justify her choices (if she was asked to justify them) by uttering (37)a, but also (38)a-b; the botanist, on the other hand, could only utter (37)a to justify her choices. Table 1 illustrates this point.

Table 1

Artist	**Botanist**
(37) a. This leaf is green.	(37) a. This leaf is green.
(38) a. This leaf is greener than that one.	(38) # a. This leaf is greener than that one.
b. This leaf is not as green as that one.	# b. This leaf is not as green as that one.

These facts, according to K&M, indicate that 'green' is ambiguous, as discussed, between a gradable 'quality' meaning, paraphrasable as *green in looks*, and a non-gradable 'classificatory' one, paraphrasable as *naturally green*. Indeed, according to these theorists, the fact that both (37)a and (38)a-b are acceptable in the photographer scenario shows that, in this context, the colour term being used is the gradable one; by contrast, the fact that, in the botanist scenario, (37)a is acceptable while (38)a-b is not, shows that, in this context, the colour term being used is the non-gradable one.

This analysis, I claim, is wrong.[30] To start with, I do not know that the botanist cannot utter (38)a: she can, as long as 'this leaf' refers to a naturally green leaf (one that was chosen) and 'that one' also refers to a naturally green leaf (one that was not chosen); indeed, she could be interested in selecting the greenest possible (naturally green) leaves available, because, for example, these may conveniently have a very high concentration of natural pigments. It is true, however, that the botanist cannot pick a

[30] Almost 50 years ago, and not without indignation, Zwicky and Sadock (1975: 4) complained: 'Philosophers perennially argue for ambiguities on the basis of a difference in understanding alone, and linguists are not immune either'.

naturally green leaf and then go on to justify her choice by uttering 'this leaf is greener than that one', where 'that one' refers to a green-painted leaf; the artist, on the other hand, can clearly do this.

The reason for such a contrast, however, is no mystery: given common ground knowledge, the fact that a naturally green leaf may be greener than a green-painted leaf is not something that the botanist would be expected to take into account (assuming that she is sane) when making a decision concerning which leaves to pick: if she is interested in running an experiment with naturally green leaves, then she is obviously not interested in a green-painted leaf, irrespective of whether there is a green-painted leaf that is greener (or less green) than a naturally green leaf. Utterances such as (38)a-b, therefore, are expected to be perceived as odd if produced by the botanist: these utterances, indeed, simply miss the mark.

K&M make a second point: (37)a, in the botanist scenario, could be used to justify the choice of a naturally green leaf over a green-painted one. I do not think that this is in fact the case. Consider, for example, the contrast between (39)a and (39)b. It is pretty clear to me that it would be weird for the botanist to utter (39)a if the two leaves being talked about were equally green-looking (the difference being solely that one of the leaves is naturally green and the other is not): one could of course *make sense* of what the botanist is trying to say, but that is a different matter. By contrast, (39)b, in the stipulated context, is a perfectly acceptable utterance.

(39) Context: The botanist is asked to justify why he chose a naturally green leaf over a green-painted leaf.

a. ?? I'm choosing this leaf, as opposed to that leaf over there, because *this leaf is green*.
b. I'm choosing this leaf, as opposed to that leaf over there, because this *leaf is naturally green.*

The point is that (39)a does not have (39)b as a reading, as clearly shown by the perceived contrast between the two examples. This contrast, in fact, runs against the very idea that 'green' can ever mean something like *naturally green*.

6 IMPLICIT DOMAIN RESTRICTION

It is a well-known fact that quantified NPs are context-dependent. For example, 'every school', in the examples below, means something different in each case: in (40), 'every' quantifies over the set of UK state schools, whereas in (41), it does over the set of French state schools.

(40) Context: The UK Government has decided to reduce the budget for education.
Every school is going to suffer.
⇒ Every *UK state school* is going to suffer.

(41) Context: The French Government has decided to reduce the budget for education.
Every school is going to suffer.
⇒ Every *French state school* is going to suffer.

The most influential proposal to handle this sort of data is what came to be known as the *C-variable account*: context provides a suitable domain of quantification, formalised as a predicate variable *C*.[31] Classically, *C* is introduced by the determiner (or quantifier) and the value of *C* is intersected with the denotation of the noun phrase. This account is associated with von Fintel (1994)'s seminal work[32], and I shall refer to it as the *classical account*. The competing proposal, due to Stanley (2000; Stanley and Szabó 2000), attaches *C* to the head noun instead, and has the value of *C* intersected with the denotation of the head noun (because of this, Stanley's account is known as 'Nominal Restriction Theory', NRT henceforth). Table 2 illustrates these two proposals.

Table 2

Classical account	$_{S}[_{DP}[_{D}\,[\text{Every}_{C}]\ _{NP}[\ _{N}[\text{school}]\ ]\ ]\ _{VP}[\ldots]]$ Interpretation rule: $[\![\text{Every}_{C}\ \ \text{NP}]\!]^{g} = [\![\text{Every}]\!]^{g}\ (\ [\![\text{NP}]\!]^{g} \cap \text{g}(C)\ )$
NRT	$_{S}[_{DP}[_{D}\,[\text{Every}]\ _{NP}[\ _{N}[\text{school}_{C}]\ ]\ ]\ _{VP}[\ldots]]$ Interpretation rule: $[\![\text{N}_{C}]\!]^{g} = [\![\text{N}]\!]^{g} \cap \text{g}(C)$

Stanley and Szabó's version of the *C*-variable account is, at bottom, a way (one among many) of executing the LN proposal (in the nominal domain); indeed, what NRT does, as shown in the table above, is to strengthen the denotation of the actual noun (by intersecting it with g(*C*)).[33] In the classical account, by contrast, no claim is made as to whether individual words can be strengthened: indeed, in von Fintel's (1994) algorithm, domain restriction does not affect the interpretation of the noun (i.e. on this account, g(*C*) is intersected with the NP, which is a *phrasal* constituent).

It should be noted that whether g(*C*) is intersected with the head noun (as in NRT) or the NP (as in the classical account) is not an arbitrary choice: it follows from where in the syntax *C* is stipulated to be

[31] To allow for binding, domain variables are often assumed to have a more complex structure: the standard proposal is that *C* is the composition of a functional variable *f* (of type $\langle e,\langle e,t\rangle\rangle$) and an argumental variable *x* (of type *e*) (cf. von Fintel 1994, Stanley and Szabó 2000, Stanley 2002).

[32] See also Westerståhl (1984) and Martí (2003).

[33] This is not the way LN advocates typically think of LN: indeed, in the writings of Carston (2002), Recanati (2004, 2010), and Ludlow (2014), LN is presented as a manifestation of the more general mechanism of lexical modulation (which also includes lexical broadening). Furthermore, for these authors, content words can be strengthened in context because these words are *inherently* context-sensitive (and not because they come with a context-sensitive element such as *C*). For the purpose of this discussion, these conceptual differences are irrelevant. It will be shown that *lexical* strengthening (irrespective of the question of how it exactly comes about) is not involved in the process of delivering domain restriction.

located and the compositionality principle (a basic desideratum that any semantic interpretation rule should satisfy). Indeed, if *C* is introduced at the D level, then g(*C*) cannot be intersected with ⟦N⟧ (i.e. if it was intersected with ⟦N⟧, then the meaning of the NP would not be exclusively determined by the meaning of its immediate daughters); conversely, if *C* is introduced at the level of N, a terminal node, then g(*C*) cannot be intersected with the ⟦NP⟧ (i.e. the meaning of the NP must be result of computing the meaning of N*c* with that of its sister—if it has a sister).

In sharp contrast with the *C*-variable approach is the *flexible universe* strategy;[34] according to it, the truth of an utterance is evaluated relative to a restricted universe of discourse: for example, given a first order model ⟨𝔇, I⟩, it is 𝔇, the domain (or universe) of discourse, that is amenable to contextual pressures. Whereas the *C*-variable approach is local (i.e. each DP accesses its own contextually supplied domain), the flexible universe solution is global.

The evidence favoring the local over the global approach comes from examples such as the following:

(42) Context: The MIT Linguistics & Philosophy department consists of linguists and philosophers, naturally. On one particular year, three academics are shortlisted for a teaching position, a linguist and two philosophers, but the department is only allowed to recommend one of them.[35]

Every linguist voted for the linguist.
⇒ Every linguist *from the MIT Linguistics & Philosophy department* voted for the linguist *from the pool of job candidates.*

What this example shows is that 'every linguist' and 'the linguist' can be evaluated relative to a different (contextually provided) set of entities—what von Fintel (1994) calls the *resource domain*; if this was not the case, the uniqueness presupposition triggered by the definite article would not be satisfied. Globally restricting the domain discourse can account for the interpretations in (40) and (41), but it is hopeless when it comes to (42); by contrast, having individual domain variables associated to each DP, can account for (40) and (41), but also for (42). Because of observations such as this (see also Soames 1986), the global approach fell out of favour, and the *C*-variable (local) account became the gold standard for doing domain restriction.

[34] As von Fintel (1994: 29, ftn. 15) remarks, this is a proposal with a venerable pedigree, which includes names such as John Wallis, August de Morgan, and George Boole. The name 'flexible universe strategy' comes from Westerståhl (1984).
[35] Adapted from Cooper (1996).

6.1 Lexical versus phrasal restrictions

The unsettled question is therefore whether domain restriction, which almost everyone accepts to be a local procedure,[36] involves strengthening at the lexical level (as in NRT), whether it comes about via a phrasal mechanism that originates higher up in the structure (as stipulated in the classical account), or whether it is a hybrid mechanism.[37] The question is important: if evidence can be found that domain restriction is computed at the lexical level (always or sometimes), then one could use this evidence to argue for the existence of LN (or, at least, for the existence of something like LN).

Stanley (Stanley 2002; Stanley and Szabó 2000) puts forward three arguments in support of NRT, arguments which turned out to be very influential; these arguments involve considerations about cross-linguistic anaphora, superlative modifiers and the pragmatic behaviour of relative adjectives. Soon afterwards, Kratzer (2004) presented the *Fake Philosopher* argument—ascribed to Breheny (2003); this argument has been taken to provide evidence that that *C* must be introduced by the determiner, *contra* NRT and as stipulated in the classical account. The *Fake Philosopher* argument has also been an influential one. In view of the conflicting evidence, the debate came to a deadlock, and some theorists have since suggested that the *C*-variable approach may have to be abandoned.[38]

The goal of the following section is to show that there is no evidence indicating that NRT (or something like NRT) exists. I will first show that Stanley's arguments in support of NRT are unsuccessful; though voices have been raised against NRT (e.g. Giannakidou 2004; von Fintel 2014) a thorough refutation of these arguments has, to my knowledge, not yet been provided. Second, I will examine the *Fake Philosopher* argument. This argument, though it provides evidence for the classical account, cannot be taken to be conclusive evidence against NRT (or so will I argue). To conclude, I will present a novel empirical argument based on the observed pragmatic behaviour of bare nouns that refutes NRT and indicates that the classical account is on the right track.

6.2 Reassessing the arguments

6.2.1 Stanley's arguments in support of NRT

[36] The exception being Schwarz (2009) who, building upon Kratzer (2007), has proposed a hybrid, situations-based system that performs both global and local restrictions. The minimal pairs reported in § 6.4.2, as far as I can tell, represent a serious challenge to this account.

[37] Indeed, the classical account and NRT are not mutually exclusive, though much of the literature is tinted with the presupposition that they are. It could be proposed, for example, that both determiners and nouns introduce a *C* variable. Etxeberria (2005), in fact, has claimed that domain restriction *á la* von Fintel as well as NRT are needed to handle Basque nominal quantification data.

[38] 'We currently have a number of good arguments supporting conflicting conclusions about where in the structure domain restriction variables are introduced. Unless we can debunk one set of these arguments, the outlook for this type of approach is not very promising.' (Schwartz 2009: 108-9)

The first argument, put forward in Stanley and Szabó (2000), is as follows. S&S observe that (43) can have two possible readings, depending on how cross-sentential anaphora is resolved, and claim that this sort of anaphora requires antecedents to be constituents (nodes) of a preceding LF.

(43) Context: Talking about a certain village…

Most people regularly scream. They are crazy.

Reading A: The people in the village are crazy.
Reading B: The people in the village that regularly scream are crazy.

S&S then note that placing *C* on the noun (*Most [people$_C$]*) allows a simple derivation of Reading A, since the pronoun *they* can have *people$_C$* as its antecedent.[39] However, if *C* were to be attached to the determiner (*Most$_C$ [people]*), there would be no antecedent constituent denoting the set of people in the village.

S&S's argument rests on the premise that the antecedent of cross-sentential anaphora must be a constituent. But Reading B, which S&S (rightly) claim to be a possible reading of (43), already refutes this premise.[40] Indeed, in the first sentence of (43), as illustrated in Fig. 1, there is no constituent anywhere in sight that denotes people who regularly scream. [41]

Figure 1

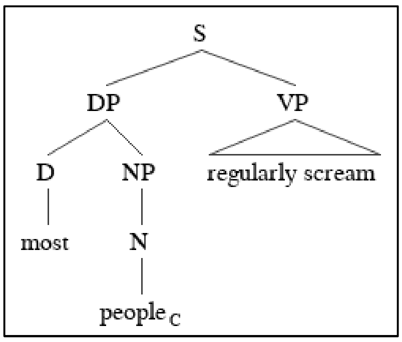

It should be noted that the existence of so-called 'complement anaphora' provides further evidence that S&S's argument is untenable (Schwarz 2009). Consider, for example, the following sentence: *Few congressmen admire Kennedy. They think he's incompetent* (Moxey and Sanford 1993). The most salient reading of this sentence is one in which the pronoun *they* picks out the 'non-admirers', despite there being no constituent in sight that has the 'non-admirers' as its denotation. Given these considerations, it is clear that that the principle that S&S put forward (i.e. the antecedent of cross-

[39] Stanley (2002) stresses that, on his account, *C* occupies the same node as the head noun.
[40] Thanks to Manuel Križ for pointing this out to me.
[41] Syntactic constituents are represented graphically as nodes in a syntactic tree.

sentential anaphora must be a constituent), principle which serves as premise in their argument, cannot be upheld.

The second argument, reported by Stanley (2002, ascribed to Delia Graff Fara), centres on the pragmatic behaviour of superlative noun phrases. Consider, for example, (44).

(44) Context: Talking about Cornell students…

The tallest student is nice.
⇒ The tallest student from Cornell is nice.

Stanley (2002) observes that, under the assumption that *tallest* takes the head noun as its argument and returns a set consisting of the tallest individual in the set denoted by the head noun, placing *C* on the determiner yields an odd result: namely, that (44) could only be truthfully uttered if the tallest student in the world happened to be a student at Cornell. If, on the other hand, *C* were to be attached to the noun, as stipulated in NRT, the correct interpretation would be obtained, as *student* would be restricted to *Cornell student* prior to combining with *tallest*.

This is not a strong argument in support of NRT: any modern semantic treatment of *-est* can generate the correct truth-conditions for (44), and it can do so without stipulating that domain restriction is computed on the noun.[42] Let's consider, for example, Heim's (1999) semantics for the superlative morpheme, which, in fact, builds upon von Fintel's (1994) account of domain restriction.

Heim (1999) assumes that a gradable adjective denotes a function from degrees to $\langle e,t\rangle$ functions (following Seuren 1973, Cresswell 1976, and others).[43] On this account, *tall* has the following meaning:

(45) For any degree d and individual x, $[\![\text{tall}]\!](d)(x) = 1$ iff x is tall to degree d (i.e. iff x's maximal height includes d).[44]

Modified nominal expressions such as *tall student*, where the modifier is a gradable adjective, also denote functions from degrees to $\langle e,t\rangle$ functions.

As is standard in semantics, Heim (1999) treats the superlative morpheme *-est* as a degree quantifier and, following von Fintel (1994), stipulates that this morpheme, by virtue of being a quantifier,

[42] See, for instance, Heim (1999), Farkas and Kiss (2000), and Sharvit and Stateva (2002).
[43] Expressions of type $\langle e,t\rangle$ denote functions that map entities to truth-values. Functions of this sort characterise sets of entities.
[44] One consequence of (45) is that the meaning of gradable adjectives becomes *downward monotonic*:
If R is a gradable adjective meaning, then $\forall x,d,d'$ $((\mathrm{R}(x,d) \wedge d > d') \rightarrow \mathrm{R}(x,d'))$. For example, if John is exactly 1.80m tall, then he is also 1.79m tall, 1.78m tall, etc., but not 1.81m tall or 1.82m tall.

introduces a domain variable *C*.[45] The upshot of this is that *-est* ends up taking three arguments: a gradable adjective, an individual, as well as a domain argument (that is, the value of *C* supplied by the context). Heim's (1999) semantics for *-est* are given in (46).

(46) Let *x* be an individual, R a gradable adjective, and K the relevant domain:

$[\![\text{-est}]\!]$ (x, R, K) is defined only if $x \in \mathrm{K} \wedge \forall y\, [y \in \mathrm{K} \rightarrow \exists d\, [\mathrm{R}(y,d)]]$. Whenever defined,
$[\![\text{-est}]\!]$ (x, R, K) = 1 iff $\exists d\, [\mathrm{R}(x,d)\ \&\ \forall y\, [y \neq x \wedge y \in \mathrm{K} \rightarrow \neg\, \mathrm{R}(y,d)]]$.

Let's now return to *tallest student* which, on Heim's (1999) account, is analysed as having the LF-constituency *[-est [tall student]]* (rather than *[[-est tall] student]*); according to the definition in (46), $[\![\text{-est}]\!]$ takes an individual *x*, a gradable adjective meaning (in the case at hand, the modified nominal *tall student*), and a domain argument (in this case, the set of Cornell students); it returns 1 iff *x* is a tall student to degree *d* and no other individual in K (the set of Cornell students) is a *d*-tall student. This semantics predicts the correct denotation for *tallest student* in the context stipulated in (44).

It is worth noting that Heim's (1999) treatment of superlative noun phrases has far greater empirical coverage than NRT. Indeed, whereas both accounts can handle (44), NRT has no resources to deal with (47), for example.

(47) <u>Context</u>: Talking about three brothers…
The tallest is nice.
⇒ The tallest (among the three brothers) is nice.

In (47), there is no set-denoting noun that can be restricted and, as a result, there is no way to derive (47)'s truth-conditions via NRT.[46] Heim's (1999) semantics for *-est*, by contrast, can handle (47) in a straight-forward manner.

The third argument, put forward in Stanley (2002), is as follows. Stanley observes that (48) may be true if talking about Smith's piano-playing at a dinner party, but not true if talking about Smith's piano-playing at formal concert setting.

[45] Heim (1999: 3) writes: 'Following von Fintel (1994), we may localize the context-dependency of quantifiers like *every* in an extra argument, a phonetically unrealized predicate variable that appears next to the determiner at LF and receives a value from the context of utterance. Adapted to the case at hand, this suggests that *-est* likewise takes an additional argument'.
[46] Any attempt to argue that, in (47), there is an unpronounced noun at LF is bound to fail. Indeed, such an account would predict (iii) to be felicitous, and it is not.

(iii) <u>Context</u>: Talking about three brothers…
\# The tall is nice [⇏ The tall (brother) is nice]

(48) Smith is a remarkable pianist.

Next, he argues that the perceived context-dependency of (48) can be captured if one assumes that *C* is located on the noun: in one case *pianist* would be intersected with the set of dinner party musicians and, in the other, with the set of professional musicians.

This is not an argument in support of having domain restriction computed on the noun though. (49), in the same way as (48), may be true if talking about Smith's piano-playing at a dinner party, but not true if talking about Smith's piano-playing at a formal concert setting: in (49), however, there is no set-denoting noun to be restricted.

(49) Smith is remarkable.

The standard assumption that relative adjectives access a comparison class (e.g. Bartsch and Vennemann 1972, Klein 1980) is enough to account for the context-dependency of (48) and (49). No appeal to domain restriction is needed.

The fact that none of S&S's arguments goes through does not prove NRT wrong, however. To rule out NRT, one needs to provide direct or indirect evidence that nouns cannot induce domain restriction.

6.3 The fake philosopher argument

On the basis of (50), Kratzer (2004) argues that nouns, *contra* S&S, cannot induce domain restriction.

(50) Every fake philosopher is from Idaho. (Kratzer 2004)

Kratzer (2004) assumes that the contextually provided domain for the DP *every fake philosopher* is the set of Americans (or US citizens). Given this domain, (50) should be interpreted as in (51)a, according to the classical account, and as in (51)b, according to NRT. (51), however, appears to be impossible.

(51) a. Every_$_{\text{C [American]}}$ fake philosopher is from Idaho.
⇒ Every American (fake philosopher) is from Idaho.
b. Every fake philosopher_$_{\text{C [American]}}$ is from Idaho
⇒ Every fake (American philosopher) is from Idaho.[47]

[47] ⟦*fake* (*American philosopher*)⟧ = {*x* : *x* is either a fake American and real philosopher OR an American and fake philosopher OR a fake American and a fake philosopher}. Evidence that this the denotation of 'fake (American philosopher)' comes from the fact that one can truly say 'she is a fake American philosopher' if/when *she* refers to (a) a real philosopher who pretends to from the US, (b) a US citizen who pretends to a philosopher, or (c) a person who pretends to be from the US and also pretends to be a philosopher.

This seems right. Let's imagine a context in which US citizens are being talked about. If (50) were to be uttered in such a context, it would surely be interpreted as meaning (51)a, and not (51)b. This becomes clear if one considers the existence of a genuine German philosopher who pretends to be American. Such a person would not be a counterexample to (50), in the stipulated context, nor would it be a counterexample to (51)a, but it would count as a counterexample to (51)b.

The *Fake Philosopher* argument has two sides to it: (i) NRT predicts (51)b, a reading that does not exist; (ii) NRT fails to predict (51)a, a reading that exists. (ii) is a good argument, and shows that *C*, the resource domain, can be intersected with the NP-restrictor, as stipulated in von Fintel (1994); however, this argument is silent on the question of whether nominal restriction *á la* S&S is (ever) possible. If one assumes that *C* appears at one, and only at one, syntactic position, as is generally assumed in the literature, then the mere existence of (51)a rules out the NRT proposal; however, there is no justification for making such an assumption: *C* variables could, in principle, be present at multiple locations (cf. ftn.37). (i), on the other hand, is not a conclusive argument, because it presupposes that, if nouns introduced *C* variables, domain restriction would be computed below a non-intersective adjective. This, as will be shown below, cannot be presupposed.

Consider (52):

(52) Every school in the UK is free.

Such an utterance would strike anyone as false; notably, if the interpreter of (52) could domain-restrict *every school in the UK* to *every* ***state*** *school in the UK*, the sentence would be saved from falsehood. This is not how domain restriction works, however: it appears not to be possible to restrict the domain of the quantifier to an arbitrary set of entities, even if such a restriction would keep afloat the presumption that the speaker follows the maxim of Quality (Grice 1975). The standard assumption, which I think is in-keeping with the domain restriction data, is that *C* is valued relative to those entities that are being talked about (or are relevant) in a given context (e.g. von Fintel 1994, Gillon 2006, Schwarz 2009).[48] From this perspective, domain restriction is a relevance-guided mechanism, one that connects what has been uttered to what is being talked about.

---

[48] If (52) was to be uttered in a context that made it clear that British state schools (and only these schools) are being talked about, then the set of entities that *every* quantifies over would get restricted to the set of UK state schools. Consider (iv) below, for example.

(iv) Context: Two British civil servants and an education specialist from the World Bank are having a discussion around specific problems facing the British state-schooling system. The education specialist points out that, according to her sources, several schools in the North of the country have started to charge fees, which imposes a heavy burden on many poor families. She advises that Government funds should be allocated to help the affected families. One of the British civil servants, visibly confused, replies:

But… is this even possible? *Every school in the UK is free.*

Thus, for *C* to be valued to the set of US citizens, (50) should be uttered in a context in which the US citizens are being talked about. Now, if *C* was attached to the head noun, as stipulated in NRT, the resulting NP would be 'fake philosopher_$_{g(C)\,=\,\{x\,:\,x \text{ is American}\}}$', whose denotation includes, for example, German philosophers who pretend to be Americans. In other words, in (50), *every* would end up quantifying over the set of fake (American philosophers), which has among its members individuals who are not being talked about (or are not relevant at that particular point in the discourse).

Let's assume that the following constraint is in place: domain restriction can be performed insofar as the resulting denotation of the NP-restrictor is either the set of entities that are being talked about or a strict subset of this set. Once this constraint is in place, domain restriction will never be computed below a non-intersective adjective (even if the head noun introduced a *C* variable). The reason for this is straightforward: as illustrated in (53) below, the set of entities that results from combining the domain-restricted noun with a non-intersective adjective cannot be a subset of the set of entities that is being talked about (on the assumption, of course, that the set of fake Americans who are philosophers or pretend to be philosophers isn't empty):[49]

(53) $$\underbrace{[\![\textbf{\textit{fake}} \textbf{ (American philosopher)}]\!]}_{\textit{NP denotation}} \not\subseteq \underbrace{\{\boldsymbol{x : x} \textbf{ is American}\}}_{\textit{what is being talked about}}$$

I am, quite clearly, stipulating that there is such a constraint. Notice, however, that it is not unreasonable to think that a constraint of this sort exists: if it did not, domain restriction could have the effect of introducing entities that are not being talked about (and are therefore irrelevant) in the discourse, while the function of this mechanism appears to be restricting quantifier domains to those entities that are being talked about.

In sum, the *Fake Philosopher* argument shows that value of *C* can be intersected with the NP, as stipulated in the classical account. These data, however, is not conclusive on the question of whether NRT is ever possible. First, showing that *C* can be intersected with the NP does not show that *C* cannot be intersected with the head noun. In addition, if *C* could be intersected with the head noun, there are reasons to think that, in (50), considerations about relevance would prevent this operation from being performed.

In the next section, I shall present a novel empirical argument in support of the classical account, one that rules out NRT conclusively.

[49] Cf. ftn. 47 for a break-down of the meaning of *fake* (*American philosopher*).

6.4 *The argument from bare nouns*

If one was able to show that, in the absence of an overt determiner, domain restriction is not possible, then one could reject NRT: this account, indeed, predicts restrictions *to be independent* from there being an overt determiner present in the structure.

To this end, I shall look at two well-researched DP constructions that lack an overt determiner: (i) bare plurals, which can be used generically or existentially (i.e. 'pigeons are nice' vs. 'I feed pigeons on Sundays'), and (ii) bare mass nouns, which can also be used generically or existentially ('water is good for you' vs. 'I've just drunk water'). Whether a covert determiner is present in these constructions is a controversial matter; for example, it has been proposed that the generic reading of bare nouns involves covert generic quantification (Lewis 1975).

6.4.1 *Domain restriction and generics*

Minimal pairs such as (54)a-b appear to indicate that generic DPs do not support domain restriction.

(54) Context: There are lions and tigers in the cage…[50]

a. Every lion has a mane.
⇒ Every lion (in the cage) has a mane.
b. The lions have a mane.
⇒ The lions (in the cage) have a mane.
c. # Lions have a mane.
d. ?? Lions **in the cage** have a mane.

(54)a and (54)b are perfectly felicitous and, in the context stipulated, appear to mean *every lion in the cage has a mane* and *the lion in the cage has a mane*, respectively. (54)c, on the contrary, feels odd. On the basis of these data, one could be tempted to argue that, in the absence of an *overt* quantifier, domain restriction is not possible. If this is in fact the case, then LN and NRT cannot be right. I am not sure, however, that domain restriction is not possible in (54)c. Indeed, in (54)d, 'lions' is *explicitly* restricted and, despite this, the example continues to be odd—perhaps a bit less odd than (54)c, but degraded relative to (54)a-b. Thus, it is unclear to me whether (54)c is odd because domain restriction is not

[50] (54)a-c are due to Croft (1986) and Krifka (1987).

possible, or whether it is odd for the same reason as (54)d, which is explicitly restricted, is odd (whatever this reason might be).[51]

Furthermore, examples have been reported indicating that generic DPs can be domain-restricted; consider (55), for example, which comes from Condoravdi (1994).

(55) Context: In 1985 there was a ghost haunting the campus . . .

a. Students were aware of the danger.
b. The students were aware of the danger.
c. There were students who were aware of the danger.

(55)a seems to be more or less equivalent to (55)b, but not to (55)c. Condoravdi (1994) argues that (55)a is not a *real* generic reading, but a separate reading, which she calls 'functional'. Whether bare plurals do in fact give rise to a third reading is a controversial issue (and not a particularly relevant one in the context of this discussion). If (55)a is a generic reading, then generic DPs, in certain contexts, must support contextual domain restriction; if it is not, then there must be some determinerless (non-existential) DPs that that, in certain contexts, support contextual domain restriction.

To sum up: one way to refute NRT would be to show that domain restriction is correlated to the presence (or absence) of an overt determiner; however, as discussed above, it is not clear from the pragmatic behaviour of bare plurals in non-existential uses that there is such a correlation.

6.4.2 Existential bare nouns and domain restriction

To circumvent the problems encountered in the generics data, I shall test whether domain restriction is possible with mass nouns and bare plurals (in non-generic, existential uses). To my knowledge, this has not been done before.[52]

Let's consider the following example, which comes from Stanley (2002):

[51] It has long been noted that generics do not like to be explicitly restricted to a time or a place—or to a time and a place (Carlson 1982). Consider, for example, the contrast between (v)a and (v)b.

(v) a. Strawberries are delicious.
b. # Strawberries in the box are delicious.

[52] I am only aware of some examples (*not* minimal pairs) reported in Arregui (2008), which suggest that indefinites without an overt determiner may fail to take restricted domains in the scope of negation. The minimal pairs presented here reveal that determinerless indefinites do not support DR, irrespective of whether negation is present.

(56) Context: Pastor Hannah is concerned about the fact that someone has been drinking the holy water in her church on warm summer days. In a discussion with John, John confesses:

I drank a little water last week.

What John expresses, as Stanley (2002: 19) remarks, is the proposition that John drank a little of the church's holy water. Stanley takes (56) to indicate that *C* is attached to the noun, as opposed to the determiner (I am not sure why): (56), I shall argue, provides evidence of the exact opposite.

Let's consider (57), against the context stipulated in (56): if *water* could be domain-restricted to *holy water from the church*, then (57)b and (57)c should express the same proposition: this is clearly not the case, however.

(57) a. I drank a little water last week.
⇒ I drank a little holy water from the church last week.
b. ? I drank water last week.
⇏ I drank holy water from the church last week.
c. I drank holy water from the church last week.

Furthermore, the contrast between (57)a and (57)b indicates that the presence of *a little* makes domain restriction possible (notice that (57)a works too if *a little* is replaced by *some*).

The pattern in (57) can be reproduced using a bare plural noun like *apples*, as illustrated in (58).

(58) Context: Pastor Hannah is concerned about the fact that someone has been eating the apples from the churchyard's apple tree on warm summer days. In a discussion with John, John confesses:

a. I ate some apples last week, please forgive me.
⇒ I ate some apples from the churchyard's apple tree last week, please forgive me.
b. ? I ate apples last week, please forgive me.
⇏ I ate apples from the churchyard's apple tree last week, please forgive me.
c. I ate apples from the churchyard's apple tree last week, please forgive me.

Once again, if *apples* could be domain-restricted to *apples from the churchyard's apple tree*, then (58)b should be as acceptable as (58)c and, interestingly, it is not. The contrast between (58)a and (58)b, furthermore, shows that the presence of *some* makes domain restriction possible.

More or less the same results are obtained if the target sentence is negated, as illustrated in (59).

(59) Context: Pastor Hannah is concerned about the fact that someone has been eating the apples from the churchyard's apple tree on warm summer days. In a discussion with John, John protests:

a. I didn't eat any apples last week, so I'm not sorry.

⇒ I didn't eat any apples from the churchyard's apple tree last week, so I'm not sorry.
b. ? I didn't eat apples last week, so I'm not sorry.
⇏ I didn't eat apples from the churchyard's apple tree last week, so I'm not sorry.
c. I didn't eat apples from the churchyard's apple tree last week, so I'm not sorry.

The contrasts in (59) are a bit less sharp than the one reported in (58), the reason for this being two-fold. First, *any* is typically associated with domain-widening effects (Kadmon and Landman 1993); as a result, it is expected to interfere with the mechanism of domain restriction. In addition, (59)b is not as pragmatically infelicitous as (58)b. Indeed, assume that the (implicit) question that John is trying to addresses is 'have you eaten apples from the churchyard's apple tree last week?': then (59)b entails a negative answer to this question (i.e. 'I didn't eat apples last week' entails 'I didn't eat apples from the churchyard's apple tree last week'), whereas (58)b is clearly irrelevant (that is, it does not settle the question).

Given these data, it should be clear that LN (in its NRT version, or in any other version) is not involved in the process of domain restriction; if the pragmatic readings of (57)a, (58)a and (59)a were the result of the noun (*water* or *apples*) being context-sensitive (by virtue of having *C* attached, or by virtue of being itself context-sensitive), then it would be a complete mystery why in (57)b, (58)b and (59)b the very same noun does not exhibit this context-sensitivity. These contrasts, in this regard, do not just show that domain restriction cannot be characterised in terms of LN: it also shows (or at least strongly suggests) that LN is not a thing: if such a mechanism existed, then (57)b, (58)b and (59)b should be felicitous in their respective contexts (and they are not).

As stipulated in von Fintel's (1994), it must be the determiner, as opposed to the noun, that induces domain restriction. Two scenarios compatible with von Fintel's proposal must be considered, however. In (57)b-(59)b, the bare noun could be combining with a silent existential quantifier or, alternatively, with no quantifier at all, the D position being ∅. If the former is the case, then it means that covert existential quantification does not introduce domain variables; alternatively, if the bare noun combines with no quantifier at all, the D position being ∅, then a type-shifting rule must be posited to derive (57)b-(59)b's existential force.

7 CONCLUSION

From the data and analyses presented in this chapter, two things can be concluded. First, those pragmatic considerations that can, at the compositional stage, disambiguate an ambiguous word, assign a referent to a pronominal element, fix a contextual parameter such as the standard of comparison for a relative adjective or, indeed, provide a suitable domain of quantification, do not appear to be able to strengthen the denotation of individual content words. Second, there are processes and mechanisms operative in natural language interpretation that *create the illusion* that the meanings of content words are massively context-sensitive. In this chapter, I have identified seven of them, in the following order: world-knowledge inferences (§ 2), local exhaustification (§ 3), intensification of threshold-sensitive gradable adjectives via stress placement (§ 4.1), non-literal (rhetorical) usage of *and*-contradictions (§ 4.2.1 & 4.3.2), implicit precisification (e.g., a silent *definitely* operator) (§ 4.2.2 & 4.3.1.), misleadingness (§ 5), and domain restriction (see § 6).

The LN account (in its standard formulation) does not seem very plausible; whoever wants to argue for a constrained version of LN has a two-fold challenge ahead of them. This challenge would involve, one the one hand, putting forward an account of LN that makes sense of the negative facts: that is, an account that explains, in a principled way, why, for example, 'cut' does not modulate to *cut with a knife* in (1)b, why, in (24), 'upset' does not modulate to *terribly upset* (or to *only mildly upset*), why, in (57)b, 'water' does not modulate to *holy water from the church*, or why 'green', in (39)a, does not modulate to *naturally green* (just to mention some cases). On the other hand, it should be made explicit what data motivates the constrained LN account (the data here reviewed, as has been discussed, can hardly been taken as indicative of the existence of LN).

Chapter III
THE CHARACTER OF IMPRECISION

1 OVERVIEW

The lexical modulation hypothesis, as discussed, can be broken down into two sub-hypotheses: lexical narrowing (i.e. content words can be logically strengthened) and lexical broadening or loosening (i.e. content words can be logically weakened). When contextualists talk about lexical broadening, they appear to have in mind two phenomena: cases of what looks like loose (or imprecise) interpretations of absolute adjectives, numerals and geometrical terms, and cases of what looks like individual words being interpreted hyperbolically. In the first part of this chapter, I will argue that different interpretative mechanisms are at work across these cases and, furthermore, that none of these mechanisms can be adequately characterised in terms of lexical broadening. In the second part, I will prepare the ground for the next chapter: I will characterise the phenomenon of imprecise interpretation, discuss the relationship between imprecision and vagueness, as well as review Lasersohn's (1999) influential account of imprecision.

2 LEXICAL BROADENING

There are two phenomena which have been argued to be just instances of lexical broadening (LB henceforth): on the one hand, seemingly imprecise interpretations of absolute adjectives (e.g. 'empty', 'straight'), geometrical terms, and numerals; on the other hand, seemingly hyperbolic uses of words (e.g. 'boiling', 'impossible')—see, for example, Carston (2002, 2012), Sperber and Wilson (2008), and Wilson (2011).[53] By means of illustration, consider the following two cases:

(1) *Imprecision*

Context: John has just moved into Ann's house.

a. *John*: Where can I put my books?
b. *Ann*: The bookcase in the living room is *empty*.

(2) *Hyperbole*

Context: Julius is struggling to do his math homework; he complains to his mother…

This exercise is *impossible.*

[53] I write 'seemingly' because, as will soon be discussed, there is no evidence suggesting that individual words can be interpreted imprecisely or hyperbolically.

(1)b, at least intuitively, can be used to describe a situation in which there are two or three books left in the living room's bookcase; (2), likewise, can be used to describe an exercise that is very difficult (yet not impossible). The standard LB account of these examples can be summarised as follows: in (1)b, the word *empty* is LB-adjusted to *more or less empty*, whereas, in (2), the word *impossible* is LB-adjusted to *very difficult*; the mechanism of LB adjustment feeds composition and, as a result, has an impact on the truth-conditional profile of (1)b and (2).

On this account, imprecision and hyperbole are subsumed into the general category of LB phenomena; Wilson (2011), for example, writes: 'the interpretation of virtually any utterance involves some form of lexical broadening, in which the concept communicated by use of a word is more general than the lexical meaning' (p. 200); '[t]here is a continuum of cases of broadening, ranging from strictly literal use, through various shades of approximation to hyperbole and metaphor, with no sharp cut-off point between them' (p. 199); 'the distinction between approximation and hyperbole has no theoretical significance: an utterance does not have to be recognised as an approximation or hyperbole to be understood, no special interpretive mechanisms are needed in either case, and both are understood in the same way' (p. 200).

In what follows, I will show that this account, popularised by relevance theorists, is incorrect: imprecise interpretation and hyperbole are, as a matter of fact, different phenomena; moreover, none of these phenomena can be adequately characterised in terms of LB.

2.1 Imprecision versus hyperbole

The main datum that indicates that imprecision and hyperbole are different phenomena is the following: imprecise interpretation can be blocked by words like *completely,* known in the literature as *slack regulators* (Lasersohn 1999), whereas hyperbole cannot. For example, if, in (1), Ann's reply had been 'the bookcase in the living rooms is *completely* empty', her utterance would have been judged false if the bookcase had one or two books in it (that is, her utterance would have been interpreted as entailing that there were no books whatsoever in the bookcase). Conversely, in (2), if Julius had said 'this exercise is *completely* impossible', his audience wouldn't have inferred that the exercise in question was, as a matter of fact, unsolvable: indeed, they would have inferred more or less the same as if Julius had just uttered (2)—that is, that the exercise was very difficult. The contrast is clear: whereas the introduction of *completely* has the effect of removing imprecision, it does not have the effect of removing hyperbole (if anything, it appears to intensify the hyperbolic effect).

There is another important difference between these two phenomena: hyperbolic statements such as (2) (consider, also, 'the water in the bath is boiling', 'He's smarter than Einstein', 'Those shoes cost me

10,000 dollars') can be used to communicate affect (either positive or negative) about some state of affairs; for example, by uttering (2), Julius appears to convey how frustrated he is about the fact that he has been asked to solve such a difficult exercise (for the connection between hyperbole and affect, see, for example, Colston and Keller 1998; McCarthy and Carter 2004; Kao et al. 2014). Imprecise interpretation, by contrast, does not seem to be correlated in any way to the expression of affect; indeed, (1)b means what it means but, intuitively, does not disclose anything about the speaker's affective state in relation to the bookcase not having many books in it. This, once again, suggests that hyperbole and imprecision are phenomena of very different nature.

2.2 Imprecision, hyperbole, and lexical broadening

Is there any evidence indicating that imprecise interpretation and/or hyperbole come about via LB? No; in fact, the data that I have been able to gather suggest that these phenomena are post-compositional in nature and not the result of weakening the meaning of a lexical constituent at the compositional stage. Let's first consider (3).

(3) Context: John has just moved into Ann's house.

a. *John*: Where can I put my books?
b_1. *Ann*: # The bookcase in the living room is *empty*, though the first shelf has some books in it. (Maybe use the second shelf).
b_2. *Ann*: The bookcase in the living room is *more or less empty*, though the first shelf has some books in it. (Maybe use the second shelf).

If *empty* in (3)b_1 could be LB-adjusted to *more or less empty*, then (3)b_1 should mean the same as (3)b_2: the fact that (3)b_1 does not get to mean (3)b_2, that is, the fact that (3)b_2 is (or at least feels like) a contradiction, indicates that the mechanism responsible for delivering (1)b's imprecise interpretation is not LB. Indeed, there is no reason I can think of as to why *empty* in (3)b_1 would not be logically weakened if it supported LB: such an operation would prevent the sentence from being a contradiction and, by so doing, save both Quality and Relevance (cf. Chapter I, § 1). I will return to examples such as (3)b_1 in § 3.1, as there is more to say about them.

The same test and reasoning can be applied to (2), as illustrated in (4) below.

(4) Context: Julius is struggling to do his math homework; he complains to his mother…

a. The exercise is *impossible* (to do).
b. The exercise is *very difficult* (to do).
c. # The exercise can be done but it is *impossible* (to do).
d. The exercise can be done but it is *very difficult* (to do).

If (4)a in fact expressed the same proposition as (4)b (because *impossible* was LB-adjusted to *very difficult*, or any other weaker alternative), then why is it that (4)c does not express the same proposition as (4)d? That is, if hyperbole interpretation was the result of LB being applied on a some word or other, then why is that LB can be applied on *impossible* in (4)a but not on *impossible* in (4)c? It is not clear to me that a sensible answer to this question can be given.

I can think of two additional arguments against an LB-treatment of hyperbolic utterances; the first argument is based on (5)a, whose literal meaning is identical (logically equivalent) to that of (2).

(5) a. This exercise isn't doable.
b. This exercise is *very difficult*.

What is interesting about (5)a is that, just like (2), it can be interpreted hyperbolically: in fact, (2) and (5)a could be used interchangeably in the context stipulated in (2). In (5)a, however, there isn't a word which, if appropriately LB-adjusted, would generate (5)b; indeed, to generate (5)b, *phrasal* weakening must be performed. Given (5)a, therefore, it should be clear that hyperbole, as phenomenon, cannot be appropriately described in LB terms.[54]

There second argument is as follows. Let's assume that the perceived (hyperbolic) interpretation of (2) came about via LB: that is, let's assume that *impossible* were to be weakened to *very difficult* at the compositional stage. That would predict (2) to be truth-conditionally equivalent to (6) below.

(6) This exercise is very difficult.

Now, it is clear to me that (2) and (6) differ in terms of the net communicated content: indeed, the former expresses much more negative affect (e.g. frustration, anger) than the latter. LB fails to account for this important fact: indeed, LB predicts (2) and (6) to have the same literal meaning: but, if this is the case, then why is it that, if uttered in the same context, (2) and (6) do not convey the same information? In short LB, equates hyperbole with its non-hyperbolic paraphrase but, by so doing, excludes any possibility of explaining what is distinctive about hyperbolic utterances. Note that this point holds irrespective of whether *impossible* is LB-adjusted to *very difficult*, *extremely difficult*, or any other logically weaker modulation.

Hyperbole, as far as I can tell, is a speech-act-level phenomenon, one that involves the speaker *dramatising* certain situation, while relying on the listener to realise that a dramatisation is being played

[54](5), to be clear, indicates that hyperbole is not a lexical mechanism, though it is silent on the question of whether hyperbole feeds composition. The datapoint that suggests that hyperbole does not feed composition is (4) (and, I believe, (6), which is discussed next).

out. The success of hyperbole in communicating affect is, I suspect, a consequence of this dramatisation: when Julius, in (2), utters *the exercise is impossible* (or, in (5)a, *the exercise isn't doable*), he seems to be conveying that he feels as frustrated *as if* the exercise was in fact impossible (or as frustrated *as if* the exercise was in fact not doable).[55] Hyperbole, in this regard, appears to involve something like a trade-off: truth is sacrificed for the expression of affect (cf. Bergen 2016: 149). No more will be said about hyperbole, neither in this chapter nor anywhere else in the thesis. In the next section, I will be concerned with characterising the phenomenon of imprecision.

3 IMPRECISION

Imprecise interpretation is not just observed in the presence of an absolute adjective such as *empty*; there are other expressions/constructions that appear to support imprecise interpretation, including numerals (and measure phrases in general), geometrical terms, definite plurals, and habituals (see Chapter IV, § 2.1), just to mention a few. To illustrate the main features of imprecise interpretation, I will use absolute adjectives and definite plurals; in § 3.2, however, I will discuss other constructions that can also be interpreted imprecisely.

3.1 The defining features of imprecise interpretation

Lasersohn's (1999) *Pragmatic Halos* was, to my knowledge, the first attempt to make sense of the phenomenon of imprecision in a systematic way. His theory does not fare well against the negation data, as will be discussed in § 3.4; notwithstanding this, he made some important observations concerning the general character of the phenomenon which are worth reviewing. To start with, Lasersohn observed that imprecise interpretation is context-dependent; consider (1) again—given as (7) below.

(7) Context: John has just moved into Ann's house.

a. *John*: Where can I put my books?
b. *Ann*: The bookcase in the living room is *empty*.

As noted, in the context stipulated, (7)b is an acceptable sentence even if the bookcase in question turns out to have a few books in it; indeed, though perhaps not strictly true, the sentence is nonetheless *true enough* of such a situation (cf. Lewis 1979). This is thus a context that allows (7)b to be used (and to

[55] This picture is of course compatible with the observation that *completely* does not remove but intensifies hyperbole: for example, the addition of *completely* in (2) would make the dramatisation of the exercise's level of difficulty even more dramatic and, by so doing, the hyperbolic gesture even more explicit.

be interpreted) imprecisely. Let's now see what happens if (7)b were to be uttered in a context such as the one stipulated in (8).

(8) Context: John is looking for a book (any book) for a photo production.

a. *John*: Do you know if there are any books in the living room?
b. *Ann*: The bookcase in the living room is empty.

If John were to go to the living room and find that there are, in fact, three books in first shelf, he would inevitably conclude that Ann said something false: that is, in the context stipulated in (8), 'the bookcase in the living room is empty' is interpreted strictly or precisely and, as a result, does not tolerate exceptions: it is true if no book is in the bookcase and false otherwise.

The same context-dependency is observed with definite plurals; let's first consider (9).

(9) Context: George has just been given a new drum kit for Christmas. It's 11pm.

a. George: Can I play the drums mum?
b. Ann: Of course not. *The townspeople* are asleep.

Quite clearly, the truth of (9)b admits exceptions: that is, if a couple of people are still awake, then (9)b can still be truly uttered. This permissiveness of meaning, however, vanishes in the context stipulated in (9).

(10) Context: The US army is (secretly) testing a chemical weapon that can kill in a matter of seconds; the weapon, however, has a special 'selection' property: it does not affect people who are asleep, only people who are awake. To make sure that the weapon works correctly, the plan is to release it in a small town far from big cities late at night, when the town dwellers are asleep. No one must be killed during the testing procedure. The army captain, after inspecting each of the town houses, instructs the weapon technician…

You can release it now. The townspeople are asleep.

Indeed, in (10), 'the townspeople are asleep' does not tolerate exceptions: if there are two or three people still awake by the time the army captain uttered this sentence, what he said is quite plainly false. Thus, whether a precise or imprecise interpretation obtains depends on features of the context: as shown, the sentences 'the bookcase in the living room is empty' and 'the townspeople are asleep' can be interpreted imprecisely in (7) and (9), respectively, but only a precise interpretation is admissible in (8) and (10).

Another important point, also due to Lasersohn (1999), concerns the possibility of what he calls (overt) *slack regulation*; for both 'the bookcase in the living room is empty' and 'the townspeople are asleep', language seems to have lexicalised resources that prevent these sentences from being interpreted imprecisely, as illustrated in (11):

(11) a. The bookcase in the living room is *completely* empty.
b. *All* the townspeople are asleep.

The sentences in (11) cannot be interpreted imprecisely; for example, (11)b is false if one person in the town is awake, *even* in the context stipulated in (9). Likewise, (11)a is false if there is one book in the bookcase, even in the context envisaged in (7). To better visualise the imprecision removal effect of *all* and *completely*, consider the following minimal pairs:

(12) Context: George has just been given a new drum kit for Christmas. It's 11pm.

I. a. *George*: Can I play the drums mum?
b. *Ann*: No dear. The townspeople are asleep.
c_1. *George*: ?? That's not true. The neighbours downstairs are still awake.
c_2. *George*: Strictly speaking, that's not true. The neighbours downstairs are still awake.

II. a. *George*: Can I play the drums mum?
b. *Ann*: No dear. ***All*** the townspeople are sleep.
c_1. *George*: That's not true. The neighbours downstairs are still awake.
c_2. *George*: ? Strictly speaking, that's not true. The neighbours downstairs are still awake.

(13) Context: John has just moved into Ann's house.

I. a. *John*: Where can I put my books?
b. *Ann*: The bookcase in the living room is empty.
c_1. *Paul (Ann's husband)*: ?? That's not true: there are two books in the first shelf.
c_2. *Paul (Ann's husband)*: Strictly speaking, that's not true: there are two books in the first shelf.

II. a. *John*: Where can I put my books?
b. *Ann*: The bookcase in the living room is ***completely*** empty.
c_1. *Paul (Ann's husband)*: That's not true: there are two books in the first shelf.
c_2. *Paul (Ann's husband)*: ? Strictly speaking, that's not true: there are two books in the first shelf.

Let's focus on (12)I; (12)I.c_1 is an odd answer to (12)I.b: it feels as if George was somehow missing the point; (12)I.c_2, by contrast, does not have this problem and, though it sounds pedantic, is not pragmatically deviant. The difference between (12)I.c_1 and (12)I.c_2 is that the latter, but not the former, is headed by *strictly speaking*; what this expression appears to do is to signal that George is answering to a strict (or precise) interpretation of Ann's turn (an interpretation which is clearly not intended by Ann). By signalling what he is doing, George eludes pragmatic deviance: he is perceived as playing some sort of metalinguistic game (i.e. one that involves being pointlessly pedantic). In (12)II, the introduction of a slack regulator (*all* in this case) reverses the situation: (12)II.c_1 is perfectly fine, whereas the addition of *strictly speaking* in (12)II.c_2 feels redundant. This contrast is expected if *all* has the effect of forcing a strict interpretation of (12)II.b: indeed, on the one hand, (12)II.c_1 now becomes a legitimate objection (Ann said that *all* the townspeople are sleeping and this happens not to be true); one the other hand, it is not clear what *strictly speaking* is doing in (12)II.c_2, since (12)II.b can only be interpreted strictly. Exactly the same pattern is observed in (13).

So far, it has been shown that 'the bookcase in the living room is empty' and 'the townspeople are asleep' can be interpreted either imprecisely or strictly/precisely (depending on the context in which they are uttered); furthermore, it has been shown that the addition of a slack regulator prevents these sentences from being interpreted imprecisely. Next, I will comment on an additional feature that sentences that support imprecise interpretation exhibit: this feature (which Križ (2015), in the context of plurals sentences, calls 'unmentionability of exceptions') is here referred to as 'strict behaviour in contradiction test'.

Remember that, in § 2.2., I used example (3), given as (14) below, to argue against an LB-based account of imprecision.

(14) Context: John has just moved into Ann's house.

a. *John*: Where can I put my books?
b_1. *Ann*: # The bookcase in the living room is *empty*, though the first shelf has some books in it. (Maybe use the second shelf).
b_2. *Ann*: The bookcase in the living room is *more or less empty*, though the first shelf has some books in it. (Maybe use the second shelf).

The same test can be applied to a sentence containing a definite plural; consider, for example, (15).

(15) Context: George has just been given a new drum kit for Christmas. It's 11pm.

a. *George*: Can I play the drums mum? It looks like Smith next door is awake.
b_1. *Ann*: # Of course not. The neighbour next door is awake, but the neighbours are asleep.
b_2. *Ann*: Of course not. The neighbour next door is awake, but *most of the neighbours* are asleep.

In addition to (severely) calling into question the LB thesis, the contradictoriness of (14)b_1 and (15)b_1 appears to indicate that, as far truth-conditional content goes, the first conjunct of (14)b_1 and the second conjunct of (15)b_1 are equivalent to 'the bookcase in the living room is completely empty' and 'all the neighbours are asleep', respectively. Now, this raises an obvious question: if this is true, then why is it that 'the bookcase in the living room is empty' and 'the neighbours are asleep' support imprecise interpretation, while 'the bookcase in the living room is completely empty' and 'all the neighbours are asleep' do not? *All* must do something to *the townspeople*, and *completely* must do something to *empty*… but, what, exactly? The answer to this question will have to wait until the next chapter.

On the basis of the observations made above, imprecision as a phenomenon may be characterised in terms of the following three features:

(i) Context-sensitivity: Imprecise interpretation is permitted in certain contexts but not in others.

(ii) Overt slack regulation: Modification of an imprecision-friendly construction by a slack regulator has the effect of preventing the sentence that contains the said construction from being interpreted imprecisely.

(iii) Strict behavior in contradiction test (SBCT): S is a sentence that admits either a precise or imprecise interpretation (e.g. 'the bookcase in the living room is empty'); S_{imp} is S's imprecise paraphrase (i.e. 'the bookcase in the living room is more or less empty') and S_{str} is S's strict/precise paraphrase (i.e. 'the bookcase in the living room completely empty'); SBCT can then be stated as follows: if 'S_{str} and not P' is a contradiction, then 'S and not P' is also a contradiction (irrespective of the context in which it is uttered).

This characterisation will be (slightly) revised in Chapter IV; for current purposes, it will suffice.

From now onwards, the term 'imprecision' should be understood as referring to a specific phenomenon that affects certain sentences in certain contexts, a phenomenon that can be identified by looking at whether the sentence under examination possesses features (i), (ii) and (iii).

3.2 Numerals and geometrical terms

Sentences containing numerals, such as (16), just like 'the bookcase in the living room is empty' and 'the townspeople are asleep', can also be interpreted imprecisely: indeed, (16) possesses features (i), (ii) and (iii).

(16) 100 visitors were inside the building.

Let's first establish that (16) possesses (i):

Context A: The second floor of an old museum collapsed. The box office attendant is asked to check the precise number of visitors that were inside the building at the time that the accident occurred; she replies: '**100 visitors were inside the building**'.

Context B: A broken aristocrat decides to open his house to the public in return for a small fee; one of the aristocrat's sons asks his father whether the initiative is paying off; the aristocrat replies: '**100 visitors were inside the building** (today)'.

In Context A, (16) is interpreted precisely: it cannot be used to describe a situation in which there were 99 (or 101) visitors inside the museum at the time that the second floor collapsed; in Context B, conversely, (16) is interpreted imprecisely: the sentence can be used (and not be judged false) if there were 99 (or 101) visitors inside the aristocrat's house.

Let's now establish (ii):

(17) *Exactly* 100 visitors were inside the building.

The addition of *exactly*, quite clearly, blocks imprecision: indeed, although (16) supports imprecise interpretation in Context B, (17) does not.

Finally, let's establish (iii):

(18) a. #100 visitors were inside the building, but there wasn't an even number of them.
b. Approximately 100 visitors were inside the building, but there wasn't an even number of them.

(18)a is a contradiction, irrespective of the context in which it is uttered; that is, there is not context in which (18)a can be interpreted as meaning (18)b.

The same results are obtained with sentences containing geometrical terms, such as (19).

(19) The figure on the screen is a square.

(19) can be interpreted precisely or imprecisely depending on the context (I let the reader to think of possible contexts and convince themselves that this is true) [feature (i)]; (20), which incorporates the modifier *perfect*, cannot be interpreted imprecisely [feature (ii)]; (21) is a contradiction (no matter the context in which it is uttered) [feature (iii)].

(20) The figure on the screen is a *perfect* square.

(21) # The figure on the screen is a square, but it is not a plane figure with four equal straight sides and four right angles.

3.3 Imprecision versus vagueness

Vague terms lack sharp boundaries; consider, for example, the term *young*: there does not seem to be a fact of the matter as to at which point a young person stops being *young* and begins to be *not young*; there are, of course, clear cases of youngness and clear cases of non-youngness: however, it does not seem possible to identify one point at which '*x* is *young*' flips from truth to falsity (a condition that persists even if the adjective is combined with a PP that overtly specifies the relevant comparison class: indeed, *young for a professional football player,* in the same way as *young*, lacks sharp boundaries). In addition, vague terms admit so-called 'borderline cases': each of these predicates (e.g. *young*, *heap*, *expensive*, *city*, etc.) is such that there are things to which one is not sure whether or not the predicate applies, irrespective of how much knowledge one has about the thing. Vagueness is often discussed as being a lexical phenomenon (e.g. the extension of *young* is vague); however, sentences that contain vague terms are also vague: the sentence 'his cousin is *young*', by virtue of having a vague constituent, is vague (i.e. it lacks a sharp truth-falsity boundary and admits borderline cases).

Imprecision, it must be noted, is not the same phenomenon as vagueness, a point that has been argued for extensively in the literature (see, for example, Pinkal 1995; Lasersohn 1999; Kennedy 2007). Perhaps the most fundamental feature of imprecision is that it comes and goes with the pragmatic wind: in certain contexts, certain sentences are interpreted imprecisely, but in other they are not—cf. the interpretation of sentence (16) in Context A vis-à-vis Context B, for example. By contrast, it does not seem possible to cook up a context in which the non-sharp boundaries of a vague sentence are rendered sharp and its borderline cases are eliminated (cf. Keefe 2007). Consider (22), for example: if in her desire worlds Charlotte lives in Cambridge, is (22) true or false? Not clear, precisely because Cambridge is (at least for me) a borderline case of *city* (i.e. a thing out there which it is unclear whether the predicate *city* applies to).[56] Now, can one come with a context in which (22)'s borderline cases *all* disappear, rendering (22) either clearly true or clearly false: the answer is *no*, there does not seem to be such a context.

(22) Charlotte wants to live in a city.

Of course, if one sets up a context in which *city* is stipulated to mean *urban agglomeration with at least 1,000,002 people*, then the term will no longer be vague and, as a result, (22) will not be vague either (at least not vague by virtue of having *city* as a constituent). However, that would amount to changing the meaning of the word *city* by equating it with the meaning of a non-vague measure phrase (Kennedy

[56] If, for the reader, Cambridge (UK) is not a borderline case of city, then replace Cambridge by another city (one that, according to their intuitions, would count as a borderline case of city).

2007). The expression *natural context* (from Pinkal 1995, 1996) is meant to exclude contexts such as those.

Another argument against identifying imprecision with the general phenomenon of vagueness (one that, to my knowledge, has not been made, at least not explicitly) is this. Let's assume that imprecision and vagueness are indeed the same thing. 'His cousin is *young*' lacks a sharp semantic boundary (or has vague truth-conditions) by virtue of having 'young', a vague term, as a lexical constituent; likewise, 'the bookcase is *empty*' lacks a sharp semantic boundary by virtue of having 'empty', a (by assumption) vague lexical constituent. On this account, if the latter sentence supports imprecise interpretation, it is because it has vague truth-conditions. (Such an account, of course, would not explain why 'my cousin is *young*' is vague irrespective of the context in which it is uttered, unlike 'the bookcase is *empty*'; for the sake of the argument, let's leave this issue aside). As has been noted, sentences that can be interpreted imprecisely, such as 'the bookcase is *empty*', exhibit 'strict behavior in contradiction test' (SBCT); but, if 'the bookcase is *empty*' had vague truth-conditions, then it should not exhibit SBCT (that is, (23) below should not be a contradiction):

(23) # The bookcase is empty, but there are two books in the first shelf.

In order to predict the contradiction in (23), it seems to me, one has no choice other than saying that 'the bookcase is *empty*' is true only if there are no objects in the bookcase; but, if that is the case, then nothing is vague about the truth-conditions of 'the bookcase is *empty*'. In sharp contrast to 'the bookcase is *empty*', 'his cousin is *young*' never exhibits SBCT, not even in a context that 'artificially' forces a non-vague interpretation of *young*, as illustrated below.

(24) <u>Context</u>: In order to count as young, one must have less than 12 years old.
His cousin is young, but she is 13.

The contrast between (23) and (24) strongly suggests that 'the bookcase is *empty*' has non-vague truth-conditions, as opposed to 'his cousin is *young*', whose meaning is obviously vague. The thesis that has is that 'the bookcase is *empty*' supports imprecise interpretation *because* it is vague (in the same way as 'his cousin is *young*' is) must be wrong: what the data seem to be telling us is that 'the bookcase is *empty*' supports imprecision despite having perfectly delineated truth-conditions.

To argue that imprecision and vagueness are different phenomena, a different approach can be taken: instead of showing that central features of imprecision (such as context-dependence or SBCT) are not features of vagueness, it can be shown that a central feature of vagueness is not a feature of imprecision. As is well-known, vague terms/sentences, precisely because they lack sharp boundaries, can be used to induce the sorites paradox, for many the empirical landmark of vagueness; consider, for example, (25) below.

(25) a. If John's height is190cm, then 'John is tall' is true; if John's height is 150cm, then 'John is tall' is false.
b. For any height *h*, if John's height is *h* and 'John is tall' is true, then 'John is tall' is also true at *h*-1cm.
c. 'John is tall' is true if John's height is 150cm.

Though (25)a and (25)b both seem intuitively true, they lead to the paradoxical conclusion in (25)c. Let's now consider (26).

(26) a. If exactly 100 visitors were inside the building, then '100 visitors were inside the building' is true; if no visitors were inside the building, then '100 visitors were inside the building' is false.
b. ?? For any number *n* of visitors, if there were *n* visitors inside the building and '100 visitors were inside the building' is true, then '100 visitors were inside the building' is also true at $n - 1$.
c. '100 visitors were inside the building' is true if exactly 0 visitors were inside the building.

In (26) the paradox does not go through, and it does not because (26)b cannot be true: if $n = 100$, then '100 visitors were inside the building' is true; however, if $n = 99$, then '100 visitors were inside the building' is, as a matter of fact, false. The sorites paradox, then, may be invoked to distinguish imprecision from vagueness, a point already made in Kennedy (2007).

To sum up, imprecision and vagueness cannot be the same thing; imprecision is a condition that affects certain sentences only in certain contexts; vagueness, on the other hand, seems to be a semantic feature of a great deal of lexical items which percolates up to the level of sentences and is largely independent from contextual factors. I do not think this characterisation is entirely accurate but, for our immediate purposes, it will suffice; in Chapter VI, I will revisit the vagueness/imprecision conundrum and present a slightly more complex picture.

3.4 Lasersohn's (1999) Pragmatic Halos

Though Lasersohn (1999) made important observations, his account of imprecision is wrong. Lasersohn represents imprecision at the predicate level; on his account, each expression in the language is assigned a denotation (relative to a model), which is used to calculate the truth-conditions of sentences, as well as a set of objects of the same logical type as the denotation (relative to a context). Each object in this set, which Lasersohn calls the *halo*, differs from the denotation only in some respect that is pragmatically ignorable in context. The denotation itself is included in the halo.

Lasersohn's system is (or at least attempts to be) compositional: if a complex expression consists of several predicates, each with its own halo, all those halos will 'add up' to produce the halo for the complex expression. Slack regulators, within this framework, serve to shrink or tighten the pragmatic halo of the expressions they combine with. A distinctive ingredient of Lasersohn's (1999) account is its watered-down version of the maxim of Quality: in uttering a given sentence S, the speaker does not commit to the truth of S, but to the truth of S being somewhere in the halo.

This account works in cases such as (1)a—given as (27) below—in which the observed (imprecise) interpretation is logically weaker than the utterance's literal meaning.

(27) The bookcase in the living room is empty.

As discussed, there are contexts in which (27) can be interpreted imprecisely; in such contexts, on Lasersohn's account, the interpretation of (27) would consist of a pragmatic halo that would include the utterance's literal meaning as well as other propositions. Simplifying for the sake of exposition, let's take the VP *is empty* to have $\{x : x$ has 0 books$\}$ as denotation, and its halo to be something like $\{\{x : x$ has 0 books$\}$, $\{x : x$ has exactly 1 book$\}$, $\{x : x$ has exactly 2 books$\}$,$\{x : x$ has exactly 3 books$\}\}$; further, let's assume that the NP 'the bookcase in the living room' has no halo. The halo of the complex expression—namely, of (27)—would then be the set of propositions {*the bookcase in the living room has 0 books, the bookcase in the living room has exactly 1 book, the bookcase in the living room has exactly 2 books, the bookcase in the living room has exactly 3 books}*. Given this halo, (27) could be truly uttered to describe a situation in which there are exactly 3 books in the bookcase: indeed, there is one proposition in (27)'s halo that is true of that situation, namely, *the bookcase in the living room has exactly 3 books*.

Lasersohn's account breaks, however, as soon as (27) is embedded under negation—as noted in Carter (2016, 2019); consider, for example, (28).

(28) Context: Charlotte and Gerard rented a flat for the weekend in the outskirts of London. The flat was advertised as 'having everything one could wish for' from a full-packed fridge to kitchenware, from TV to board games, etc. However, much to their surprise, the couple found that there was almost nothing to eat in the fridge, no board games, one plate only (and no cutlery). In this context, Gerard points out:

The bookcase in the living room isn't empty.

If in the bookcase there is only one book, (28) is likely to strike Charlotte as being false; indeed, in the stipulated context, (28) appears to mean something like *there is a reasonable number of books in the living room bookcase* (a proposition that is logically stronger than (28)'s literal meaning).

This interpretation, however, cannot be generated by Lasersohn's system. The halo of the VP *is empty* must include its actual denotation, that is, $\{x : x$ has 0 books$\}$: the halo of (28), therefore, is bound to include the proposition *the bookcase in the living room doesn't have 0 books* (= *the bookcase in the living room has at least 1 book in it*). Lasersohn's maxim of Quality, then, permits (28) to be used as soon as there is 1 book in the bookcase, no matter the context: in other words, (28) is predicted to be always interpreted precisely (clearly a bad prediction. It is essential to the mechanics of Lasersohn's account that the literal meaning is included in the halo; as a result, as Hoek (2018) observes, the account is structurally incapable of predicting global strengthening.

It should also be noted that that it is not clear whether Lasersohn's system accounts for SBCT, a defining feature of the phenomenon of imprecision. Consider, for example, the sentences in (29).

(29) a. # The theatre was empty, and only two people were sitting on the front.
b. # The townspeople were asleep, but some of them were awake.
c. # Sweden has 10 million inhabitants, but it doesn't have an even number of inhabitants.

Lasersohn's account enables speakers to use a sentence even if its literal meaning is a false: the only requirement is that truth lies somewhere in the halo. The sentences in (29) are all literally false (because they are contradictions); however, given how the system is set up, nothing prevents (29)b, for example, to be uttered in a context in which (29)b's pragmatic halo includes one or more contingent propositions. Lasersohn is aware of this issue, and makes the following proposal:

> One could felicitously assert the second clause [(29)b] only if there were some pragmatic relevance to the fact that some of the townspeople were awake—that is, if there were some pragmatic relevance to the distinction between the set of townspeople as a whole and the set of all the townspeople but the awake ones. But in that case, this distinction would not be ignorable in context, so the latter set would not be in the halo of the townspeople. Hence the first clause of [(29)b] will never be close enough to true for its context if the context is one in which the second clause is assertable' (Lasersohn 1999: 531).

Whether such move would work depends on a number of factors, in particular: (i) the precise formulation of the notion of relevance that is used in the model, and (ii) whether the pragmatic considerations that regulate the size of the halo operate at the level of lexical items, at the level of embedded clauses or, alternatively, at the level of whole sentences. In the absence of an explicit formal model, Lasersohn's proposal cannot be appropriately assessed.[57]

4 SUMMARY

In this chapter, I have done a number of things. First, I have argued that neither imprecision nor hyperbole can be made sense of in terms of LB (a hypothesised lexical modulation procedure that involves weakening the root meaning of the expression). Second, I have provided an empirical characterisation of imprecision, a characterisation that suggests that the mechanism behind it is pragmatic in nature and operates on the meaning of whole utterances. Third, I have argued, and I am not the first one to do so, that imprecision and a vagueness are different phenomena. Finally, I have discussed Lasersohn's influential account of imprecision as well as Carter's (2016, 2019) (fatal) argument against it.

In the next chapter, I will introduce Križ's (2015) account of imprecision; this account was originally developed to account for the fact that unquantified plural sentences such as 'the townspeople are asleep' can, in some contexts, be interpreted imprecisely. This account, although not without problems, is rich in theoretical insights and, as will be discussed, provides a principled explanation for SBCT and the phenomenon of slack regulation.

[57] Lasersohn may have at his disposal an alternative account of SBCT (one that he does not himself consider): he could say that a contingent proposition never differs from a logical contradiction in pragmatically ignorable ways. On this account, if the literal meaning of the utterance is a contradiction, the halo of the utterance will not contain any contingent proposition.

Chapter IV
(NON-)MAXIMALITY REVISITED

1 OVERVIEW

Imprecision appears to have two legs: a pragmatic and a semantic one. Indeed, certain contexts allow for a sentence to be interpreted imprecisely, but others, as discussed in the previous chapter, do not. Likewise, not all sentences can be interpreted imprecisely: for example, 'the townspeople are asleep' and 'the bookcase is empty' can, but 'all the townspeople are asleep' and 'the bookcase is completely empty' cannot. Križ's (2015) major theoretical insight was to identify a connection between imprecise interpretation and a semantic property of certain constructions/sentences, known in the linguistics literature as *homogeneity*. For a sentence S to be interpreted imprecisely, S needs to be uttered in a special kind of context (more on this to come) but, also, S has to be homogenous.

2 HOMOGENEITY

2.1 The phenomenon

The application of a predicate to a plurality (as denoted by a definite plural) creates an extension gap, an observation that dates back to Fodor (1970). Let me illustrate with (1).

(1) Context: There are 10 books on the table; pointing at them, George utters…

a. John wrote the books.
True iff John wrote all the books.

b. John didn't write the books.
True iff John wrote none of the books.

It is clear that (1)a is true of a situation (or world) in which John wrote all the books; it also clear that (1)b, its negation, is true of a situation in which John did not write any of the books: but what about situations in which John wrote some but not all books? (1)a and (1)b, as it happens, are neither true nor false of such situations. For example, take a situation in which John wrote half of the books: (1)a is not true (cf. 'John wrote some of the books', which is clearly true in that situation), but it is not false either (cf. 'John wrote all the books', which is clearly false in that situation). Following standard practice, I will call the set of worlds where (1)a is true the *positive extension* (or the truth-conditions) of (1)a (which I will often represent as $[\![(1)a]\!]^{+}$); the set of situations where (1)a is false the *negative extension* (or the falsity-conditions) of (1)a (also, $[\![(1)a]\!]^{-}$); and the remainder set of situations where (1)a is neither true nor false the *extension gap* of (1)a (or $[\![(1)a]\!]^{\#}$).

Definite plurals, it should be noted, are 'hidden' in a number of constructions, such as possessive or demonstrative plural NPs, as illustrated below.

(2) a. John wrote those books. (= John wrote <u>the books</u> over there.)
b. John read her books. (= John read <u>the books</u> that she wrote.)

(2)a and (2)b, naturally, also exhibit an extension gap.

Sentences such as (1)a-b and (2)a-b are said to have the *homogeneity* property: the name homogeneity is meant to reflect the fact that the predicate (in this case, the derived unary predicate *John wrote*) is neither true nor false of a plurality (in this case, *the books*) if it is true of some of the parts of the plurality and false of others, that is, if the plurality of books is not *homogeneous* with respect to the property of having been written by John.

Not all predicates, when applied to a plurality, induce gappiness, however; for example, *numerous* and *few in number* (also *heavy* and *light* in their collective readings) do not. Consider (3), for example.

(3) The students are numerous.

If (3) were to be homogenous, that would mean that if (3) is true, *numerous* cannot be false of any subgroup of the children, which is of course absurd. In the light of cases such as (3), Križ (2015, 2019) has argued that homogeneity should be characterised as a property of lexical predicates, rather than as a property of definite plural noun phrases. Križ's reasoning is as follows: if homogeneity was somehow rooted in the meaning of the definite plural—a thesis that has been defended in, for example, Breheny (2005) and Magri (2014)—the existence of non-homogenous sentences such as (3) would be hard to make sense of. A competing view, put forward in Bar-Lev (2018), has it that it is a logical operator that introduces homogeneity. In this chapter, I will refrain from endorsing any particular view on the question of why, and at which, level homogeneity originates, and will limit myself to speak of homogeneity as a property of sentences: homogenous sentences are those that, when negated, exhibit the sort of gap that the sentences in (1) and (2) exhibit.

The addition of *all* to a definite plural NP has the effect of removing homogeneity, an observation that dates back to Löbner (2000). This is illustrated in (4) below.

(4) a. John wrote all the books.
True iff John wrote all of the books.

b. John didn't write all the books.
True iff John didn't write all of the books.

It should be noted that all quantifiers have this effect: indeed, none of the sentences in (5) have an extension gap. *All* happens to be special, as Križ (2015) observes, in that the removal of homogeneity appears to be its sole semantic contribution.

(5) a. John wrote some of the books.
True iff John wrote at least 1 book.

b. John didn't write some of the books.
True iff John didn't write any book.

c. John wrote most of the books.
True iff John wrote most of the books.

d. John didn't write most of the books.
True iff John didn't write most of the books.

In the previous chapter, I characterise *all* as a slack regulator (a particle responsible for disabling the potential for imprecise interpretation). Here, I am presenting a (slightly) different analysis: *all* is being characterised as a homogeneity remover. Soon it will be argued that *all*, as a result of removing homogeneity, also disables the potential for imprecise interpretation.

Homogeneity, it should be noted, can be observed in a number of domains, and not just in the individual domain. Habitual sentences, for example, which arguably involve reference to plural events, also exhibit the homogeneity property, which can be appropriately removed by the introduction of a quantifier such as *always*, as illustrated in (6) below.

(6) a. John reads the newspaper on Saturdays.
b. John doesn't read the newspaper on Saturdays.
c. John *always* reads the newspaper on Saturdays.
d. John doesn't *always* read the newspaper on Saturdays.

If John reads the newspaper on some (but not all) Saturdays, (6)a is neither true nor false (the same can be said of (6)b, its negation). However, in such a situation, (6)c is false whereas (6)d is true.

It is also worth pointing out that homogeneity effects are not tied to pluralities *per se*: they appear to arise whenever a property is ascribed to a mereologically complex object—another important observation due to Löbner (2000). Consider the following example (from Križ 2015):

(7) The book is intelligently written.

Books are mereologically complex objects (i.e. they have parts: chapters, paragraphs, sentences, etc.). Faced with a book that contains some excellent chapters and some mediocre ones, (7) seems neither true nor false: in other words, such a situation falls into the extension gaps of (7). Just like with unquantified plural sentences and habitual sentences, (7)'s homogeneity property can be removed, as illustrated in (8).

(8) The book is intelligently written *throughout*.

Indeed, the application of *throughout*, an adverbial quasi-quantifier, makes homogeneity disappear: if the book in question happens to contain some excellent chapters and some mediocre ones, then (8) is false (whereas (7), as noted, is neither true nor false).

2.2 The nature of homogeneity

There are at least two other phenomena in natural language for which truth-value gaps have been invoked: vagueness and presupposition. For example, if x is a borderline case of mountain (neither clearly a mountain nor clearly not a mountain), then, according to some theories (e.g. supervaluationism; see Fine (1975) and Keefe (2000), among others) the sentence 'x is a mountain' is neither true nor false: x falls into the extension gap of the predicate *mountain*. Likewise, the sentence 'the king of France is bold', according to the standard semantic account of presupposition, is neither true nor false if evaluated relative to a world in which there is no king of France: the world that we live in, for example, falls into the extension gap of this sentence.

In what follows, it will be shown that homogeneity is, as a phenomenon, unrelated to both vagueness and presupposition.

2.2.1 Homogeneity and presuppositions

Homogeneity has often been called a presupposition (see, for example, Löbner 2000 and Gajewski 2005). Claiming that homogeneity is a presupposition amounts to stipulating that (1)a and (1)b, given below as (9)a and (9)b, presuppose (9)c:

(9) a. John wrote the books.
b. John didn't write the books.
c. Either John wrote all the books, or he wrote none of them.

Let's assume that *the books* just means *all the books*. If this is the case, then (9)a is true if John wrote all the books (in this case, the presupposition is trivially satisfied); it is false, if John didn't write all the books and the presupposition, (9)c, is satisfied (i.e. it is false if John wrote none of the books); and it is undefined if the presupposition is not satisfied (i.e. if John wrote at least 1 but not all the books). Alternatively, it could be stipulated that *the books* means *some of the books*; if this is the case, (9)a is true if John wrote some of the books and (9)c is satisfied, false if John didn't write any of the books (in this case, (9)c is trivially satisfied), and undefined whenever the presupposition is not satisfied.

Homogeneity is not a presupposition, however; below, I will consider three empirical observations that set these two phenomena apart.

2.2.1.1 Homogeneity violations and the 'wait a minute!' test (Spector 2013)

Spector (2013) points out that homogeneity violations cannot be objected to in the same way as presupposition failures. Consider, for example, (10) and (11).

(10) A: Does John know that Mary either bought all the jewels or none of them?
B: Wait a minute! I didn't know she can't have bought just some of them.

(11) A: Did Mary buy the jewels?
B: # Wait a minute! I didn't know that she can't have bought just some of them.

(10)A, due to the presence of *know*, triggers the presupposition that Mary either bought all the jewels or none of them. Presuppositions can be objected with *Wait a minute!*, as famously noted in von Fintel (2004): the fact that (10)B is licit illustrates this point. If (11)A were to trigger the same presupposition as (10)A, then it would be rather mysterious why (11)A cannot be objected with (11)B.

2.2.1.2 Projection (Križ 2015)

Presuppositions, as is well-known, project from the antecedent of a conditional; for example, (12) presupposes that John is dead.

(12) If Mary knew that John is dead, she would not be coming for dinner tonight.

If homogeneity were the presupposition that either every or no member of the plurality fulfills the predicate, then one would expect (13) to presuppose that either all the subjects or none of them are asleep. This is clearly not the case, however.

(13) If the subjects were asleep, the study could start.

2.2.1.3 Weeell… (Križ 2015)

An additional reason to distinguish presupposition failure from a homogeneity violation is that, whereas the former admits a hesitant *weeell* as a reply, the latter does not. For example, in (14), B has the option of uttering *weeell*, as if expressing doubt regarding the appropriateness of what has been asserted.

(14) Context: There are 10 books on the table. John has written exactly five of them.

A: John wrote the books on table.
B: Weeell…

By contrast, in the face of presupposition failure, a hesitant *weeell* is odd; indeed, B's reply in (15) does not seem at all appropriate.

(15) Context: Adam has never smoked.

A: Adam has stopped smoking.
B: # Weeell…

2.2.2 Homogeneity and vagueness

It is not immediately obvious how to make sense of the idea that homogeneity is some kind of vagueness: for one thing, 'the numbers which are evenly divisible by 2 are even' is homogenous yet none of its constituents are vague, as far as I can tell. To make the link, one would need to claim that plural predication is itself vague (see, for example, Scha 1981) and, further, that the undefinedness of (16)b-c relative to the situation in (16)a is of the same kind as the undefinedness of (17)b-c relative to the situation in (17)a.

(16) *Borderline case*

a. Situation: Paul is 1.77m.

b. Paul is tall (for a British adult).
c. Paul isn't tall (for a British adult).

(17) *Homogeneity violation*

a. Situation: Half of the books on the table are red, and half are blue.

b. The books on the table are red.
c. The books on the table aren't red.

In what follows, I will show that this link cannot be made: homogeneity violations and borderline cases dramatically come apart in linguistic tests, which indicates that these are different phenomena. The link, it seems, is that both in the case of homogeneity and in the case of vagueness, truth-value gaps are a useful tool to make formal sense of what is going on: however, *what is going on* is different in each case.

2.2.2.1 The Sorites paradox (Križ 2015)

One of the distinctive features of vague predicates (such as *rich*, *tall* or *heap*) is that they can be used to trigger (fallacious) soritical reasoning; consider (18), for example.

(18) a. If John has \$ 5 million, then 'John is *rich*' is true.
b. If John has no money whatsoever, then 'John is *rich*' is false.
c. If John has an \$ *n* amount of money and, as a result of having this amount, 'John is *rich*' is true, then, if John has *n* – \$ 1, then 'John is *rich*' is also true.

(18)c, the conditional premise, enables ones to infer that if John has not money at all, then he is rich, which contradicts (18)b.

If homogeneity were to be just a symptom of vagueness, then a sentence of the form 'the kids got excellent grades' (which is homogenous) should lead to the sorites paradox. I do not think it does, however—neither does Križ (2015). Consider (19), for example.

(19) a. If all of the all the kids got excellent grades, then 'the kids got excellent grades' is true.
b. If none of the kids got excellent grades, then 'the kids got excellent grades' is false.
c. If *n* kids got good grades makes 'the kids got excellent grades' true, then *n* – 1 kids getting excellent grades also makes the kids got excellent grades' true.

One has no issue in accepting the conditional premise for 'John is rich' in (18)c; (19)c, by contrast, is hard; indeed, imagine that there are 10 kids, 9 got excellent grades and 1 got poor grades… is 'the kids got excellent grades' true? I don't think it is, frankly. There may be some special cases, as Križ (2015) remarks, in which (19)c could be accepted: these special cases, as will be discussed, are contexts in which, for current purposes, it is irrelevant whether one of the children did not do well. In Križ's (2015) theory, and in the reformulation of his theory that I will present, such imprecise or non-maximal interpretations are not a direct manifestation of semantic vagueness but the outcome of a *pragmatic* mechanism (to which the trivalent semantics of homogeneous plural predication is conceptually prior).

2.2.2.2 Negated disjunctions (Križ 2015)

As has been observed in Chapter II, § 4.2, sentences of the form $\ulcorner\neg(\mathrm{F}a \vee \neg(\mathrm{F}a))\urcorner$, despite being logical contradictions, are not perceived as contradictory when F is a vague predicate, and appear to entail that *a* is a borderline case of F. For example,

(20) Bakewell is neither a city nor not a city.
⇒ Bakewell is neither definitely a city nor definitely not a city.

(21) John is neither rich nor not rich.
⇒ John is neither definitely rich nor definitely not rich.

If a sentence such as 'John read the books' was vague (due to the presence of plural predication), then one would expect this sentence to escape contradictoriness when embedded in a negated disjunction, just like (20) and (21) do. This is not the case, however, as illustrated in (22).

(22) # John neither read the books nor didn't he read the books.
⇏ John neither read all of the book nor did he read none of the books.

2.2.2.3 Disagreement

One thing that borderline cases lend themselves to is disagreement; for example, I may be of the opinion that Cambridge is a city, whereas you may insist that Cambridge is not a city, that it is just a town. Likewise, I may be inclined to think that someone who owns a nice house (no mortgage) and drives a BMW is a rich person, whereas you may dispute that, and claim that, in order to count as rich, one has to own many properties (one is not enough).

With this in mind, let's consider (23).

(23) Context: There is a pile of 20 books on the table; 17 are clearly blue, and 3 are clearly red.
John: [*pointing at the book pile*] The books on this pile are blue.
Jane: Really? These books are not blue.
John: Well, I think many people would be willing to say that these books are blue.

The exchange above is so weird that I can only imagine two lunatics having it; compare with the perfectly natural (24).

(24) Context: John is visiting Cambridge for the first time.
Tour guide: Cambridge is a city.
John: Really? This isn't a city.
Tour guide: Well, I think many people would be willing to say that Cambridge is a city.

Quite clearly, whereas disagreement about what counts (and does not count) as a borderline case of a vague predicate is perfectly possible, disagreement about whether a situation in the extension gap of a homogeneous sentence is in fact in the extension gap of that sentence does not seem possible. Vague terms, as discussed in the previous chapter, lack sharp semantic boundaries and, as a result of this, it is often the case that people disagree about (or feel compelled to negotiate) where these boundaries should be drawn. Plural sentences, by contrast, have perfectly sharp boundaries (with respect to whether the predicate is true, false, or neither true nor false of the plurality): just because unquantified plural sentences exhibit an extension gap does not mean that these sentences lack sharp semantic boundaries.

2.2.2.4 *A word of caution: definitely*

Križ observes that homogeneity and vagueness appear to have a common feature (which may be used to argue that these, after all, are one and the same phenomenon); in his own words, '*definitely* functions very much like *all* in that it modifies a vague adjective so that it becomes simply false of a borderline case' (Križ 2015: 41).

On this point, I beg to differ with Križ: *definitely* and *all* do not function alike. To begin with, the sole semantic purpose of *all*, as far as I can tell, is to remove homogeneity; if a sentence is not homogenous, then the application of *all* to a plural definite NP results in infelicity.

(25) a. The students are numerous.
b. # *All* the students are numerous.

Definitely, by contrast, can felicitously modify a non-vague predicate, as illustrated in (26) below, which indicates that its core function cannot be that of removing vagueness.

(26) a. 11 is a prime number.
b. 11 is *definitely* a prime number.

Definitely, it seems to me, works as a general indicator of epistemic certainty, like *certainly*, *clearly*, or *without a doubt*. These epistemic adverbs, by virtue of signaling certainty, have the effect of pushing the borderline cases of a vague predicate into its negative extension; however, as (26)b indicates, this appears to be a side-effect of these adverbs rather than their *raison d'être*.

The fact that epistemic adverbs have vagueness-trimming effects, it must be stressed, does not mean that these adverbs *remove* vagueness; as it has long been noted, 'the vagueness of a vague predicate is ineradicable' (Dummett 1978/1959: 182). This is another difference between homogeneity and vagueness: the former can effectively be removed; the latter can, via the application of an epistemic adverb such as *definitely*, be eradicated at some level (for example, the borderline cases of *young* may be pushed into the predicate's negative extension) but it cannot be eradicated at all levels (indeed, the complex predicate *definitely young* is vague: there isn't a fact of the matter as to at which point a person stops being *definitely young* and begins to be *not definitely young*). Another way of making this point is to say that, whereas higher-order vagueness is a thing,[58] there does not appear to be such a thing as higher-order homogeneity.

Last but not least, it is worth reflecting on the fact that epistemic adverbs such as *definitely* do not seem to interact at all with the homogeneity kind of gap. Consider, for example, (27).

[58] For discussion, see, for example, Sainsbury (1991),Williamson (1999) and Keefe (2000).

(27) a. John read the books.
b. John *definitely* read the books.
c. John read the books, *definitely*.

These sentences are true if John read all the books, false if he read none, and undefined otherwise: that is, the indication of epistemic certainty, which, as discussed, interacts with the phenomenon of vagueness, does absolutely nothing to homogeneity. This indicates, quite strongly I think, that these two phenomena should be kept apart.

2.2.2.5 Homogeneity, vagueness, and imprecision

In Chapter III, § 3.3, I argued that vagueness and imprecision are different phenomena; in this section, I did something different: I argued that homogeneity and vagueness are different phenomena. In the sections to come, it will be shown that there is a link between homogeneity and imprecise interpretation—a link which, to my knowledge, was identified for the first time in Križ (2015): homogeneity (a semantic phenomenon of its own) is a *pre-requisite* for a sentence to be interpreted imprecisely.

2.2.3 Final remarks

The introduction of a third truth-value is a popular strategy to deal with presupposition and vagueness. Following Križ (2015), I have used this strategy to model yet another phenomenon, namely, *homogeneity*. If presupposition failure, borderline cases, and homogeneity violations are all dealt with via the introduction of an additional truth-value, a trivalent semantics will not suffice: in addition to truth and falsity, three truth-values are going to be needed. In the present work, since I am only interested in formalising homogeneity, I will keep the semantics trivalent; notice, however, that a more complex system with more than three truth-values could also be implemented (cf. Spector 2016).

3 KRIŽ'S THEORY OF (NON-)MAXIMALITY

One of Križ's (2015) major insights was to offer a precise diagnosis of why (28)a can be interpreted imprecisely while (28)b cannot: according to his account, only homogenous expression can be interpreted imprecisely; as a result, by removing homogeneity via the application of *all*, the potential for imprecise interpretation is disabled.

(28) a. The townspeople are asleep.
b. All the townspeople are asleep.

But why does a sentence need to have the homogeneity property in order to be interpreted imprecisely? This question is answered by Križ's (2015) theory of (non-)maximality, which is introduced in the next section. The key idea is this: 'an undefined sentence, but not a false sentence, can be used when the situation described is, for current purposes, equivalent to a situation where the sentence is literally true.' (Križ 2017: 22).

3.1 Preliminaries

As discussed in the previous chapter, imprecision 'comes and goes with the pragmatic wind'[59]: focusing on sentences with definite plural NPs, the general fact seems to be that, in contexts in which it does not matter whether *all Xs are blah* or whether only some of them are, *Xs are blah* is interpreted imprecisely. By contrast, in contexts in which it does matter whether *all Xs are blah* or whether only some of them are, then *Xs are blah* is interpreted precisely. What one wants, therefore, is a theory of (*im*)*precision*: a theory that explains why, in certain contexts, certain sentences are interpreted imprecisely and why, in other contexts, the same sentences are interpreted precisely. This is exactly what Križ (2015) offers to us.

In the more technical literature, to refer to imprecise and precise interpretations of plural sentences, the terms *non-maximality* and *maximality* are used (cf. Brisson 1998; Malamud 2012; Križ 2015; Bar-Lev 2018). From now onwards, I will switch to the technical terminology. The term (*non-*)*maximality* will be reserved to refer to the general phenomenon of (im)precision.

3.2 The theory

Since the emergence of QUD theory (Roberts [1996] 2012), it has become a mainstream to model what is and isn't relevant at a given point in discourse is via set partitions (cf. Groenendijk and Stokhof 1984). Križ (2015) follows this trend. A QUD (or Question Under Discussion) is a partition on a set of possible worlds; when modelling human linguistic interactions, this set is typically identified with the Stalnakerian Context Set, namely, the set of worlds that are compatible with all the propositions in the common ground ('the common ground', in this framework, stands for the set of propositions that are mutually believed by the conversation participants).[60] Križ (2015), it should be noted, uses the term *issue*, as opposed to QUD, perhaps because it's less theoretically loaded than the latter. Following Roberts ([1996] 2012), Križ (2015) stipulates that assertions are always interpreted relative to an issue,

[59] The expression comes from Lewis (1986).
[60] If the common ground is the set of mutually believed propositions $\{p_1, p_2, p_3, \ldots p_n\}$, then the Context Set is $\bigcap\{p_1, p_2, p_3, \ldots p_n\}$.

and that (fully) resolving an issue amounts to determining in which of the issue cells the actual world is.

Križ's theory of (non-)maximality is built on three pillars: (i) the notion of *sufficient truth*; (ii) a revised (weaker) maxim of Quality (which depends on the notion of *sufficient truth*); and (iii) a constraint that imposes a particular alignment between issues and sentences. The consequence of (i) and (ii) is that, given a declarative sentence S and an issue *I*, S will be interpreted as 'something is the case that, relative to *I*, is equivalent to S being true'; (iii), in turn, rules out the possibility of S being used if false. To illustrate the workings of the theory, let's consider the following (highly idealised) scenario, which comes from Križ (2015).

<u>*Sue's talk*</u>. Sue has given a talk and we are interested in how Sue's talk was received: the issue in question, let's call it *I*, can be described as follows: a cell i_1 (where Sue's talk counts as well-received), a cell i_2 where the reception is mixed, and finally i_3 (where it was ill-received). Further, let's assume that the reception of Sue's talk is judged based on the facial expressions of the professors in the audience, and that the Context Set consists of four worlds: w_1 (where all the professors smiled), w_2 (where all the professors smiled except for Smith), w_3 (where only half of the professors smiled), and w_4 (where none of the professors smiled): w_1 and w_2 are contained in i_1 (the positive reception cell), w_3 is contained in i_2, whereas w_4 is contained in i_3. Fig. 1 illustrates the issue just described.

Figure 1. A depiction of issue *I*.

i_1: positive reception	w_1: all the professors smiled w_2: all the professors smiled except for Smith
i_2: mixed reception	w_3: only half of the professors smiled
i_3: negative reception	w_4: none of the professors smiled

The reason why both w_1 and w_2 are in i_1 is because, whether Smith smiled or not, is irrelevant as to whether Sue's talk was well-received (by stipulation): Smith hardly ever smiles, so even if he did not smile, it seems safe to conclude that the talk went well (and that Smith was just being Smith). Smith, therefore, is an irrelevant exception: whether he smiled or not does not influence one's judgement as to in which cell w_2 should be placed (as Križ likes to put it, Smith is irrelevant *for current purposes*).

Let's now imagine that (29) is uttered as an answer to *I* (that is, with the intent to resolve this issue).

(29) The professors smiled.

(29), quite clearly, can be uttered to describe a state of affairs compatible with w_2 if the underlying issue is whether or not Sue's talk was well-received. (29), however, is not true in w_2: it is true if all the professors smiled, false if none of them smiled, and undefined otherwise (cf. §2). Križ's intuition is this: (29) can be uttered in w_2 despite not being true in w_2 because, relative to issue I, w_2 and w_1 are equivalent (and (29) is true in w_1). Let's now introduce Križ's notion of *sufficient truth*.

(I) Sufficient truth (Križ 2015)
We write $\approx_I$ for the equivalence relation that holds of two worlds u, v iff u and v are in the same cell of I. A sentence S is *sufficiently true* in world w with respect to an issue I iff there is some world w' such that $w' \in [\![S]\!]^+$ (S is literally true in w') and $w \approx_I w'$.

Thus, though (29) is not true in w_2 (semantically speaking), with respect to issue I, is sufficiently true in w_2, because w_2 is in the same cell as w_1, and (29) is literally true in w_1.[61]

The next ingredient of Križ's theory is a change in the maxim of Quality.[62] In Grice (1975), the maxim of Quality instructs speakers to make only true statements (to the best of the speaker's epistemic abilities). However, under the QUD conception of communication, the main goal of assertion is to resolve a certain issue (that is, to learn in which cell of the partition the actual world is). In the light of this, Križ proposes to weaken the maxim of Quality so that it requires only that the speaker should say sentences that are sufficiently true for current purposes.

(II) Križ's (weak) maxim of Quality
A speaker may say only sentences which, as far as she knows, are *sufficiently true*.

Križ's weakened maxim of Quality has the following effect: the information that is communicated by a sentence is not its literal truth-conditions, but rather the union of every issue cell that the sentence's positive extension intersects. Why so? Let's consider (29) again. Even if the speaker knows that the actual world is w_2 and not w_1, the (revised) maxim of Quality still permits her to utter (29), because w_2 'lives' in the same cell as w_1 and, in w_1, the sentence is literally true. The hearer, in turn, can infer no more than the actual world is somewhere in i_1: indeed, according to Križ's weakened maxim of Quality,

[61] Križ uses the expression 'true enough', as opposed to 'sufficiently true'; I have opted for the latter denomination, because 'true enough' is typically used in the literature to describe how a sentence *feels* when it is not strictly speaking true, yet it is nonetheless usable (for example, Lewis 1979). Throughout the thesis, I'm using the expression 'true enough' in the latter sense.

[62] This move is reminiscent of Lasersohn's (1999) Pragmatic Halos (see Chapter III).

(29) does not need to be true, just sufficiently true;[63] thus, with respect to issue *I*, (29) communicates the proposition i_1, as illustrated in Fig. 2.[64]

Figure 2. (29)'s communicated content relative to issue *I*.

i_1: positive reception	w_1
	w_2 ↓
i_2: mixed reception	w_3
i_3: negative reception	w_4

Note that, on this account, a sentence such as (29) will be interpreted maximally if the issue under consideration is such that w_2 is in a different cell from w_1; this would be an issue in which Smith not smiling would be relevant for current purposes. For example, let's imagine that, if all the professors smile, including Smith, then Sue gets a special award; further, let's assume that the issue under consideration is the one represented in Fig. 3 (call this issue *II*).

Figure 3. A depiction of issue *II*.

ii_1: Special award	w_1 : all the professors smiled
ii_2: No award	w_2: all the professors smiled except for Smith w_3: only half of the professors smiled w_4: none of the professors smiled

If (29) were to address *II*, then (29) would be interpreted maximally: indeed, $[\![(29)]\!]^+$ only intersects ii_1, and ii_1 entails that every single professor smiled and Sue obtained the special award. Thus, with respect to issue *I*, (29) is predicted to be interpreted as meaning i_1 (a non-maximal interpretation) and, with respect to *II*, it is predicted to be interpreted as meaning ii_1 (a maximal interpretation). That these predictions are correct can be corroborated by comparing how the sentence 'the professors smiled' is interpreted in (30) vis-à-vis (31).

(30) Jane believes (and thus wants to communicate to John) that either w_1 or w_2 are the case.

Was Sue's talk well-received? [issue *I*]
Jane: a. The professors smiled. So yes, it was.
b. All the professors, or all the professors except for Smith, smiled. So yes, it was.
c. # All the professors smiled. So yes, it was.

(31) Jane believes (and thus want to communicate to John) that w_1 is the case.

Did Sue get the special award? [issue *II*]

[63] Notice that, since the equivalence relation $\approx_I$ is reflexive (by virtue of being an equivalence relation), being true entails being sufficiently true; the opposite, however does not hold: a sentence can be sufficiently true without being true.
[64] Križ refers to i_1 as a *Quality* implicature, i.e.: (29), relative to issue *I*, triggers the implicature that i_1.

Jane: a. The professors smiled. So yes, she did.
b. # All the professors, or all the professors except for Smith, smiled. ~~So yes, she did.~~
c. All the professors smiled. So yes, she did.

Let's first focus on example (30): (30)a and (30)b seem to communicate the same message, which indicates that (30)a is being interpreted non-maximally. (30)c, by contrast, commits Jane to a stronger proposition than the one that she is in a position to assert: indeed, Jane does not know whether w_1 or w_2 is the case, and (30)c communicates that w_1 is the case. In short, in order to communicate what she wants to communicate, Jane can choose either (30)a and (30)b, but cannot choose (30)c, the explicitly maximal sentence: this indicates that (30)a is not interpreted maximally.

The situation is reversed in (31); indeed, in this case, (31)c, the explicitly maximal sentence, and (31)a appear to communicate the same message: that Sue got the special award. (31)b, by contrast, does not communicate this message: (31)b is true in w_1 , but also in w_2, and the latter world is one in which Sue did not get the special award; as a result, (31)b does not resolve issue *II*. In fact, in (31)b, Jane cannot add 'So yes, she did (get the special award)' at the end, because this sentence, given the stipulated common ground, does not follow from 'all the professors, or all the professors except for Smith, smiled'.

Križ's system, therefore, captures the fact that (non-)maximality is *issue* or *partition*-sensitive. Changing the issue (without removing or adding worlds from/to the Context Set) can have the effect of forcing either a maximal or non-maximal interpretation. Indeed, *I* and *II* are different partitions on the same set of worlds, that is, $\{w_1, w_2, w_3, w_4\}$: the only thing that changes from *I* to *II* is how this set of worlds is partitioned.

The last ingredient of Križ's theory is the *Addressing* constraint. Križ notes that his account, unless constrained, predicts that a sentence such as (32) can be uttered in w_2 as an answer to issue *I* (cf. Fig. 1): indeed, given such an issue, (32) complies with the (revised) maxim of Quality, as w_2 is in the same cell as w_1, where (32) is literally true.

(32) All the professors smiled.

This seems like a bad prediction: (32) is just false in w_2 and cannot be used to describe a state of affairs compatible with w_2. To avoid this problem, Križ (2015) posits the following constraint, which he calls *Addressing*.

(III) Addressing

A sentence S may be used to address an issue *I* only if there is no cell $i \in I$ such that i overlaps with both the positive and the negative extension of S, i.e. S is true in some worlds in i and false in others.

What this constraint does is to force a particular alignment between sentences and issues: in order for a sentence S to address an issue, the worlds in any given issue cell cannot fall on different sides of S's true-false boundary. From *Addressing*, as Križ (2015: 80) remarks,

> it follows that no sentence can be used when it is literally false. For assume that the actual world w is in the question cell i_1, and [sentence] S is false in w. Then either S is not true in any world in i_1 and therefore eliminates i_1 as a possible answer to the current issue, in which case it is obviously inappropriate because the right answer shouldn't be eliminated. Or alternatively, S is true in some of the worlds in i_1, but then it is false in others in the same cell (including w). This means, [by (III)], that S cannot be used to address the issue at hand.

A direct consequence of this is that a sentence without an extension gap (such as an *all*-sentence) cannot be used non-maximally: such a sentence can be either literally true or literally false, and *Addressing* prevents speakers from using it when false; a sentence without an extension gap, therefore, can only be used when it is literally true. By contrast, a homogenous sentence such as (29) can be used when it is true, of course, but also when it is undefined (provided that it is sufficiently true).

If applied to our examples, (III) entails that (32) cannot be used to address issue I, because (32)'s positive and negative extension, as illustrated in Fig. 4, both intersect i_1.

Figure 4. Sentence (32) against issue I: a violation of *Addressing*.

i_1: positive reception	w_1 (**T**), w_2 (**F**)
i_2: mixed reception	w_3 (**F**)
i_3: negative reception	w_4 (**F**)

In Križ's framework, in fact, due to the *Addressing* constraint, an *all*-sentence such as (32) can only be used to address an issue in which every exception matters; there is something counterintuitive about this, because it seems to me that both *all the professors smiled* and *the professors smiled* are admissible ways of addressing I. In § 4.1, in fact, I will claim that the *Addressing* constraint is far too restrictive and could be dispensed with if Križ's (weak) maxim of Quality was refined.

3.3 Non-maximality and negation

Križ's (2015) account, it should be noted, makes the right predictions for negated plural sentences; consider the contrast between (33)a and (33)b.

(33) a. The students didn't do well in the exam.
b. None of the student did well in the exam.

It is easy to imagine a context in which (33)a is true even if one or two of the students did fine; (33)b, by contrast, does not tolerate exceptions. This is expected: *none* is the negative analogue of *all*: it removes homogeneity and, as a result, blocks (non-)maximality. (33)a can be interpreted non-maximally for the same reason as its positive counterpart can: it has the homogeneity property and is therefore susceptible to (non-)maximality.

3.4 The defining features of (im)precision

In Chapter III, § 3.1, it was shown that (im)precision—which I now refer to as (non-)maximality—is associated with three features: (i) *context-sensitivity*; (ii) *susceptibility to overt slack regulation*; and (iii) *strict behaviour in contradiction test*. Let's see how Križ's theory helps us to make sense of all this. I shall first comment on (iii), then (ii), and, finally, (i).

3.4.1 Strict behaviour in contradiction test

Sentences such as (34)b and (34)c are contradictions; even in a context in which (34)a could interpreted non-maximally, the first conjunct in (34)b and (34)c still behaves as if it meant the same as 'all the professors smiled'.

(34) a. The professors smiled.
b. # The professors smiled, but one of them did not.
c. # The professors smiled, but not all of them did.

In Križ's theory, the contradictoriness of (34)b and (34)c is predicted. As discussed, (34)a is true if all the professors smiled, false if none of them smiled, and undefined otherwise; as a result of this, (34)a is true in exactly the same worlds in which 'all the professors smiled' is true. From this is follows that (34)b and (34)c can never be true: as soon as the first conjunct is true, then the second will be false, and thus the conjunction will be false.[65] Now, Križ's algorithm is global, that is to say, it takes the meaning of full sentences as input. For a sentence to be interpreted non-maximally, the sentence must be *sufficiently true*: (34)b and (34)c, however, can never be true and, as a result, there is no way that they will ever be *sufficiently true*. The upshot is that non-maximality cannot be computed on such sentences.

Now, as soon as a sentence such as (34)b or (34)c is split into two different sentences, the situation changes. Consider, for example, (35).

[65] To model the meaning of complex sentences such as (34)b and (34)c, Križ relies on Strong Kleene connectives; in the case of conjunction, $p \wedge q$ is true if both p and q are true; undefined if p is true and q is undefined (or vice versa); and false if p is false and q is undefined/false/true (or vice versa).

(35) A: How was Sue's talk received?
B: Very well! *The professors smiled*. Of course, one of them didn't… Smith… but he never does…

(36) A: How was Sue's talk received?
B: # Very well! *All the professors smiled*. Of course, one of them didn't… Smith… but he never does…[66]

It is clear that 'the professors smiled' in (35)B is interpreted non-maximally—if it was not, (35)B would entail that all the professors smiled and, as a result, it would give rise to a cross-sentential contradiction, in the same way as (36)B does. This is expected on Križ's account: (35)B's first sentence can be evaluated relative to A's question, whereas 'one of them didn't', the next sentence, can be evaluated relative another issue (e.g. *did all the professors smile?*). ☐

☐

It cannot be denied that *of course* plays some role in 'rescuing' (35)B; however, its role should not be overestimated. For one thing, *of course* is also present in (36)B, and (36)B is a cross-sentential contradiction. Furthermore, the addition of *of course* in the middle of (34)c does not make the sentence any better, as shown in (37) below.

(37) # The professors smiled, but *of course* one of them didn't.

If (35)B is good, it is because it consists of two different sentences, each of which can be evaluated independently. What *of course* appears to do, as Križ (2015) observes, is to somehow signal that a shift to a more fine-grained issue must be performed; this, if true, would explain why *of course* feels like a helpful hand in this sort of case (i.e. it signals that a change of issue is in order).

In short, on Križ (2015) account, the contradictoriness of (34)b-c is predicted on the basis of 'the professors smiled' being homogenous and (non-)maximality being a global mechanism that takes full sentences as input.

3.4.2 From slack regulators to homogeneity removers

In Chapter III, it was noted that *all* has the function of blocking imprecision in definite plural NPs (or forcing a precise interpretation)—following Lasersohn (1999), I referred to it as a *slack regulator*. Križ's (2015) theory, however, leads to a slightly more complex characterisation of *all*: yes, it does prevent sentences containing definite plurals from being interpreted non-maximally (or imprecisely),

[66] This intuition is slightly confounded by the fact that *all the professors smiled* can be interpreted hyperbolically; however, prefacing the sentence with *as a matter of fact* seems to control for this (i.e. *as a matter of fact*, *all the professors smiled*).

but it does so by virtue of removing homogeneity: since homogeneity is a pre-requisite for non-maximality, the addition of *all*, naturally, prevents non-maximal interpretation.

3.4.3 Context-sensitivity

Križ's account explains why a sentence (with the homogeneity property) is interpreted non-maximally in certain contexts and maximally in others: according to this account, a sentence S is interpreted as 'something is the case that, relative to the issue under consideration, is equivalent to S being true'. Thus, if the issue changes (and issues are features of context), then the result of the interpretation process is liable to change too.

4 KRIŽ (2015) PUT TO TEST

Križ's (2015) theory is very beautiful and provides a great deal of insight into the nature of the phenomenon of (im)precision; and it works, but it works as long as highly idealised issues are considered (for example, in Sue's talk case, *all the professors (or all except Smith) smiled* entails the 'positive reception' cell (cell i_1, see Fig. 1) and vice versa. As will be shown below, as soon as more realistic issues are considered, Križ's (2015) account runs into trouble.

4.1 The overt question and the current issue

Consider the following case:

(38) Context: Gerard has ten kids; he and his kids are at a party.

a. *Friend*: Hey, are you staying at the party for a bit longer?
b. *Gerard*: The kids are tired… (so I'm not sure).

Note that (38)b is interpreted non-maximally: indeed, (38)b would be readily judged true even if only some of Gerard's children were tired. Let's identify the question in (38)a with the issue; since (38)a is a binary question, the issue, which I will designate as Q, is bound to have two cells, i.e.: q_1 (*staying at the party*) and q_2 (*not staying at the party*). Which worlds feature in q_1 and q_2 will depend on the set of worlds that one is partitioning; if Q is a partition on the Context Set, then it seems completely reasonable to assume that in q_1 there is a world w_1, in which Gerard is staying at the party despite all his kids being really tired (in this world, he cannot come back to his house until 1am anyway, because he lost his keys and has to wait for his wife to get to the house first—and she will be getting there at 1am), as well as a world w_2, in which Gerard is staying at the party because his kids are not tired at all and are having a good time. Further, it is entirely reasonable to suppose that, in q_2, there is a world w_3 in which Gerard does not stay at the party (in this world, all his children are tired), a world w_4 in which Gerard does not stay at the party (in this world, 9/10 of his children are tired), as well as a world w_5 in which Gerard does not stay at the party because he has to get up early the next day (although in this world, none of his kids are tired). This issue is illustrated in Fig. 5.

Figure 5. A depiction of issue *Q*.

q_1: Staying at the party	w_1: Gerard is staying at the party & all his kids being really tired. w_2: Gerard is staying at the party & none of his kids are tired.
q_2: Not staying at the party	w_3: Gerard is not staying at the party & all his kids are tired. w_4: Gerard is not staying at the party & 9/10 of his kids are tired. w_5: Gerard is not staying at the party & none of his kids are tired.

Notice that the positive extension of (38)b intersects q_1 ($w_1 \in [\![(38)b]\!]^+$) as well q_2 ($w_3 \in [\![(38)b]\!]^+$); in addition, the negative extension of (38)b intersects both q_1 ($w_2 \in [\![(38)b]\!]^-$) and q_2 ($w_5 \in [\![(38)b]\!]^-$). According to Križ's *Addressing*, a sentence S may be used to address an issue *I* only if there is no cell $i \in I$ such that *i* overlaps with both the positive and the negative extension of S; in the case under consideration, *Addressing* is violated, as q_1 as well as q_2 overlap with both the positive and the negative extension of (38)b. (38)b, therefore, cannot be used to address *Q*. This is the first obstacle one runs into; *Q* seems like a perfectly reasonable issue: if (38)b is not licit as an answer to *Q* (because it violates *Addressing*), then, given that non-maximality is being computed, should one conclude that (38)b is addressing some other issue? And, if so, *which* issue?

Addressing, it seems to me, extraordinarily limits the kind of issues that a sentence can address. But is this constraint really needed? I am not sure that it is; one could drop *Addressing,* and, furthermore, add to Križ's maxim of Quality an extra condition: *A speaker may say only sentences which, as far as she knows, are sufficiently true*, ***and, moreover, not false***. The net effect of this addition would be that of removing from the Context Set the worlds in which the sentence is false. The revised system now prevents *all*-sentences from being interpreted non-maximally (i.e. upon uttering an *all*-sentence, the *not-all* worlds will be removed), and it does so without the need of having *Addressing*.

Now, even if *Addressing* were to be dropped and the maxim of Quality tweaked, the resulting algorithm would still fail to assign (38)b a non-maximal interpretation (if *Q* is the issue). Let's see why; worlds w_2 and w_5, as soon as (38)b is uttered, will be removed from the Context Set (due to the revised maxim of Quality just proposed): however, the positive extension of (38)b intersects both q_1 and q_2. On Križ's account, the information that is communicated by a sentence is not its literal truth-conditions, but rather the union of all issue cells that the sentence's positive extension intersects. Since (38)b intersects every *Q*-cell, it cannot be used to address it (if it could, then the communicated content of (38)b would be weakened to triviality). This is the second problem that the account faces; this problem, it should be noted, is not exclusive to Križ's theory: rather, Križ's theory inherits this problem from QUD theory (Roberts [1996] 2012), something which I will discuss in more detail in the next chapter.[67]

The *Addressing*-related problem, as mentioned, can be avoided by revising Križ's maxim of Quality along the lines suggested. However, if one sticks to the original account, then (38)b, due to *Addressing*, cannot be used to address *Q*: thus, (38)b has to be addressing a different issue (not the one associated with the overt question). Regarding the second problem (the fact that (38)b intersects every *Q*-cell), there is no much room for maneuvering: because (38)b intersects every cell of *Q*, (38)b cannot be used to address *Q* : thus, once again, there is no choice but to say that (38)b is addressing a different issue (not the one associated with the overt question).[68] Thus, to handle (38)b, the *current* issue (the issue against which the non-maximal interpretation is calculated) cannot be identified with the issue induced by the overt question.

This raises the question: which issue should one take (38)b to be addressing? If what one wants is to derive a non-maximal interpretation, one would need, for example, the issue *are any of the kids tired?*, which consists of two cells, i.e., (i) *at least 1 of the kids is tired* and (ii) *none of the kids are tired*. Then (38)b would only intersect (i), *Addressing* will be satisfied, and (38)b would be predicted to mean (i). Of course, there is still the question of how to derive the fact that (38)b, if uttered against (38)a, triggers the inference that Gerard is not likely to stay at the party; however, how this inference is derived, Križ could argue, is a separate issue, and not something that a theory of (non-)maximality should be concerned with.

But something isn't right here: in order to derive the non-maximal interpretation, one needs to stipulate that the issue *are any of the kids tired?* (or some other issue that returns the desired result) is the issue that the sentence is in fact addressing, an issue which isn't induced by the overt question. But, stipulating that the issue *are any of the kids tired?* is the current issue amounts to stipulating that (38)b is to be interpreted non-maximally (and this is what needs to be explained, and not stipulated). Indeed, nothing stops one from stipulating that (38)b is addressing the issue *are all of the kids tired?*; given such an

[67] In fact, there's a sense in which Križ did the best job possible given the structural limitations imposed by the QUD framework. More on this to come.

[68] Notice that it would not work to claim that (38)b, if uttered against (38)a, presupposes that *if all the kids are tired, then the speaker is not staying at the party*. If such a presupposition were to be triggered and accommodated (and note that there is no evidence whatsoever that something of this sort is taking place), then the positive extension of (38)b would not intersect every *Q*-cell: it would only intersect q_2. But, then, (38)b's communicated content would be predicted to be q_2, and this is a bad prediction: (38)b is not interpreted as meaning *the speaker is not staying at the party* (it is interpreted, ignoring the probabilistic component of the inference, as meaning *the speaker is not staying in the party as a result of some of his kids being tired*)—cf. § 4.2.2. I write 'ignoring the probabilistic component of the inference' because, as a matter of fact, the inference to the effect that the speaker is not staying in the party is probabilistic in nature: indeed, (38)b triggers the inference that Gerard is more likely than not to leave the party, and not that he *is* leaving the party. Such an inference, for obvious reasons, cannot be derived in this framework.

issue, *Addressing* will be satisfied, and (38)b would be interpreted maximally, and, as a result, entail that *all the kids are tired*. That this is not the interpretation that (38)b has is obvious; what is not obvious is why (38)b is interpreted non-maximally, as opposed to maximally. Križ's (2015) theory, as far as I can tell, is silent on this issue.

Križ (2015), it must be noted, is aware that, for his theory to work, the current issue can rarely be identified with the overt question. Križ regards this as an unfortunate fact of the relationship between overt questions and issues; in his own words, 'it is not possible to just ask a single overt question and pretend that the meaning of this question just *is* the current issue, not even as a hypothesis in a thought experiment' (2015: 86). If Križ were to be right about this, then his theory could hardly be put to test, something that he himself acknowledges: '[this] weakens the predictive power of the theory, since one cannot set up a context so precisely as to fully constrain the current issue and put a prediction to the test […]. One is therefore forced to restrict oneself to considerations of plausibility' (2015: 86).

The overt question does not always set up the current issue; that much is true. I can think of at least two configurations in which this is obviously the case; let's first consider (39) below.

(39) a. *John*: Does Bill like Mary?
b. *Paul*: MARY likes BILL [captialisation signals focal stress]

(39)b, quite clearly, does not answer (39)a: John is asking about whether Bill likes Mary, and not about the reverse. However, (39)b does not feel like an irrelevant answer: the reason for this seems to be that (39)b is not understood as an answer to (39)a, but rather to the question *who likes whom?* (39), therefore, is an instance in which the overt question cannot be identified with the issue that the answer is in fact addressing. Paul manages to elude addressing the overt question by means of a specific linguistic device, that is, focal marking. It is well-known that prosodic focus often acts as a cue for issue identification, and that speakers, when it is unclear what the question is or when the wrong question has been posited, can signal which question they are actually addressing by means of this device (cf. Roberts [1996] 2012; Büring 2003).[69]

(40), below, is another instance in which the overt question cannot be identified with the issue.

(40) Context: John and Paul are friends; John is in love with Jane, Paul's sister (Paul is very much aware of this). John organised a party in his house and invited both Paul and Jane. Paul confirmed his attendance, but Jane did not. Jane's silence is making John anxious. John decides to ring Paul, his friend, in an attempt to get some information from him.

[69] Thanks to Manuel Križ for discussion on this point.

a. *John*: Hi Paul, how are you?
b. *Paul*: I'm great, thanks
c. *John*: May I ask how is your sister?
d. *Paul*: She's not coming to the party John.

(40)d, quite clearly, is not addressing (40)c, but the question that John really cares about, namely, *is Jane coming to the party?* This is an instance of mind-reading, something that close friends are usually quite good at.

This granted, I suspect that, in most cases, it is possible to identify the current issue with the overt question. Let's return to (38), given as (41) below:

(41) Context: Gerard has ten kids; he and his kids are in a party...
a. *Friend*: Hey, are you staying at the party for a bit longer?
b. *Gerard*: The kids are tired.

(41)b, as far as I can tell, addresses the issue induced by (41)a and no other issue: indeed, the truth of (41)b increases the probability of the actual world being one in which Gerard is not staying at the party, as opposed one in which he stays. I find hard to believe that (41)b is uttered to address an issue other than the one induced by (41)a, or that (41)b addresses (41)a via some intermediate issue (an intermediate issue that has to be stipulated).

In the next section, I will show that, if Križ's (2015) theory, and the QUD framework more generally, are reformulated, it becomes possible to derive non-maximal interpretations on the basis of the sentence uttered and the overt question alone. This is good for at least two reasons: first, it is good because the systematic reliance on hypothetical issues is, from a theoretical standpoint, undesirable; in addition, it is good because it makes the theory testable.

A small digression is in order. It is well-known that, while it makes complete sense to treat issues as partitions of Logical Space (or the Context Set),[70] treating linguistic questions in an analogous way—that it, as denoting partitions—is problematic: indeed, linguistic questions show an asymmetry between positive and negative information, an asymmetry that partitions cannot capture (see, for example, Heim 1994; Klinedinst and Rothschild 2011; Spector and Egré 2015, among others). To circumvent this problem, a well-known approach is to represent linguistic questions as denoting non-mutually exclusive sets of propositions (so-called Hamblin sets).

[70] As Fox (2018: 404) beautifully puts it, entertaining an issue 'involves a concern with locating oneself in a space of possibilities, where certain distinctions matter and others don't. A partition is useful for describing such a concern.'

The point that I want to make is this: that linguistic questions and issues cannot be identified with the same formal object is not in tension the view that I am defending here: namely, that issues can be identified with the overt (linguistic) question. As pointed out by many, partitions can be derived from question denotations (if the latter are formalised as Hamblin sets) in a straightforward manner; to derive a partition given a Hamblin set H, one just needs to partition Logical Space to sets of possibilities that agree with each other on the truth-value of members of H—for a formal definition of this procedure, see, for example, Fox (2018).

4.2 Bad predictions

Insofar as Križ's (2015) account can be tested, it does make incorrect predictions in at least two cases, which Križ himself identifies and discusses.

4.2.1 Problematic case I

(42) a. *John*: Did some of the kids come?
b. *Paul*: The kids came.

(42)b, if uttered against (42)a, and if (42)a is identified with the issue, would be predicted to communicate the proposition *at least 1 of the kids came*. Let's see why. (42)a is a binary question and, as a result, the issue, let's call it *F,* is bound to have two cells f_1 (*at least 1 of the kids came*) and f_2 (*none of the kids came*). The positive extension of (42)b is only going to intersect f_1 while its negative extension will only intersect f_2 (*Addressing* is therefore satisfied). Since Križ's maxim of Quality enables the speaker to use (42)b if it is sufficiently true, the listener can infer no more than the actual world must be in f_1: as a result, (42)b ends up communicating f_1. This is, of course, a bad prediction: (42)b is not interpreted as meaning *some of the kids came*.

4.2.2 Problematic case II

(43) Context: In order to pass, Peter has to solve either all the math problems, or at least half of them and write an essay on mathematical Platonism.

a. Did Peter pass the exam?
b. He solved the math problems.
c. He passed the exam.

(43)b can only be interpreted as meaning that Peter passed the exam, and that he did so by means of solving all the math problems. This is not the interpretation that Križ's account predicts. The question is *did Peter pass the exam?*, which has two cells, cell *i* (*Peter passed the exam*) and cell *ii* (*Peter didn't pass the exam*). All the worlds where (43)b is literally true are, given the context stipulated, worlds where he passed, and all the worlds where it is false are worlds where he didn't pass. *Addressing* is therefore satisfied. Križ's algorithm generates the final interpretation by taking the union of all the cells that the sentence intersects: since (43)b only intersects cell *i*, the prediction is that, in the context stipulated, (43)b should be interpreted as meaning (43)c. This prediction is incorrect.

5 SUMMARY

In this chapter, I have done three things. First, I have introduced the semantic condition known as *homogeneity*; this condition has been shown to be unrelated both to vagueness and presupposition failure. Second, I have reviewed in quite some detail Križ's (2015) account of non-maximality and showed how it makes sense of the empirical reflexes of imprecise interpretation, namely, *context-sensitivity*, *overt slack regulation*, and *strict behaviour in contradiction test*. Finally, I have pointed out some limitations Križ's (2015) account.

In the next chapter, a probabilistic account of non-maximality will be put forward, one that builds upon Križ's (2015) theoretical insights.

Chapter V
PROBABILISTIC (NON-)MAXIMALITY

1 OVERVIEW

In the previous chapter, Križ's (2015) account of (non-)maximality was introduced and discussed at length; in addition, I pointed out some of its problems. In the present chapter, I shall develop a probabilistic account of (non-)maximality, one that builds upon Križ's main theoretical insights. This account, as shall be discussed, overcomes the limitations of the original framework.

2 ADDRESSING AN ISSUE

In order to reformulate Križ's (2015) account, I first need to some conceptual groundwork; in what follows, I will characterise the common ground in probabilistic terms, present a formal definition of *relevance*, and revisit some technical notions.

2.1 The classical notion of relevance

According to the 'classical' framework, a proposition p is relevant to a *QUD* or issue I (a partition of C, the Context Set), if p has the effect of 'eliminating' issue cells (or i-cells). This approach to modelling relevance has been popularised by Roberts (1996/2012), though its origins can be traced back to Groenendijk and Stokhof's (1984) seminal work on questions and answers. In this framework, a proposition p eliminates i-cells (and hence it is **relevant** to I) iff p is a partial answer to I: that is, iff $p \cap C \neq \emptyset$ and, furthermore, there is an $X \subset I$ such that $\{p \cap C\} \subseteq \bigcup X$ (such that p contextually entails[71] $\bigcup X$); p is a complete answer to I iff $p \cap C \neq \emptyset$ and, furthermore, it contextually entails one (and only one) i-cell (and thus eliminates all i-cells but one).[72] Note that a complete answer is a partial answer, but not vice versa. A proposition p is **irrelevant** if it intersects every i-cell or if it does not intersect any: in either case, p does not eliminate any i-cell (and hence it is not a partial answer). These possibilities are exemplified in Fig. 1 below.

[71] Given two propositions p and q, p contextually entails q iff $\{p \cap C\} \subseteq q$, where C is the Context Set.
[72] There is a more restrictive notion of (complete/partial) answer, that excludes over-informative propositions. Here I am using the more liberal notion, which corresponds to Groenendijk and Stokhof's (1990) notion of *giving* a (partial/complete) pragmatic answer.

Figure 1. Relevance and irrelevance: the ‘classical’ notions.

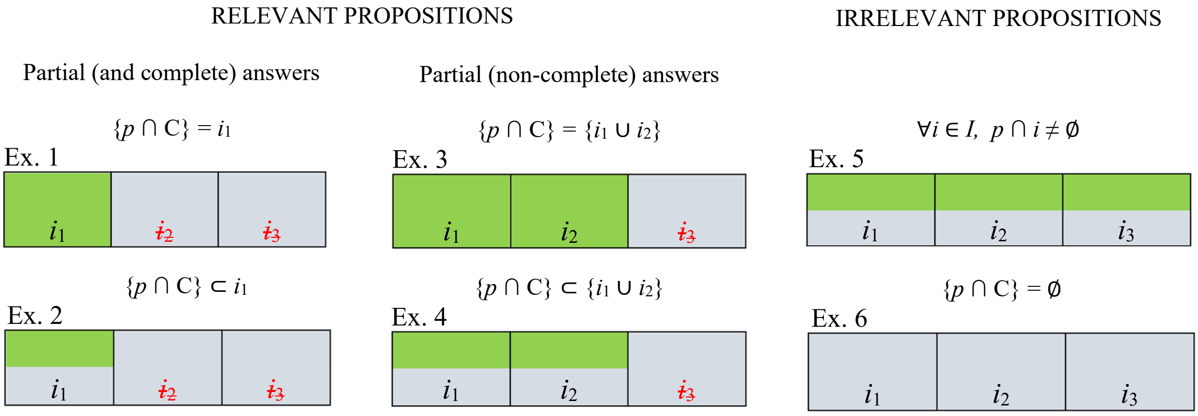

Let’s consider the following exchange:

(1) A: Is Mary coming to the party?
B: She is tired.

A’s question is of the yes/no kind and can thus be identified with a two-cell issue, let’s call it K: k_1 (*Mary is coming to the party*) and k_2 (*Mary is not coming to the party*). Given a non-idealised Context Set (i.e., a Context Set such that B’s utterance does not contextually entail either k_1 or k_2), B’s utterance, as per the definitions above, is bound to be irrelevant to A’s question. B’s answer, however, does not *feel* irrelevant: it does not eliminate k_1, that is true, but it nonetheless conveys relevant information, namely, that it is unlikely for the actual world to be in k_1.

Within the classical tradition, two strategies have been pursued to deal with cases such as (1): either B’s utterance does not address the issue induced by A’s question (but some other issue, which needs to be stipulated) or, as has been suggested in Roberts (1996/2012:12), B’s utterance presupposes something that, if accommodated into the common ground, would make the proposition expressed by B’s utterance a complete answer. The former strategy was discussed in the previous chapter; I have nothing else to add on the matter. Regarding Roberts’s (1996/2012) suggestion, I do not think it is correct. If a presupposition of the sort that Roberts has in mind were to be accommodated, then A would infer that Mary is not coming to the party. If I was A, however, I would not infer this: I would infer that Mary is (very) unlikely to come.

What seems to be going on here is this: it is common ground that, if someone is tired, then s/he is not likely to be going to parties; upon learning that Mary is tired, A infers that Mary is not likely to go the party (i.e. A infers that the actual world is much more likely to be in k_2 than in k_1). Because B’s contribution feels relevant, it is imperative to have a definition of relevance that accommodates cases in which the answer does not eliminate any issue cell yet provides an indication as to where in the partition the actual world is more likely to be. What is in fact needed is a gradient (probabilistic) notion of relevance.[73]

2.2 Towards a probabilistic notion of relevance

To characterise relevance in probabilistic terms, a probabilistic notion of common ground is required. Thus, instead of identifying the common ground with the Context Set (a set of possible worlds), I will model common knowledge as a probability measure function P that maps each subset of Logical Space (represented as Ω) to a probability score (i.e. a real number between 0 and 1). Issues or QUDs will no longer be thought of as partitions of the Context Set (no such notion exists in this this framework): issues are partitions of Ω or Logical Space.

The idea of characterising relevance in probabilistic terms is of course not new, at can be traced back at least to Carnap (1950). The basic intuition, succinctly stated in Büring (2003), is as follows: given an issue *I* and a prior distribution P_0, *p* is going to be relevant if updating P_0 with *p* shifts the probabilistic weights among *i*-cells. For example, in (1), and assuming a flat prior distribution (i.e. $P_0(k_1) = P_0(k_2) = 0.5$), updating P_0 with the positive extension of B's utterance generates a posterior P_1 such that $P_1(k_1) \neq P_0(k_1)$ and $P_1(k_2) \neq P_0(k_2)$ (e.g. $P_1(k_1) = 0.2$ and $P_1(k_2) = 0.8$).

It is worth noting that, unlike the classical notion of relevance, the probabilistic notion does not require a proposition to be conducive to cell elimination; to visualise this, consider example (2).[74]

(2) a. *John*: Is Mary coming to the party?
b. *Jane*: She's tired.
c. *Lucy*: But she really wants to come.

As before, let's identify (2)a with the issue, represented as *K*, which consist of two cells, i.e. k_1 (Mary is coming to the party) and k_2 (Mary is not coming to the party). Furthermore, let's assume a flat prior distribution, i.e.: $P_0(k_1) = P_0(k_2) = 0.5$. Jane's turn, (2)b, which feels relevant, increases the probability of the actual world being in k_2; however, Lucy's turn, (2)c, which also feels relevant, takes us back to a flat (or near-flat) distribution. Fig. 2 illustrates the sequence.

[73] This is not an original point; in fact, Simons et al. (2010: 316-7) have made the exact same point: '[Roberts' (1996)] definition of Relevance is overly restrictive and should be weakened at least to allow for discourse moves which merely raise or lower the probability of some answer to the QUD being correct. Consider for example the sequence: Q: "Is it going to rain?" A: "It's cloudy." A's utterance does not contextually entail an answer to the QUD (at least not in Pittsburgh, PA). Intuitively, it is relevant because it somewhat raises the probability of an affirmative answer to the QUD.'
[74] Thanks to Manuel Križ for the example!

Figure 2. *k*-cells as (2) unfolds.

(i) Is Mary coming?

Yes	**NO**
50%	50%

(2)b →

(ii) Is Mary coming?

Yes	**NO**
20%	80%

(2)c →

(iii) Is Mary coming?

Yes	**NO**
50%	50%

The transition from (i) to (ii) gets John closer to resolve the issue that he himself raised; indeed, on the basis of the information communicated by (2)b, John can make the informed guess that Mary is likely not to come. However, when moving from (ii) to (iii), John is brought back to where he was: indeed, at (iii), he is as uncertain as to whether Mary will be coming as he was at (i). If, in order for a proposition to be relevant, that proposition had to push one closer to some resolution or other, (2)c should not feel relevant: after learning that (2)c is the case, John is no longer able to make an informed guess as to whether or not Mary is likely to come. John may be closer to the truth at (iii) than at (ii), but he is also further way from being able to rule out one of the issue cells. (2)c, however, does feel relevant. The take-home message is this: information does not need to be resolution-conducive in order to be relevant. The notion of relevance here proposed is in conformity with this observation: given an issue *I* and a prior distribution P_0, *p* is going to be relevant if updating P_0 with *p* shifts the probabilistic weights among *i*-cells (irrespective of whether learning that *p* is conducive to cell elimination).

The pseudo-formal definition given above is a bit too simplistic, however. Consider, for example, (3):

(3) A: Is Mary coming to the party?
B: ? She's not tired.

(3)B, quite indisputably, raises the probability of the actual world being in k_1 (Mary is coming to the party): after all, not being tired, is a reason to come (as opposed to not come) to a party. However, (3)B, from a pragmatic point of view, is borderline infelicitous and does not feel relevant. It seems that, in order for relevance to kick in, updating P_0 with *p* should result in a significant shift of the probability distribution at the cell level. But how significant?

I am of the impression that, if there is something like a significance threshold for relevance, this threshold is going to depend on many factors. Let's imagine that Mary is one the most famous actresses in the UK, that John is a journalist working for a tabloid, and that Jane is Mary's personal assistant. John needs to find out whether Mary will be attending this party; if she is, he should try to be at the door of the designated venue with his photographic equipment (a picture of Mary may translate into a considerable cash bonus). Mary is however super secretive, and never lets anyone know about her movements. John is lucky to know Mary's personal assistant, Jane, and decides to call her to ask whether Mary will be going to the party. Jane would lose her job if she gave John a 'yes' or 'no' answer;

she likes John however and, instead of telling him ‘I have no idea’, utters (3)B. Notably, in this context, (3)B feels relevant. Indeed, if I were John, upon hearing (3)B, I would pack up my camera and head to wherever the party is being held. Thus, although (3)B would typically not count as a relevant answer to the question in (3)A, if the stakes happen to be high, and information scarce, it may.

2.3 p-relevance: a definition

In what follows, I will define relevance in probabilistic terms and, to distinguish it from the classical (non-probabilistic) notion, I will refer to it as p-relevance; the definition relies on the Kullback–Leibler (KL) divergence (Kullback and Leibler 1951), an information-theoretic measure of how one probability distribution is different from another. From now onwards, and for the ease of exposition, I will assume that Logical Space, represented as Ω, is finite.

(I) p-relevance of a proposition

Notation:

I is a partition on Ω: $I = \{i_1,\ldots,i_n\}$
i is partition cell: $i \in I$
w is a singleton: $|\, \mathrm{w} \,| = 1$
p is a proposition: $p \in \wp(\Omega)$
P is the common ground: a probability distribution over Ω
U^* and P^* are probability distributions over I

Definition:

For any $i \in I$,

$$\mathrm{P}^*(i) = \sum_{\mathrm{w} \subseteq i} \mathrm{P}(\mathrm{w})$$

$$\mathrm{U}^*(i) = \mathrm{P}^*(i \mid p)$$

Applied to our purposes, the Kullback-Leibler (KL) divergence looks as follows:

$$\mathrm{D_{KL}}(\mathrm{U}^* \| \mathrm{P}^*) = \sum_{i \in I} \mathrm{U}^*(i) \times \log\left(\frac{\mathrm{U}^*(i)}{\mathrm{P}^*(i)}\right)$$

Then, ***p* is p-relevant relative to partition *I* and a probability distribution P, if $\mathbf{D_{KL}}$ ($\mathbf{U^*} \,\|\, \mathbf{P^*}$) > $\boldsymbol{\alpha}$**, where α is a contextual threshold.

$\mathrm{D_{KL}}$ ($\mathrm{U}^* \,\|\, \mathrm{P}^*$) should be read as ‘the Kullback–Leibler divergence from P^* to U^*, where P^* is the prior distribution over *i*-cells and U^* the posterior distribution.’ If $\mathrm{D_{KL}}$ ($\mathrm{U}^* \,\|\, \mathrm{P}^*$) = 0, then, for any $i \in I$, $\mathrm{P}^*(i)$

$= U^*(i)$ (that is, the probabilistic weights among *i*-cells has not shifted from P^* to U^*). If $D_{KL}(U^* \parallel P^*) > 0$, then there must be some $i \in I$ such that $P^*(i) \neq U^*(i)$; thus, if $D_{KL}(U^* \parallel P^*) > 0$, *p* is p-relevant, as updating P with it has the net effect of shifting the probabilistic weights among *i*-cells.[75] The variable α is a contextual threshold, whose value (a non-negative number) is going to depend on how high the stakes are and how scarce information is. With the notion of p-relevance defined, a precise formulation of the maxim of Relevance can be given.

(II) Maxim of Relevance
A speaker should say sentences whose positive extension is p-relevant.

In this framework, propositions that are classically irrelevant (propositions that eliminate no issue cell) can be broken down into two kinds: those that are nonetheless p-relevant (i.e. $D_{KL}(U^* \parallel P^*) > \alpha$), and those that are p-<u>ir</u>relevant (i.e. $D_{KL}(U^* \parallel P^*) < \alpha$). The latter kind of propositions (but not the former) constitute violations of (II); (4)B below is an example of p-<u>ir</u>relevance.

(4) A: Is George in love again?
B: # It's raining in Toronto.

Thus, a constraint seems to be needed here, one that prevents the common ground from being updated with the truth-conditions of replies such as B's (i.e. with irrelevant, and hence worthless information). Such a constraint is given below.

(III) Ban on p-<u>ir</u>relevance
I is a partition on Ω: $I = \{i_1,\ldots,i_n\}$ P is a probability distribution S is a sentence S^+ is the positive extension of S: $S^+ \in \wp(\Omega)$ If S^+ is not p-relevant relative to I and P, then, as long as I remains the issue, S cannot be used to update P.

[75] The distance between P* and U* (and not between P and U) is being measured because, for our purposes, it does not matter whether other subsets of Ω other than the partition cells have gained or lost probability after updating P with *p*: the relevance of *p* is measured relative to a partition and, as a result, all that matters is the probabilities of the cells.

2.4 From p-irrelevance to p-relevance: another case of 'black magic'?[76]

Let's recall Grice's (1975) famous example:

(5) A: Smith doesn't seem to have a girlfriend these days.
B: He has been paying a lot of visits to NY recently.

A's utterance, although not an overt question, does bring to the fore the issue of whether or not Smith has a girlfriend. B's utterance, in turn, raises the probability of Smith having a girlfriend being the case. The question is why. It is common ground that, if a person is tired, she is much more likely to stay home than to go partying. Now, is it common ground that, if a man pays regular visits to NY, he is much more likely to have a girlfriend than if he does not pay regular visits to NY? I would not think that it is. On this point, Simons (2005: 330) observes:

> An additional assumption is required, one which explicitly links the issue of having girlfriends to the issue of travel to NY: perhaps, the proposition that a person who has a girlfriend somewhere travels there frequently; or that many people have long-distance relationships, and these involve frequent trips to the same place. If A can work out that B is making *this* supposition, then she can immediately see the relevance of B's response to her remark. Without it, relevance cannot be established.[77]

One should therefore keep in mind that not all propositions that fail to be p-relevant will be perceived as irrelevant in practice: it is possible that some intervening mechanism (e.g. the sort of assumption accommodation procedure that Simons appears to have in mind) is operative behind the scenes and tweaks the common ground appropriately so that so that p-relevance is ultimately secured. Further research would be needed to establish what the constraints on this mechanism are.

2.5 Answers and relevance implicatures

In the classical framework (see § 2.1), a proposition p is relevant iff p is a *partial answer* (which includes the special case of being a *complete answer*): if p is a partial answer, then at least 1 cell of the issue is going to be eliminated. In the revised (probabilistic) framework, relevance is not linked to cell elimination, as discussed; indeed, in the revised framework, partial answers constitute a special case of p-relevance: it holds when updating P with p has the effect of bring down to zero at least 1 cell of the

[76] Kratzer (1981: 311) refers to Lewis's (1979) 'Rule of Accommodation' as *black magic*: 'If the utterance of an expression requires a complement of a certain kind to be correct, and the context just before the utterance does not provide it, then ceteris paribus and within certain limits, a complement of the required kind comes into existence. This is black magic, but it works in many cases.'

[77] The assumption that seems to be needed is in fact this: *if John has been paying a lot of visits to NY recently, then he's much more likely than not to have a girlfriend*. Whether this statement can be thought of as denoting a proposition is very much an open issue.

issue. The question that I want to raise here is the following: how does one know that a proposition eliminates a cell and, as a result, constitutes a partial answer to a given issue?

In the classical framework, issues are partitions of C, the Context Set; issues, within the classical framework, could also be defined in non-pragmatic terms, that is, as partitions of Ω or Logical Space.[78] Let's be explicit about these two possible formulations of the notion of issue:

(6) Let Q be a question (an equivalence relation):[79]

a. $I_{Q/\Omega}$ is the partition on Ω induced by Q: $I_{Q/\Omega} = \{i_Q : i_Q \in I_{Q/\Omega}\}$
b. $I_{Q/C}$ is the partition on C induced by Q: $I_{Q/C} = \{i_Q \cap C : i_Q \in I_{Q/\Omega} \ \& \ i_Q \cap C \neq \emptyset\}$

The observation that I want to make (trivial but purposeful in the context of the discussion to come) is the following: provided that $p \cap C \neq \emptyset$, and that $| I_{Q/\Omega} | = | I_{Q/C} |$, then, if p is a partial answer to $I_{Q/\Omega}$ (i.e. if $p \neq \emptyset$ and, furthermore, there is an $X \subset I_{Q/\Omega}$ such that $p \subseteq \bigcup X$ [such that p <u>logically</u> entails $\bigcup X$]), then p will also be a partial answer to $I_{Q/C}$ (i.e. $p \cap C \neq \emptyset$ and, furthermore, there will be an $X \subset I_{Q/C}$ such that $p \cap C \subseteq \bigcup X$ [such that p <u>contextually</u> entails $\bigcup X$]).[80] Let's illustrate this with (7) below.

(7) a. Did Paul pass the exam?
b. Paul did pass the exam.

$[\![(7)b]\!]^+$ is a partial (and complete) answer to $I_{(7)a/\Omega}$; hence, $[\![(7)b]\!]^+$ will also be a partial (and complete) answer to $I_{(7)a/C}$ (irrespective of what the Context Set is, as long as $I_{(7)a/C}$ is defined[81]).

Now, what if p intersects every cell of $I_{Q/\Omega}$? Then, it will (typically) be hard for p to be a partial answer to $I_{Q/C}$: this is because contextual entailment, in the absence of logical entailment, is, in practice, difficult to achieve. To illustrate this, let's consider the following case.

(8) <u>Context:</u> If Paul solves half of the exercises, he will pass the exam.

a. Did Paul pass the exam? [Two cells: YES/NO]
b. Paul solved half of the exercises.

[78] In the revised (probabilistic) framework, issue are partitions of Ω; however, this does not mean that context does not constrain issue space: it does via P, the common ground. Thus, in the revised framework, the notion of contextual entailment is preserved: given two propositions p and q, p contextually entails q iff $P(q \mid p) = 1$. By contrast, in the classical framework, if issues are defined as partitions of Ω, then contextual information is effectively taken out of the picture (and, as a result, the notion of contextual entailment is lost).
[79] For Groenendijk and Stokhof's (1984), questions are equivalence relations on possible worlds (which correspond to partitions of possible worlds); if formalised as Hamblin sets, questions can also be used to induce partitions (see Chapter 4, § 4.1).
[80] Thanks to Manuel Križ for discussion on this.
[81] If $I_{(7)a/C}$ is a partition (i.e. it is a collection of at least two non-empty cells), then $p \cap C \neq \emptyset$ and $| I_{Q/\Omega} | = | I_{Q/C} |$.

$[\![(8)b]\!]^+$ is, quite clearly, not a partial answer to $I_{(8)a/\Omega}$, as it intersects every cell of $I_{(8)a/\Omega}$: put it differently, $[\![(8)b]\!]^+$ does not logically entail any of $I_{(8)a/\Omega}$'s cells. However, if the proposition if Paul solves half of the exercises, he will pass the exam is common ground, $[\![(8)b]\!]^+$ should be a partial answer to $I_{(8)a/C}$: it should contextually entail $I_{(8)a/C}$'s YES cell (and hence eliminate $I_{(8)a/C}$'s NO cell). Thus, whether $[\![(8)b]\!]^+$ is an answer (whether its acceptance leads to cell elimination), crucially depends on whether the proposition if Paul solves half of the exercises, he will pass the exam, just by virtue of being given as context, gets into the common ground.

To determine whether a given sentence denotes an answer to an issue, I propose to use the following test, which I shall call *the should test*. First, consider (9) below:

(9) <u>Context:</u> If Paul solves half of the exercises, he will pass the exam.

a. Did Paul pass the exam? [Two cells: YES/NO]
b. Paul did pass the exam.
c. # Paul did pass the exam. So, he should have passed.

The positive extension of (9)b, as discussed in (7), is a partial answer to $I_{(9)a/C}$ (it contextually entails $I_{(9)a/C}$'s YES cell): this follows from the fact that $[\![(9)b]\!]^+$ logically entails $I_{(9)a/\Omega}$'s YES cell. The oddness of (9)c is thus expected: 'So, he should have passed' is redundant: indeed, the first sentence of (9)c already settles the question.

Let's now consider (10); if $[\![(10)b]\!]^+$ were to contextually entail $I_{(10)a/C}$'s YES cell, then (10)c should be as odd as (9)c; (10)c, however, is perfectly felicitous. This indicates that $[\![(10)b]\!]^+$ does not contextually entail $I_{(10)a/C}$'s YES cell: the proposition *if Paul solves half of the exercises, he will pass the exam*, though provided in the context, does not make it into the common ground.

(10) <u>Context:</u> If Paul solves half of the exercises, he will pass the exam.

a. Did Paul pass the exam? [Two cells: YES/NO]
b. He solved half of the exercises.
c. He solved half of the exercises… So, he should have passed.

What seems to be going on here is that (10)b, given the stipulated context, has the net effect of (massively) increasing the probability of the actual world being in the YES cell: however, (10)b does not fully settle the question (i.e. it eliminates no cell), as exposed by *the should test.*

There are cases in which contextual entailment can be achieved despite the absence of logical entailment; consider (11) below, for example.

(11) a. Is Paul living in France? [Two cells: YES/NO]
b. He is living Paris.
c. # He is living Paris… So, he should be.

'X is living in Paris' does not logically entail that 'X lives in France'; consider, for example, a world in which Paris is a British dependency: the first sentence is true relative to that world, while the second is false. Now, 'X is living in Paris' must contextually entail that 'X is living in France': if it didn't, the oddness of (11)c would be rather mysterious. Contextual entailment, therefore, appears to be achievable in cases in which contextual information is deeply entrenched, such as the well-known fact that Paris is in France: however, the provision of written context, as in (10), is not enough to secure entailment (and, consequently, cell elimination).

It should thus be clear that whether a given proposition is a partial answer is an empirical question that should be addressed by empirical means: for example, with the application of *the should test*, as proposed here. The fact (11)b is an answer (it eliminates the NO cell, as exposed by *the should test*) confirms that answerwood should be defined in terms of contextual entailment (as stipulated in the classical framework and as presupposed in the revised—probabilistic—framework).[82] Now, in the classical framework, one has no other option than concluding that (10)b is irrelevant relative to (10)a: as disclosed by *the should test*, $[\![(10)\text{b}]\!]^{+}$ intersects every cell of $I_{(10)\text{a}/\text{C}}$. This is a structural limitation of the classical framework, in which something is either a partial answer, and hence relevant, or not a partial answer, and hence irrelevant. In the revised (probabilistic) framework, one can treat (10)b as a relevant contribution ($[\![(10)\text{b}]\!]^{+}$ is p-relevant) despite it not being a partial answer (despite it not eliminating any cell).

In what follows I will often use the term *relevance implicature*; this term is meant to designate the probabilistic inference that the speaker *intends* the hearer to derive in cases in which the positive extension of the uttered sentence fails to provide a partial answer to the issue yet succeeds in securing p-relevance. For example, in (8)b/(10)b, the speaker implicates that the actual world is almost certain to be in the YES cell; in (1)b, in turn, the speaker implicates that the actual world is more likely to be in k_2 than in k_1; in (5)b, B implicates that Smith (perhaps via the accommodation of some contextual assumption) is more likely than not to have a girlfriend. In (7)b/(9)b, by contrast, the speaker asserts that the actual world is in the YES cell, but she does not implicate it.

[82] In the classical framework, p is a partial answer to I (a partition of C) iff $p \cap \text{C} \neq \emptyset$ and, furthermore, there is an $\text{X} \subset I$ such that $\{p \cap \text{C}\} \subseteq \bigcup\text{X}$. In the revised (probabilistic) framework, this can be translated as follows: p is a partial answer to I (a partition of Ω) and P (the common ground) iff there is an $\text{X} \subset I$ such that $\text{P}(\bigcup\text{X} \mid p) = 1$.

3 PROBABILISITC (NON-)MAXIMALITY: A DEVELOPMENT ON KRIŽ (2015)

Probabilistic (non-)maximality (PNM for short) has two ingredients: a (very slight) revision of Grice's (1975) original maxim of Quality and an account of how the meaning of sentences is assimilated into the common ground.

3.1 The maxim of Quality

Grice's (1975) **original** maxim of Quality only allows speaker to say sentences that, as far they know, are true. Under the assumption that natural language sentences are bivalent, an **alternative** (and equivalent) formulation of this maxim becomes possible: speakers are only allowed to say sentences that, as far as they know, are not false. However, if natural language sentences can have more than two truth-values due to homogeneity, then these two formulations are not equivalent: the former, but not the latter, bans speakers from saying sentences which are undefined. In Križ's framework, non-maximality arises, among other things, because speakers can, under certain conditions, use undefined sentences to communicate information; if speakers were banned from using undefined sentences, then non-maximal interpretations would never be computed. Given these considerations, it is clear that the alternative formulation of the maxim of Quality is, for our purposes, the appropriate one.

(IV) PNM's Maxim of Quality
A speaker may say only sentences which, as far as she knows, are not false.

3.2 Updating the common ground with trivalent sentences

PNM implements Križ's (2015) core idea, i.e.: a sentence S is interpreted as 'something is the case that, relative to the issue under consideration, is equivalent to S being true (and S isn't false)'. In Križ's account, this idea is implemented as follows: if two worlds w_0 and w_1 happen to be live in the same cell of issue I ($w_0 \approx_I w_1$), and a sentence S is true in w_0, undefined in w_1, and there is no issue cell in which the positive and negative extension of S overlap (*Addressing* constraint), then the S will be interpreted non-maximally: S's communicated content (a proposition) will be true in w_0 but also in w_1. However, if the positive extension and the extension gap of S never overlap (if there aren't two worlds w_0 and w_1 such that $w_0 \approx_I w_1$), then the sentence will be interpreted maximally (provided *Addressing* is met). As has been discussed, Križ's account runs into trouble as soon as non-idealised issues are considered (cf. Chapter IV, § 4).

PNM's implementation of Križ's idea is as follows: if the positive extension and the extension gap of S are equivalent relative to the issue under consideration, then S will be interpreted non-maximally; if, conversely, these two propositions are not equivalent, then S will be interpreted maximally. What does it mean for S^+ and $S^\#$ to be equivalent relative to the issue under consideration? In PNM's framework, it means that, given an issue *I* and a prior distribution P, those issue cells that have their probability increased as a result of updating P with S^+ are the same issue cells that have their probability increased as a result of updating P with $S^\#$; furthermore, it means that, given any two cells *i* and *ii*, if updating P with S^+ triggers the inference that the actual world is more likely to be in *i* than in *ii*, then updating P with $S^\#$ also triggers this inference (and vice versa). This notion is defined below.

(V) Equivalence between propositions rel. to a partition and a probability distribution (v.1)

<u>Notation</u>:

P is a probability distribution over Ω
I is a partition on Ω: $I = \{i_i,\ldots,i_n\}$
i is partition cell: $i \in I$
p and q are propositions: $p, q \in \wp(\Omega)$
$\approx_{I,\mathrm{P}}$ is an equivalence relation between propositions parametrised by I and P

<u>Definition:</u>

$p \approx_{I,\mathrm{P}} q$ if, $\forall i \in I$,

$$(\mathrm{P}(i \mid p) - \mathrm{P}(i)) \times (\mathrm{P}(i \mid q) - \mathrm{P}(i)) > 0 \ \lor \ \mathrm{P}(i) = \mathrm{P}(i \mid p) = \mathrm{P}(i \mid q),$$

and for any two cells i and $ii \in I$,

$$\mathrm{P}(i \mid p) > \mathrm{P}(ii \mid q) \leftrightarrow \mathrm{P}(i \mid q) > \mathrm{P}(ii \mid q).$$

In what follows, I will introduce PNM's update algorithm, and explain what each of its steps do.[83]

> **(VI) PNM's update algorithm**
>
> Let I be a partition on Ω, i a partition cell, S a natural language sentence (S^+ = S's positive extension; S^- = S's negative extension; $S^\#$ = S's extension gap), P_0 a prior probability distribution over Ω, and $\approx_{I,P_0}$ an equivalence relation between propositions. Update with S proceeds as follows. For any $p \in \wp(\Omega)$:
>
> (i) $\mathbf{P_1(p) = P_0(p \mid \Omega \setminus S^-)}$
>
> (ii) $\mathbf{P_2(p) = P_1(p \mid S^+)}$
>
> (iii) If $S^+ \approx_{I,P_0} S^\#$, then $\mathbf{P_3(p) = \sum_{i \in I} P_1(p \mid i) \times P_2(i)}$. Otherwise, $\mathbf{P_3(p) = P_2(p)}$.
>
> P_3 is the final interpretation.

Step (i) follows from the maxim of Quality (cf. Def. (IV)): speakers are banned from saying sentences that, as far they know, are false; under the assumption that speakers in fact adhere to this maxim, then, when S is uttered, it must be the case the speaker does not believe that S is false: the maxim of Quality, therefore, has the net effect of assigning the worlds in which S is false probability zero.

Step (ii) updates P_1 (the probability distribution that results from updating P_0 with the set of worlds in which S is not false) with the positive extension of S. P_2 results from this operation. If $P_1 = P_2$, then it means that the sentence is bivalent (i.e. not homogenous); if $P_1 \neq P_2$, then it means that the sentence is trivalent (or homogenous).

Step (iii) generates P_3, the final interpretation. If $P_1 = P_2$, then $S^\# = \emptyset$ and, as a result, $S^+ \approx_{I,P_0} S^\#$ is undefined (i.e. if $S^\# = \emptyset$, then, $\forall i \in I$, $P_0(i \mid S^\#)$ is undefined). Thus, if $P_1 = P_2$, then, the final interpretation, P_3, will be identical to P_1 and P_2. This means that, if S is not homogeneous, PNM's update procedure generates the same result as if P_0 had been updated with the truth-conditions of S; indeed, if S is not homogeneous, then $P_1 = P_2 = P_3 = P_0|S^+$.

If $P_1 \neq P_2$, then it must be checked whether $S^+ \approx_{I,P_0} S^\#$ holds: if it does not hold, then it means that S^+ and $S^\#$ are not equivalent for current purposes: the sentence is therefore interpreted maximally, i.e.: $P_2 = P_3$.

[83] Many thanks to Manuel Križ and Benjamin Spector for helping me to define (VI).

If, however, $S^+ \approx_{I,P_0} S^\#$ holds, then a non-maximal interpretation is generated, which is obtained as follows. The algorithm checks what fraction of each cell is taken up by p in P_1 and, subsequently, weights that fraction by the overall probability of the appropriate cell as given in P_2.[84] This generates a distribution in which each issue cell has the same probability mass as if S was true: however, in P_3, the relative likelihood of worlds within each cell is the same as in P_1. In Fig. 3, a graphical illustration of step (iii) is given.

Figure 3. Step (iii), a graphical illustration

Let's assume issue I consists of two cells, cell i and cell ii; further, let's assume that P_1 (cell i) = 0.6, P_1 (cell ii) = 0.4, P_1 (p | cell i) = 1/3 (i.e. p takes up 1/3 of cell i's probability), and P_1 (p | cell ii) = 2/3 (i.e. p takes up 2/3 of cell ii's probability).

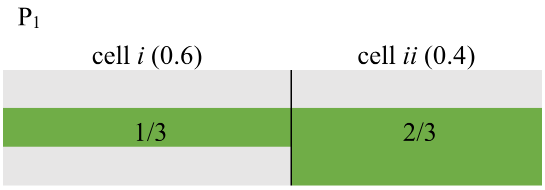

p is the union of the green sub-cells.
P_1 (p): 1/3 × 0.6 + 2/3 × 0.4 = **0.467**

Let's assume that P_2 (cell i) = 0.5 and P_2 (cell ii) = 0.5, as illustrated below:

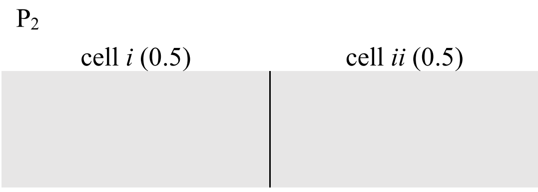

To find out P_3 (p), one proceeds as follows:
P_1 (p | cell i) × P_2 (cell i) = 1/3 × 0.5 = 0.167
P_1 (p | cell ii) × P_2 (cell ii) = 2/3 × 0.5 = 0.333
P_3 (p) = 0.167 + 0.333 = **0.5**

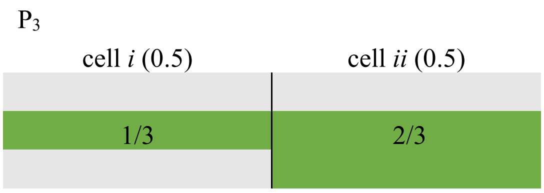

P_1 (p) ≠ P_3 (p), but…
P_1 (p | cell i) = P_3 (p | cell i) = 1/3
P_1 (p | cell ii) = P_3 (p | cell ii) = 2/3

[84] This specific procedure ('the algorithm checks what fraction of each cell is taken up by p in P_1 and, subsequently, weights that fraction by the overall probability of the appropriate cell as given in P_2') is an original idea of Benjamin Spector.

3.3 The theory put to work

In what follows, I will show how PNM derives a non-maximal and a maximal interpretation on the basis of the uttered sentence and the issue induced by the overt question.

3.3.1 The kids are tired: a non-maximal interpretation

As noted in the previous chapter, if Gerard replies to (12)a with (12)b_1, (12)b_1 is true enough even if some but not all his kids happen to be tired. Because of this, (12)b_1 is said to be interpreted non-maximally, i.e.: (12)b_1, in the stipulated context, does not appear to entail that all the kids are tired. However, if Gerard replies (12)b_2, and some but not all his kids happen to be tired, then what Gerard said is clearly false.

(12) Context: Gerard has ten kids; he and his kids are at a party.

a. Friend: Are you staying at the party for a bit longer?
b. Gerard: b_1. The kids are tired.
b_2. All the kids are tired.

(12)a, a yes/no question, induces an issue, let's call it *L*, which consists of two cells: l_1 (*staying at the party*) and l_2 (*not staying at the party*). For the sake of clarity, I will refer to (12)b_1 as ***the kids*** and to (12)b_2 as ***all the kids***. P_0, the prior probability distribution and *L*, the issue, are represented in Fig. 4 ('sc' stands for 'sub-cell', a subset of a cell).

Figure 4. Issue *L* at P_0.

l_1: Staying at the party $P_0(l_1) = 0.4$	l_2: Not staying at the party $P_0(l_2) = 0.60$
sc_1: Staying at the party & all the kids are tired $P_0(sc_1) = 0.01$	sc_4: Not staying at the party & all the kids are tired $P_0(sc_4) = 0.1$
sc_2: Staying at the party & some but not all the kids are tired $P_0(sc_2) = 0.1$	sc_5: Not staying at the party & some but not all the kids are tired $P_0(sc_5) = 0.40$
sc_3: Staying at the party & none of the kids are tired $P_0(sc_3) = 0.29$	sc_6: Not staying at the party & none of the kids are tired $P_0(sc_6) = 0.1$

the kids (due to its homogenous semantics) is false in worlds in which none of the kids are tired; thus, if ***the kids*** is uttered, these worlds, as per (VI/i), will be assigned probability zero: P_0 will be updated with $\{w$: ***the kids*** is not false in $w\}$, which will generate a posterior distribution, namely, P_1, such that $P_1(sc_3) = P_1(sc_6) = 0$. ***all the kids***, as opposed to ***the kids***, is false in worlds in which not all the kids are tired; thus, if ***all the kids*** is uttered, these worlds, as per (VI/i), will be assigned probability zero: P_0 will be updated with $\{w$: ***all the kids*** is not false in $w\}$, which will generate a posterior, P_1, such that $P_1(sc_2) = P_1(sc_3) = P_1(sc_5) = P_1(sc_6) = 0$. Thus, if ***the kids*** is uttered, P_1 is going to look as in Fig. 5, whereas if ***all the kids*** is uttered, P_1 is going to look as in Fig. 6.

Figure 5. P_1, issue L, if ***the kids*** is uttered.

l_1: Staying at the party $P_1(l_1) = 0.18$	l_2: Not staying at the party $P_1(l_2) = 0.82$
sc_1: Staying at the party & all the kids are tired $P_1(sc_1) = 0.0164$	sc_4: Not staying at the party & all the kids are tired $P_1(sc_4) = 0.164$
sc_2: Staying at the party & some but not all the kids are tired $P_1(sc_2) = 0.164$	sc_5: Not staying at the party & some but not all the kids are tired $P_1(sc_5) = 0.656$

Figure 6. P_1, issue L, if ***all the kids*** is uttered.

l_1: Staying at the party $P_1(l_1) = 0.09$	l_2: Not staying at the party $P_1(l_2) = 0.91$
sc_1: Staying at the party & all the kids are tired $P_1(sc_1) = 0.09$	sc_4: Not staying at the party & all the kids are tired $P_1(sc_4) = 0.91$

At step (VI/ii), P_1 is updated with S^+: since ***the kids*** and ***all the kids*** have the same truth-conditions, P_2 will be same irrespective of which of these sentences is uttered**,** as illustrated in Fig. 7 below.

Figure 7. P_2, issue L, irrespective of whether ***the kids*** or ***all the kids*** is uttered.

l_1: Staying at the party $P_2(l_1) = 0.09$	l_2: Not Staying at the party $P_2(l_2) = 0.91$
sc_1: Staying at the party & all the kids are tired $P_2(sc_1) = 0.09$	sc_4: Not staying at the party & all the kids are tired $P_2(sc_4) = 0.91$

If S = ***all the kids***, then $P_1 = P_2$, $S^+ \approx_{I,P_0} S^\#$ will not hold, and the sentence's final interpretation, P_3, will be identical to P_2. However, if S = ***the kids***, then $P_1 \neq P_2$: indeed, whereas P_1 is as in Fig. 5, P_2 is as in Fig. 7. If $P_1 \neq P_2$, then it must be checked whether $S^+ \approx_{I,P_0} S^\#$ holds. To determine this, P_0 must be updated with the extension gap, that is, $\{w$: ***the kids*** is undefined in $w\}$, which generates the distribution in Fig. 8.

Figure 8. $P_0|S^\#$, issue L, if ***the kids*** is uttered.

l_1: Staying at the party $P_0|S^\#(l_1) = 0.2$	l_2: Not staying at the party $P_0|S^\#(l_2) = 0.8$
sc_1: Staying at the party & all the kids are tired $P_0|S^\#(sc1) = 0.2$	sc_4: Not staying at the party & all the kids are tired $P_0|S^\#(sc4) = 0.8$

In both $P_0|S^+$ (which is identical to P_2; cf. Fig. 7) and $P_0|S^\#$ (Fig. 8), the probability of l_2 increases and the probability of l_1 decreases relative to P_0 (Fig. 4): $S^+ \approx_{I,P_0} S^\#$ is therefore the case. Thus, ***the kids***'s final interpretation, P_3, corresponds to a non-maximal interpretation.

Figure 9. P_3, issue L, if ***the kids*** is uttered.

l_1: Staying at the party $P_3(l_1) = 0.09$	l_2: Not staying at the party $P_3(l_2) = 0.91$
sc_1: Staying at the party & all the kids are tired $P_3(sc_1) = 0.0082$	sc_4: Not staying at the party & all the kids are tired $P_3(sc_4) = 0.182$
sc_2: Staying at the party & some but not all the kids are tired $P_3(sc_2) = 0.0818$	sc_5: Not staying at the party & some but not all the kids are tired $P_3(sc_5) =0.728$

As shown in Fig. 9, in P_3 the cells have the same probability as in P_2 (Fig. 7): that is, for current purposes, everything is as if ***the kids*** was true; however, in P_3, the relative likelihood of every subset within each cell is like in P_1. This makes P_3 a non-maximal interpretation: P_3 does not rule out the possibility that the actual world is one in which some but not all the kids are tired. Indeed, from P_3, it can be inferred that *at least one of the kids is tired*; in addition, it can be inferred that *it is* (*a lot*) *more likely for Gerard to leave than to stay at the party* (the relevance implicature: 91% vs 9%).

Given a reasonable P_0, Gerard would be more likely to leave the party if all the kids are tired vis-à-vis if only a few of the kids are tired (i.e. the larger the number of tired kids, the greater the urgency to leave). This continues to hold at P_3, because neither probabilistic conditioning nor step (VI/iii) alters the relative likelihood of subsets within each cell. Thus, the final interpretation that PNM derives for ***the kids*** is much more nuanced that *at least one of the kids is tired*: it may be paraphrased as *at least one of the kids is tired and it is more likely than not that a significant number of them is tired as opposed to just one or two (and it is more likely for Gerard to leave than to stay at the party)*. This seems like a good result.

It should be noticed that ***the kids***, because of the plural, will typically be associated with a *strictly more than 1* interpretation. Like Spector (2007)—see also, Sauerland, Anderssen, and Yatsushiro (2005), and Zweig (2008)—, I am inclined to treat this component of the meaning of the plural morpheme as an implicature, rather than as part of its semantics. This implicature has not been taken into account in the derivation of ***the kids***: however, if it had been, in P_3 the proposition *at least two kids are tired* would have probability 1, while the proposition *exactly 1 kid is tired* would have probability 0.

3.3.2 The patients have terminal cancer: a maximal interpretation

Consider (13) below; if the Doctor were to reply to (13)a with (13)b_1, (13)$b_1$would be judged false if some but not all the patients had terminal cancer; in other words, in the stipulated context, (13)b_1 would be interpreted maximally. PNM, therefore, should predict $P_3 = P_2$ for both (13)b_1 and (13)b_2.

(13) Context: There are 10 cancer patients in room B.

a. *Nurse*: Are all the cancer patients in room B going to die in the next few weeks?
b. *Doctor*: b_1. The patients in room B have terminal cancer.
b_2. All the patients in room B have terminal cancer.

(13)a, a yes/no question, induces an issue, let's call it *M*, which consists of two cells: m_1 (*dying in the next few weeks*) and m_2 (not *dying in the next few weeks*). For the sake of clarity, I will refer to (13)b_1 as ***the patients*** and to (13)b_2 as ***all the patients***. P_0, the prior probability distribution, and *M*, the issue, are represented in Fig. 10.

Figure 10. Issue *M* at P_0.

m_1: all dying in the next few weeks $P_0(m_1) = 0.45$	*m_2: not all dying in the next few weeks* $P_0(m_2) = 0.55$
sc_1: ∀ dying in the next few weeks & ∀ terminal cancer $P_0(sc_1) = 0.30$	sc_1: not ∀ dying in the next few weeks & ∀ terminal cancer $P_0(sc_1) = 0.05$
sc_2: ∀ dying in the next few weeks & ∃¬∀ terminal cancer $P_0(sc_2) = 0.10$	sc_5: not ∀ dying in the next few weeks & ∃¬∀ terminal cancer $P_0(sc_5) = 0.20$
sc_3: ∀ dying in the next few weeks & ¬∃ terminal cancer $P_0(sc_3) = 0.05$	sc_6: not ∀ dying in the next few weeks & ¬∃ terminal cancer $P_0(sc_6) = 0.30$

Like before, the interpretation of the *all*-sentence, ***all the patients*** in this case, is such that $P_3 = P_2$: this is because *all*-sentences do not have an extension gap; as a result, $S^+ \approx_{I,P_0} S^\#$ will not hold and the sentence's final interpretation will be identical to P_2, as stipulated in step (VI/iii). ***the patients***, by contrast, is homogenous; therefore, it must be checked whether or not $S^+ \approx_{I,P_0} S^\#$ holds. To determine this, P_0 must be updated with S^+, that it, {*w* : ***the patients*** is true in *w*}, which generates the distribution in Fig. 11, and with $S^\#$, that is, {*w* : ***the patients*** is undefined in *w*}, which generates the distribution in Fig. 12.

Figure 11. ***the patients***, issue *M*, $P_0|S^+$.

m_1: all dying in the next few weeks $P_0|S^+(m_1) = 0.86$	*m_2: not all dying in the next few weeks* $P_0|S^+(m_2) = 0.14$
sc_1: ∀ dying in the next few weeks & ∀ terminal cancer $P_0|S^+(sc_1) = 0.86$	sc_1: not ∀ dying in the next few weeks & ∀ terminal cancer $P_0|S^+(sc_1) = 0.14$

Figure 12. ***the patients***, issue *M*, $P_0|S^\#$.

m_1: all dying in the next few weeks $P_0|S^\#(m_1) = 0.33$	*m_2: not all dying in the next few weeks* $P_0|S^\#(m_2) = 0.66$
sc_2: ∀ dying in the next few weeks & ∃¬∀ terminal cancer $P_0|S^\#(sc_2) = 0.33$	sc_5: not ∀ dying in the next few weeks & ∃¬∀ terminal cancer $P_0|S^\#(sc_5) = 0.66$

As can be seen from the above, $S^+ \approx_{I,P_0} S^\#$ does not hold: learning that *all the patients have terminal cancer* raises the probability of the actual world being in m_1, whereas learning that *some but not all the patients have terminal cancer* has the opposite effect: raises the probability of the actual world being in m_2. Because $S^+ \approx_{I,P_0} S^\#$ does not hold, then then final interpretation of ***the patients*** is predicted to be identical to P_2: i.e., ***the patients*** is predicted to be interpreted maximally, as illustrated in Fig. 13.

Figure 13. ***the patients***, issue *M*, P_3.

m_1: all dying in the next few weeks $P_3(m_1) = 0.86$	*m_2: not all dying in the next few weeks* $P_3(m_2) = 0.14$
sc_1: ∀ dying in the next few weeks & ∀ terminal cancer $P_3(sc_1) = 0.86$	sc_1: not ∀ dying in the next few weeks & ∀ terminal cancer $P_3(sc_1) = 0.14$

Again, this seems like a good result: from P_3 it can be inferred that *all the patients have terminal cancer* (the maximal interpretation), and that *it is much more likely than not that they will die in the next few weeks* (the relevance implicature).

3.4 The kids looked happy: a harder case

Consider the following case:

(14) Context: Jane, a newly trained teacher, had to teach a class for the first time in her life.

a. *George*: How did the class go? [cell *i* (*class went well*); cell *ii* (*class didn't go well*)]
b. *Jane*: b_1. The kids looked happy.
b_2. All the kids looked happy.

Quite clearly, (14)b_1 is interpreted non-maximally: Jane could have felicitously uttered (14)b_1 to describe a situation in which most of the kids enjoyed the class but one or two did not (the same, as has been discussed extensively, cannot be said of (14)b_2—unless, of course, Jane was exaggerating).

Insofar as (14)b_1 is interpreted non-maximally, $S^+ \approx_{I,P_0} S^\#$ should hold: however, it seems to me that (14)b_1 can be interpreted non-maximally even if $S^+ \approx_{I,P_0} S^\#$ does not hold. Learning that *all the kids looked happy* should, given a reasonable common ground, raise the probability that the class went well; however, it is not clear to me that the same can be said of *at least 1 kid looked happy and at least 1 kid didn't look happy*. It could well be that such a proposition has not detectable effect on the issue (i.e. $P_0 \mid S^\#$ (cell i) $\cong P_0$ (cell i) and $P_0 \mid S^\#$ (cell ii) $\cong P_0$ (cell ii)), and that, in spite of this, the sentence is interpreted non-maximally. To handle cases such as (14)b_1, a slight revision of (V)—given in (VII) below—could be considered.

(VII) Equivalence between propositions rel. to a partition and a probability distribution (v.2)

Notation:

P is a probability distribution over Ω
I is a partition on Ω: $I = \{i_i, \ldots, i_n\}$
i is partition cell: $i \in I$
p and q are propositions: $p, q \in \wp(\Omega)$
$\approx_{I,P}$ is an equivalence relation between propositions parametrised by I and P

Definition:

$p \approx_{I,P} q$ if, $\forall i \in I$,

$$(P(i \mid p) - P(i)) \times (P(i \mid q) - P(i)) \geq 0,$$

and for any two cells i and $ii \in I$,

$$P(i \mid p) > P(ii \mid q) \leftrightarrow P(i \mid q) > P(ii \mid q).$$

In (VII), '>' has been replaced by '≥': according to the revised definition, two propositions are equivalent (relative to an issue and a probability distribution) if they do not pull the cells' probabilities in different directions (i.e. unlike in (V), it is not required that the propositions pull in the same direction); furthermore, just like in (V), it is required that the two propositions induce the same ordering of cells. Provided that this slight change is made, PNM may be able account for the contrast in interpretation between (14)b_1 and (14)b_2 (however, cf. discussion in the next section).

3.5 Does PNM make sense of our usability intuitions regarding homogenous sentences?

It does, to an extent. To start with, it is worth recalling that a homogenous sentence, if uttered to describe a situation in which the sentence is undefined, gives rise to a homogeneity violation, as illustrated below:

(15) <u>Context:</u> There are 10 kids: 7 look tired, and 3 do not; referring to the whole group of kids, Peter says:

a.	# The kids are tired	[Homogeneity violation]
b.	Some of the kids are tired	[True sentence]
c.	All the kids are tired	[False sentence]

Quite clearly, (15)a does neither mean (15)b nor does it mean (15)c. However, a speaker can use (15)a non-maximally (to communicate that some of the kids are tired) as long as the common ground is such that it does not rule out the possibility of the sentence being true, and if, for current purposes, it is irrelevant whether some but not all or all the kids are tired. Likewise, a speaker can use (15)a maximally (to communicate that all the kids are tired) as long as the common ground is such that it does not rule out the possibility of the sentence being true, and if, for current purposes, it does matter whether some but not all or all the kids are tired.

Let's first see how PNM distinguishes between a maximal and a non-maximal interpretation.

(16) <u>Context:</u> Gerard has ten kids; he and his kids are at a party.

a. Friend: Are you staying at the party for a bit longer?
b. Gerard: The kids are tired.

(17) <u>Context</u>: There are 10 cancer patients in room B.

a. *Nurse*: Are all the cancer patients in room B going to die in the next few weeks?
b. *Doctor*: The patients in room B have terminal cancer.

(16)b and (17)b are homogenous sentences: (16)b is undefined if some but not all of the kids are tired and (17)b is undefined if some but not all the patients in in room B have terminal cancer. But, as discussed, the result of updating the common ground with these sentences, given the issue and the prior in operation in each case, is very different. PNM assigns probability zero to both $[\![(16)b]\!]^{-}$ and $[\![(17)b]\!]^{-}$ (as per (VI/i)); however, whereas $[\![(17)b]\!]^{\#}$ ends up having probability zero, $[\![(16)b]\!]^{\#}$ doesn't (as per (VI/iii)). This accounts for why (17)b—if given as an answer to (17)a—cannot be used to describe a situation in which some but not all the patients in room B have terminal cancer, whereas (16)b—if given as an answer to (16)a—can be used to describe a situation in which most but not all the kids are tired.

There is something else that PNM must be able to explain: why is that (16)b can be used to describe a situation in which only 3/10 of Gerard's kids are tired, whereas (14)b_1, repeated below as (18)b, cannot be used to describe a situation in which only 30% of the kids in Jane's class did not look happy.

(18) Context: Jane, a newly trained teacher, had to teach a class for the first time in her life.

a. *George*: How did the class go?
b. *Jane*: The kids looked happy.

Indeed, if one were to find out that, as a matter of fact, just 30% of Jane's class looked happy, the intuition at least is that Jane said something false; by contrast, if one were to find out that, as a matter of fact, 3/10 of Gerard's kids were tired, (16)b does not feel false: the intuition, as far as I can tell, is that (16)b is true enough.

PNM is not without resources to explain the contrast in interpretation between (16)b and (18)b. (18)b triggers the relevance implicature that the class is likely to have gone well; however, if just 30% the class looked happy, then the class is unlikely to have gone well. Thus, the advocate of PNM could just say this: (18)b is not usable to describe a situation in which only 30% of the class looked happy because, against such a situation, (18)b is misleading. Let's now consider (16)b: it triggers the relevance implicature that Gerard is likely not to stay at the party; if 3/10 of Gerard's kids are tired, then Gerard is likely not to stay at the party (having three tired kids seems like a good enough reason to leave). The advocate of PNM could thus say this: (16)b can be used to describe a situation in which 3/10 of Gerard's kids are tired because, against such a situation, (16)b is not misleading.

There is something conceptually odd about making sense of the data in the manner suggested above, however. To explain the difference in usability between (16)b and (17)b, PNM invokes the fact that, in (17), $P_3([\![(17)b]\!]^{\#}) = 0$ whereas, in (16), $P_3([\![(16)b]\!]^{\#}) > 0$. By contrast, to explain the difference in usability between (16)b and (18)b, PNM appeals to the consideration of whether or not the sentences are misleading. This is odd, because, as far as I can tell, what needs to be explained here is one and the same thing: the degree of deviation from truth that a homogenous sentence is able to tolerate against an issue and a common ground.

The problem is not just conceptual: PNM, as formulated, makes an incorrect prediction. For example, it predicts that (18)b could be usable if only 2 kids in Jane's class looked happy, and the school director, who witnessed the class, said to Jane: 'The class was fantastic; don't worry about the children looking a bit unhappy; they always look unhappy on Mondays'. Relative to this situation, (18)b feels false, yet it is not misleading: if this is what in fact happened, then Jane's class went well. The conclusion appears to be this: to account for the interpretation of (18)b, PNM should be reformulated so that the set of worlds in which, say, less than half of the kids in Jane's class looked happy, ends up having probability zero in the final distribution.

It is not entirely clear to me how to achieve this result at this point; I have an idea, however. It may be that extension gap is in fact context-sensitive: for example, in (18)b, the extension gap may be *more than half but not all the kids looked happy*, as opposed to *some but not all the kids looked happy*. If this was the case, then, it would not be necessary to reformulate (V) along the lines proposed in (VII);[85] furthermore, the contrast in interpretation between (16)b and (18)b would be accounted for: indeed, on such an account, these sentences would differ in terms of the size of their respective gaps and, hence, in terms of how much deviation from truth each of them tolerates.

Coincidentally (or perhaps not so coincidentally), context-sensitive gappiness, as will be discussed in the next chapter, is a (surprisingly) fruitful conjecture when one looks outside the specific realm of plural sentences; further research should determine whether the gap of plural sentences can be made context-sensitive and at what cost.

3.6 Some predictions

3.6.1 Prior dependency

It should be clear by now that whether a homogenous sentence ends up being interpreted maximally or non-maximally depends on whether $S^+ \approx_{I,P_0} S^\#$ holds, and whether $S^+ \approx_{I,P_0} S^\#$ holds, depends on the nature of the issue but also on the nature of P_0, the prior distribution. Insofar as I am right in thinking that the current issue can (at least in most cases) be identified with the overt question, then there should not be much room for disagreement as to which issue the uttered sentence is addressing. It is less straightforward to determine, however, what the prior distribution is. For each of the examples discussed, I have just assumed what I took to be a reasonable prior: however, given more elaborate scenarios, it may be that very different priors need to be assumed and that, under such priors, PNM makes incorrect predictions.

[85] I suspect that the actual notion of equivalence between two propositions that is operative during the process of interpretation is closer to (V) than to (VII). However, this is an empirical question that I have to leave for future research.

It is left to further research to determine how robust PNM's predictions are. Such an endeavour, if pursued, should take into account at least two potential caveats. First, the choice of issue is going to impose a non-trivial constraint on what the prior distribution can be: if I has only two cells, i_1 and i_2, then it cannot be the case that $P_0(i_1) = 0$ and $P_0(i_2) = 1$, for example: indeed, if the prior assigned i_1 and i_2 those probabilities, then the expected value of I would be zero (and, if so, why would such an issue be raised in the first place?). Second, the uttered sentence is also likely to put a constraint on the range of possible priors. As discussed in § 2.4, certain sentences, when uttered as replies to certain issues, trigger inferences which may end up being accommodated into the common ground. If something of this sort takes place, then P_0 is bound to bare the mark of it.

3.6.2 *Relevance implicature strength*

Consider, again, *the tired kids* case:

(19) Context: Gerard has ten kids; he and his kids are at a party.

a. Friend:	a. Are you staying at the party for a bit longer?
b. Gerard:	b_1. The kids are tired
	b_2. All the kids are tired

(19)b_1 is interpreted non-maximally, as noted; (19)b_2, by contrast, does not admit a non-maximal interpretation (i.e. (19)b_2 is not homogenous). Now, both (19)b_1 and (19)b_2 have something in common: both trigger the same relevance implicature, namely, that Gerard is more likely than not to leave the party. PNM, in this respect, makes a very specific prediction: the *strength* of this implicature should be the same irrespective of whether (19)b_1 or (19)b_2 is uttered: this is because the probability of the cells, according to PNM's update algorithm, is always given by P_2. Thus, the theory predicts that, if (19)b_1 were to raise the probability of the *not-staying-at-the-party* cell from 0.6 to 0.91, so should (19)b_2.

This prediction should be tested experimentally. Again, two caveats must be taken into account. First, if/when tested, it would be advisable that one group of participants is given the *the*-sentence (and not the *all*-sentence), while the other is given the *all*-sentence (and not the *the*-sentence). Without such a manipulation, the experiment could have the unwanted effect of forcing the *the*-sentences and the *all*-sentences into competition: this could in turn lead participants to make task-specific inferences which may in turn contaminate the experimental results. Second, universally quantified sentences are prone to be interpreted hyperbolically (e.g. 'my mother calls me every single day', 'all my friends have gone insane'). Thus, when testing the *all*-sentences, it is crucial to make sure that participants are interpreting them literally (as opposed to hyperbolically). If this is not controlled for, then the results are likely to be affected by this confounding factor.

4 COMPARISON WITH KRIŽ (2015)

4.1 Irrelevance and the stipulation of issues

In Križ's framework, a classical QUD framework without probabilities, answers whose truth-conditions intersect every single cell of the partition have to be banned on grounds of irrelevance: indeed, if such answers were permitted, Križ's account would predict 'the kids are tried' (if uttered as a response to the question 'are you staying at the party for a bit longer?') to effectively communicate the proposition *the speaker is staying at the party for a bit longer or the speaker is not staying at the party for a bit longer* (i.e. the union of all the issue cells).

There is no much room for maneuver: within Križ's (2015) framework, if the overt question is 'are you staying at the party for a bit longer?', then 'the kids are tired' will be interpreted non-maximally insofar as the sentence is taken to address a hypothetical issue which is different from the one induced by the overt question. Such a move, as discussed in the previous chapter, undermines the predictive power of Križ's account.

PNM operates within a probabilistic framework and 'being relevant' is being p-relevant (as opposed to classically relevant): this means that, insofar as the positive extension of a sentence S satisfies p-relevance, then S can be used to address the issue (irrespective of whether S is classically relevant). PNM, therefore, can derive non-maximal and maximal interpretations solely on the basis of the sentence uttered and the issue induced by the overt question: there is no need for stipulating additional issues.

4.2 Problematic case I: an instance of over-informativity

Križ's account, as discussed in Chapter IV, § 4.2.1, predicts (20)b to be interpreted as meaning *at least 1 kid came* (a bad prediction).

(20) <u>Context:</u> There are 10 kids.

a. *John*: Did some of the kids come?
b. *Paul*: The kids came.

PNM's prediction is, leaving aside some technicalities aside, the same:[86] learning that all the kids came (the positive extension of (20)b) raises the probability that some of the kids came; likewise, learning that some but not all the kids came (the extension gap of (20)b) raises the probability that some of the kids came; hence, $S^{+} \approx_{I,P_0} S^{\#}$ holds and, as a result, (20)b is predicted to be interpreted non-maximally.

[86] For example, Križ's account predicts (20)b to communicate the proposition *at least 1 kid came*; PNM, by contrast, predicts the final distribution to be one in which the proposition *at least 1 kid came* has probability 1.

However, in the framework that I am proposing, a constraint can be put in place that prevents (20)b from ever being interpreted non-maximally (in Križ's account, as will soon be discussed, such a constraint cannot be put in place). $[\![(20)\text{b}]\!]^+$ is an *over-informative* complete answer: that is, $[\![(20)\text{b}]\!]^+$ *strictly* entails the YES cell.[87] Thus, to account for why (20)b is not interpreted non-maximally, it is enough to stipulate that over-informative complete answers cannot be used to update the common ground (cf. Spector's (2007b) notion of *strong relevance*).

There are both empirical and conceptual reasons to ban over-informative complete answers (it is not an ad-hoc fix); on the empirical side, such answers feels like challenges to the proposed issue: indeed, (20)b, see also (21)B below, appears to trigger the inference (of strong meta-linguistic flavour) that the question posited is not fully appropriate.[88]

(21) A: Did mother buy a dog?
B: She bought a black dog.

The intuitive explanation for this phenomenon, which is also the conceptual motivation to put some sort of ban on over-informativity, is the following: by being over-informative, speakers are violating the maxim of Quantity (Grice 1975), which instructs them not to say more than they are required; as a result, listeners infer that speakers are in fact addressing some other issue/question (for example, in (21), *what type of dog did she buy?*).[89]

In Križ's framework, this (very reasonable) constraint cannot be imposed for a rather interesting reason: given how Križ's system is set up, in order for non-maximality to be computed against a 2-cell (or binary) issue, the sentence's positive extension *must* be a strict subset of one of the cells. Hence, if such a configuration were to be banned on the grounds of over-informativity (in order to prevent, for example, (20)b from being interpreted non-maximally), then non-maximality would never be computed.

[87] Put in probabilistic terms: a complete answer *a* is over-informative relative to cell *i* iff P (cell *i* | *a*) = 1 and, *in addition*, P (*a* | cell *i*) < 1.

[88] I'm aware that (21)B feels more like a challenge to the overt question than (20)b; this is because (20)b is (slightly) infelicitous: (20)b lacks an overt quantificational element that can be contrastively stressed and, for some reason, prosodically marked contrast between *some* and the quantifier that cannot be stressed (because it isn't there) is required. To make the point in a less technical way, (20)b's (mild) deviance appears to be related to (i)a's deviance.

(i) a. ?? Mary saw some of the students, but John saw the students.
b. Mary saw some of the students, but John saw ALL the students.

It could be argued that (20)b is not interpreted non-maximally precisely because of this reason (i.e. some prosodic condition of addressing is violated). I suspect, however, that these data require a deeper explanation. (Thanks to Manuel Križ for extensive discussion on this point).

[89] It is less clear to me whether there are good empirical and/or conceptual reasons to ban over-informative *partial* answers. I am not even sure whether it makes sense to talk about over-informative partial answers: partial answers are, in essence, under-informative: indeed, in such cases, what the speaker does is to provide less information than the issue requires.

Consider again the sentence ‘the professors smiled’ and issue *I* (*how did Sue’s talk go?*), as depicted in Fig. 14 below: the positive extension of ‘the professors smiled’ is a strict subset of i_1; to compute the sentence’s communicated content, on Križ’s account, the union of the cells that the sentence’s positive extension intersects is taken (in this case: $\bigcup\{i_1\} = i_1$).

Figure 14. A depiction of issue *I*.

i_1: positive reception	w_1: all the professors smiled w_2: all the professors smiled, except for Smith who stayed neutral
i_2: non-positive reception	w_3: half of the professors smiled w_4: none of the professors smiled

Now, there is a clear difference between (20), and *the Sue’s talk* case. The positive extension of (20)b is, as a matter of fact, a strict subset of one of the issue cells induced by (20)a; ‘the professors smiled’, in turn, is a strict subset of i_1 insofar as one takes the issue/Context-Set to be the one depicted in Fig. 14. But this issue/Context-Set is an idealisation: phenomenologically speaking, ‘the professors smiled’, given the question ‘how did Sue’s talk go?’, is not an over-informative complete answer: in fact, it is not even an answer: ‘the professors smiled’ does not come out marked in *the should test*, which indicates that both cells retain probability after the update: ‘How did the talk go? The professors smiled. So, it should have gone well.’ Cf. § 2.5.

The conclusion that can be drawn from these observations is this: In Križ’s framework, propositions that are not answers (they do not eliminate any cell of the issue) yet are p-relevant are ‘crammed together’ with over-informative answers; as a result of this, Križ cannot impose a ban on over-informative complete answers and stipulate that such answers trigger some form of issue accommodation: if he did so, non-maximal interpretations would never be computed against a binary issue. In PNM, by contrast, banning over-informative complete answers is, as far as I can tell, harmless: such a constraint accounts for why (20)b cannot be interpreted non-maximally and, importantly, does not get in the way of any documented case of non-maximality.

4.3 *Problematic case II: an instance of* $S^{+} \not\approx_{I,P_0} S^{\#}$

In the previous chapter, Chapter IV, § 4.2.2, repeated below as (22), was discussed. On Križ’s account, (22)b is predicted to communicate the same proposition as (22)c.

(22) <u>Context:</u> In order to pass, Peter has to solve either all the math problems, or at least half of them and write an essay on mathematical Platonism.
a. Did Peter pass the exam?
b. He solved the math problems.
c. He passed the exam.

This is not the prediction that PNM makes. The first thing to note is that $[\![(22)\text{b}]\!]^{+}$ is not an over-informative complete answer: $[\![(22)\text{b}]\!]^{+}$, in fact, is not even an answer, as shown in *the should test* below.

(23) Context: In order to pass, Peter has to solve either all the math problems, or at least half of them and write an essay on mathematical Platonism

a. Did Peter pass the exam?
b. He solved the math problems. So he should have passed.

$[\![(22)\text{b}]\!]^{+}$ eliminates no cell of the partition; furthermore, by uttering (22)b, the speaker implicates that it is *almost* certain that Peter passed the exam. What needs to be explained, as far as I can tell, is why (22)b is interpreted maximally.

This can be explained. The overt question induces a binary issue: cell *i* (*Peter passed*) and cell *ii* (*Peter didn't pass*). Let's assume a flat prior: P_0 (cell *i*) = P_0 (cell *ii*) = 0.5. The crucial question is whether $S^{+} \approx_{I,P_0} S^{\#}$ holds. Updating P_0 with $[\![(22)\text{b}]\!]^{+}$ will result in cell *i* concentrating almost all the probability mass: for example, $P_0\,|[\![(22)\text{b}]\!]^{+}$ (cell *i*) = 0.99 and $P_0\,|[\![(22)]\!]^{+}$ (cell *ii*) = 0.01. Updating with $[\![(22)\text{b}]\!]^{\#}$, by contrast, will have the effect of bringing the probability of cell *i* down (e.g. $P_0|S^{\#}$ (cell *i*) = 0.25 and $P_0|S^{\#}$ (cell *ii*) = 0.75): indeed, if Peter solved some but not all the math problems, then he did not pass the exam by solving all the exercises (which is one of the ways of passing). Thus, $S^{+} \approx_{I,P_0} S^{\#}$ is bound not to hold and, as a result, P_3 is bound to be identical P_2. This is, as far as I can tell, the correct prediction: (22)b's final interpretation enables us to infer that Peter solved all the math problems and that he is almost certain to have passed the exam.

4.4 The maxim of Quality

Križ's maxim of Quality is 'issue-dependent': it does not require speakers to tell the truth, but only that they should say sentences that are *sufficiently true* for current purposes (cf. Chapter IV, § 3.2). To prevent false sentences from qualifying as sufficiently true, Križ is forced to come up with *Addressing*, a constraint which, as discussed, extraordinarily limits the number of issues that a sentence can address. By contrast, in PNM, nothing like *Addressing* needs to be stipulated: the maxim of Quality is not issue-dependent and plays its usual role, i.e.: that of instructing speakers not to say false sentences.

It should be noted that Križ's account and PNM handle the oddness of cases such as (24) differently:

(24) Context: there's a pile with 10 books, 8 are clearly red and 2 are clearly blue; while pointing at the pile, John says…

\# The books are red.

Križ (2015) predicts this sentence to be a Quality violation: indeed, since it is common ground that not all the books (those we are talking about) are red, then (10) cannot be *sufficiently true* (i.e. in order for (24) to be *sufficiently true*, $[\![(24)]\!]^{+}$ should intersect at least 1 cell of the issue; this is not going to happen, however: the Context Set contains no worlds in which all the books are red. PNM's maxim of Quality, by contrast, does not predict (24) to be a Quality violation: PNM's maxim of Quality, after all, only prevents speakers from saying false sentences, and (24), in the context stipulated, is not false, just undefined. According to PNM, (24) is a Relevance violation: the maxim of Relevance instructs speakers to say sentences whose positive extension is p-relevant; however, (24) can never be p-relevant because $P([\![(24)]\!]^{+}) = 0$.

5 SUMMARY

In this chapter, I have done a number of things. In § 2, I have provided a probabilistic definition of relevance, as well as revisit some technical notions, such as the notions of *answer* and *relevance implicature*. In § 3, a thorough re-formulation (and re-conceptualisation) of Križ's (2015) account of (non-)maximality was attempted; in particular, I have put forward a novel notion of *equivalence for current purposes* (defined on propositions), as well as an algorithm to update the common ground with trivalent sentences, an algorithm that is homogeneity-sensitive as well as issue-sensitive. In § 4, PNM and Križ (2015) have been compared: the former account, it has been argued, fares better than the later in explaining the relevant data.

Chapter VI
INVISIBLE GAPS

1 OVERVIEW

There at least two families of sentences that support (non-)maximality. One the one hand, there are sentences which, in out-of-the-blue-contexts, exhibit a truth-value gap, such as unquantified plural sentences or habitual sentences; on the other hand, there are sentences which do not exhibit a truth-value gap, such as sentences containing numerals or absolute adjectives (e.g. 'the bookcase is *empty*' or 'I paid *£1000* for this laptop'). If homogeneity is in fact a pre-requisite for (non-)maximality, then the latter group of sentences must be homogenous, just like plural sentences. Can a sentence be homogenous despite not being visibly homogenous? I will argue that it can.

In this chapter, the focus will be placed on making sense of the idea of invisible homogeneity in absolute adjectives only; the tools that I will develop here, however, could be applied, I believe, to derive invisible homogeneity across a number of constructions, including numerals and geometrical terms.

2 RELATIVE, ABSOLUTE AND PARTIAL ADJECTIVES

Gradable adjectives can be operationally defined as expressions that admit modification by a comparative construction; for example, *tall*, *bent* and *straight*, as shown in (1) below, are gradable, whereas *vegan* and *dead* are not.

(1) a. John is *tall*-er than Paul.
b. Rod A is more *bent* that Rod B.
c. Rod B is *straight*-er than Rod A.
d. # John is more *vegan* than Jeanne.
e. # Napoleon is more *dead* than John Lennon.

Since Unger (1975), it is common to distinguish between two classes of gradable adjectives; on the one hand, **relative adjectives**, such as *tall*, *rich*, or *young*; on the other hand, **absolute adjectives**, such as *empty*, *straight*, or *bent*.

Relative adjectives are called 'relative' because their interpretation is tied up to a context-dependent standard of comparison; for example, and assuming that John is 1.60m, the sentence 'John is *tall*' may be judged true if John is a child and is common ground that John is being compared to other children, but false if he is an adult and is common ground that he is being compared with other adults. This context-dependence is accounted for by positing that such adjectives come with a standard of comparison: either a comparison class variable (a variable ranging over properties), a context-sensitive threshold (a variable ranging over degrees or a context-sensitive function that returns a degree), or both (see Kennedy (2007) for a comparison between different approaches). Relative adjectives, in addition to being standard-dependent, are vague (more on this to come), a feature that has been taken to be linked to their standard-dependence (cf. Kennedy 2007; Lassiter and Goodman 2013; Qing and Franke 2014).

Absolute adjectives come in two types (cf. Yoon 1996; Rotstein and Winter 2004; Burnett 2014): **total adjectives**, such as *straight* or *closed*, and **partial adjectives**, such as *bent* or *open*, which are the complementary antonyms of the former adjectives, respectively. Not all total adjectives have a partial counterpart: for example, in English, there is no lexical item that denotes *not empty* (which, if existed, would be the complementary antonym of *empty*). In contrast to relative adjectives, the meaning of absolute adjectives, as discussed at length in Kennedy (2007), neither depends on a contextual standard of comparison nor does it exhibit the classic symptoms associated with vagueness. In Table 1 below, a list with some gradable adjectives (which includes relative, absolute and partial adjectives) is given.[90]

Table 1. Gradable adjectives.

Absolute adjectives		**Relative adjectives**	
Total	**Partial**		
Empty	-	Tall	Short
Full	-	Young	Old
Straight	Bent	Rich	Poor
Clean	Dirty	Expensive	Cheap
Pure	Impure	Big	Small
Invisible	Visible	Fast	Slow
Closed	Open[91]		
Clear	Unclear		
Perfect	Imperfect		

What is the empirical basis for the claim that, whereas the meaning of relative adjectives can be relativised a contextual standard of comparison, the meaning of absolute adjectives cannot? The strongest argument in support of this claim comes perhaps from the so-called *define description test* (e.g. Kyburg and Morreau 2000; Kennedy 2007). Consider, for example, the contrast between (2) and (3)/(4).

(2) Context: *There are two mamushkas: one if 6cm high and the other is 15cm high.*

Pass me the tall one.

[90] I am not sure whether this classification is exhaustive; in fact, I am not even sure whether this classification is fully adequate. However, for the modest purposes of this chapter, it will suffice.
[91] Not so clear how to classify 'open'; see § 4.1.2.

(3) Context: *There are two rods: one is 50° bent, and the other is 100° bent.*

Pass me the straight one.

(4) Context: *There are two rods: one is 50° bent, and the other is 100° bent.*

Pass me the bent one.

The fact that (2) is interpreted as a request for the taller of the two objects suggests that the meaning of *tall* accesses some sort of contextual standard of comparison, a standard against which one of the mamushkas comes out tall and the other does not. If absolute adjectives could be interpreted relative to a contextual standard, (3) and (4) would be felicitous: that is, (3) would be interpreted as a request for the straighter of the two, and (4) should interpreted as a request for the more bent of the two. The fact that (3) and (4) cannot be interpreted in this way, not even to avoid presupposition failure, indicates that absolute adjectives are indeed absolute: their meaning cannot be relativised to a contextual standard of comparison.

What is the basis for the claim that relative adjectives are vague whereas absolute adjectives aren't? One of the observations that has served as basis for such claim is that relative adjectives, unlike absolute ones, trigger (fallacious) soritical reasoning (see Chapter III, § 3.3). Consider for example, the contrast between (5) and (6)/(7).

(5) Relative adjective: *tall*
a. Any man whose height is 187cm is tall.
b. Any man who is 1cm shorter than a tall man is also tall.
c. Any man is tall.

(6) Total adjective: *straight*
a. A rod whose degree of curvature is 0° is straight.
b. #Any rod whose degree of curvature is 1° greater than that of a straight rod is also straight.
c. Any rod is straight.

(7) Partial adjective: *bent*
a. A rod whose degree of curvature is greater than 0° is bent.
b. # Any rod that is 1° straighter than a bent rod is also bent.
c. Any rod is bent.

Indeed, whereas one readily accepts the inductive premise in (5)b, (6)b and (7)b are much harder to accept (if possible to accept at all). Of course, if (6)b read 'any rod whose degree of curvature is 0.00000001° greater than that of a straight rod is also straight', then (6)b would likely be accepted as true and the false conclusion in (6)c would then follow. But such an inductive premise, it should be noted, brings something else into the picture: that the human senses are not sensitive to microscopic differences and, as a result, for all practical purposes, n + 0.00000001° = n. In the examples above, I am controlling for this confounding factor: the humans senses are sensitive to both 1cm-differences and 1°-differences; in spite of this, (5)b seems true, whereas (6)b and (7)b do not. This contrast, I believe, exposes what I take to be a fundamental difference between *tall* and *straight*/*bent*: whereas there does not seem to be a fact of the matter as to how tall a person has to be in order to count as tall, there clearly is a fact of the matter as to whether something is straight or bent: if its degree of curvature is zero, then it is straight; if its degree of curvature is greater than zero, then it is bent.

2.1 Definite plurals versus absolute adjectives

Though absolute adjectives are not vague, sentences that contain them can, given an appropriate context, be interpreted imprecisely or non-maximally. As discussed in Chapters III and IV, (non-)maximality is associated with three features: (i) *context-sensitivity*; (ii) *overt slack regulation* (via homogeneity removal); and *strict behaviour in contradiction test* (SBCT). Both unquantified plural sentences and sentences such as 'the bookcase is empty' possess these features, as shown in Chapter III; these sentences differ, however, in that the latter, unlike the former, do not exhibit a (visible) homogeneity kind of gap. Consider, for example, (8)a and (8)b vis-à-vis (9)a and (9)b.

(8) *There're 10 books on the table…*
a. John read the books.
b. John didn't read the books.

(9) *There's a piggy bank on the table…*
a. The piggy bank is empty.
b. The piggy bank isn't empty.

In out-of-the-blue contexts, (8)a and (8)b exhibit an extension gap: (8)a is true if John read all the books, false if he read none, and neither true nor false if he read some but not all; likewise, (8)b is true if John read none of the books, false if he read all, and neither true nor false if he read some but not all. The same, however, cannot be said of (9)a and (9)b: (9)a is true if there is no coin in the piggy bank; if there is 1 coin in the piggy bank, then (9)a (in an out-of-the-blue context) is, as far as I can tell, false; likewise, (9)b is true if there is 1 or more coins in the piggy bank and false if the piggy bank contains no coin. Hence, unlike (8)a and (8)b, (9)a and (9)b appear to have complementary truth-conditions.

This looks like a problem for an account that attempts to link (non-)maximality with homogeneity: (9)a and (9)b are not visibly homogenous and yet, as shown in in Chapter III, support (non-)maximality just like plural sentences do. But… is this really a problem? Perhaps it isn't: as Križ (p.c.) once pointed out to me, 'who said that homogeneity has to always be visible?'

3 ABSOLUTE ADJECTIVES

In what follows, I will put forward a trivalent semantics for absolute adjectives, a semantics that, in conjunction with (non-)maximality and an alignment condition, *predicts* gap invisibility. This approach, as will be discussed, has interesting theoretical consequences. First, it enables us to derive (im)precise interpretations of sentences such as 'the bookcase is empty' via PNM, just like with unquantified plural sentences. This is desirable because, as shown in Chapter III, the phenomenon of (im)precision or (non-)maximality is just *one* phenomenon and should therefore be analysed as such. Second, it gets the negation facts right for sentences containing absolute adjectives. As discussed in Chapter III, § 3.4, an imprecise interpretation of 'the bookcase is empty' is (or at least appears to be) weaker than its strict/precise interpretation, whereas an imprecise interpretation of 'the bookcase is not empty' is (or at least appears to be) stronger than its strict/precise counterpart. The semantics that I will put forward, unlike Lasersohn's (1999) account, predicts that there should be such an asymmetry. Finally, the framework that I will develop here casts a new light on the semantic nature of modifiers such as *completely* and *slightly*: such modifiers are to be thought of as homogeneity removers (or so will I argue).

3.1 Absolute adjectives: a semantics

As has become standard in degree semantics, I will assume that gradable adjectives are associated with a measure function: a function that takes two arguments, a world and an object, and returns a degree on a measurement scale. I will represent the co-domain of the measure function (the scale) as a pair of the form $\langle \mathbb{R}_+, \delta \rangle$, where $\mathbb{R}_+$ is the set of non-negative real numbers and δ a dimension of measurement (degrees are therefore a shorthand for $\langle n, \delta \rangle$, where $n \in \mathbb{R}_+$).

3.1.1 Total adjectives

<u>Definitions</u>: μ_δ is a measure function along the dimension δ; x is a variable ranging over entities, and d is a variable ranging over degrees; degrees are of semantic type d; ρ is a granularity parameter: an object-language variable ranging over degrees; g is the assignment function; for any assignment function g, granularity parameter ρ, and degree d, g$[d/\rho]$ is the function that is just like g save that the value it assigns to ρ is d. $\mathfrak{D}_t$, the domain of truth values, is {T, F, #}.

I shall formalise total adjectives (TAs) as follows:

(10)

$$[\![\mathrm{TA}]\!]^{w,g} \quad = \quad \lambda d_{d} \,.\, \lambda x_{e} \,.\, \begin{cases} \text{T iff } \mu_{\delta}\,(w, x) = 0 \\ \text{F iff } \mu_{\delta}\,(w, x) > d \\ \#\text{ iff } 0 < \mu_{\delta}\,(w, x) \leq d \end{cases}$$

Total adjectives are of type $\langle d,\langle e,t\rangle$: they denote functions from degrees to $\langle e,t\rangle$ functions. Let's use (10) to derive the trivalent meaning of 'the piggy bank is empty'.

(11)

a. $$[\![\text{empty}]\!]^{w,g} \quad = \quad \lambda d_{d} \,.\, \lambda x_{e} \,.\, \begin{cases} \text{T iff } \mu_{\delta}\,(w, x) = 0 \\ \text{F iff } \mu_{\delta}\,(w, x) > d \\ \#\text{ iff } 0 < \mu_{\delta}\,(w, x) \leq d \end{cases}$$

b. $$[\![\text{empty}]\!]^{w,g}\,([\![\rho]\!]^{w,g}) \quad = \quad \lambda x_{e} \,.\, \begin{cases} \text{T iff } \mu_{\delta}\,(w, x) = 0 \\ \text{F iff } \mu_{\delta}\,(w, x) > \mathrm{g}(\rho) \\ \#\text{ iff } 0 < \mu_{\delta}\,(w, x) \leq \mathrm{g}(\rho) \end{cases}$$

c. $$[\![\text{empty}_{\rho}]\!]^{w,g}\,([\![\text{the piggy bank}]\!]^{w,g}) \; = \; \begin{cases} \text{T iff } \mu_{\text{count}}\,(w, \iota x \,.\, \text{piggy-bank}'(w, x)) = 0 \\ \text{F iff } \mu_{\text{count}}\,(w, \iota x \,.\, \text{piggy-bank}'(w, x)) > \mathrm{g}(\rho) \\ \#\text{ iff } 0 < \mu_{\text{count}}\,(w, \iota x \,.\, \text{piggy-bank}'(w, x)) \leq \mathrm{g}(\rho) \end{cases}$$

For ease of exposition, I will define some technical terms. E is an expression whose semantic type ends in *t* (but which is not itself of type *t*) (e.g. 'empty'). The positive extension of E in *w*, notated $[\![E]\!]^{w,g,+}$, is the function that is just like $[\![E]\!]^{w,g}$ (the denotation of E in *w*) except that it conflates # and F (i.e. the function that maps a list of arguments to T just in case $[\![E]\!]^{w,g}$ maps it to T, and maps that list to F otherwise). The negative extension of E in *w*, notated $[\![E]\!]^{w,g,-}$, is the function that maps a list of arguments to T just in case $[\![E]\!]^{w,g}$ maps it to F, and maps that list to F otherwise. Finally, the extension gap of E in *w*, notated $[\![E]\!]^{w,g,\#}$, is the function that maps a list of arguments to T, just in case that $[\![E]\!]^{w,g}$ maps it to #, and maps that list to F otherwise. I will say that E doesn't have an extension gap if the image of the domain of $[\![E]\!]^{w,g}$ under $[\![E]\!]^{w,g,\#}$ is {F}.

In the case at hand: [92]

(12)

a.	$[\![\text{empty}]\!]^{w,g,+}$	=	$\lambda d_{d} \,.\, \lambda x_{e} \,.\, \mu_{\delta}\,(w, x) = 0$	[Positive extension]
b.	$[\![\text{empty}]\!]^{w,g,-}$	=	$\lambda d_{d} \,.\, \lambda x_{e} \,.\, \mu_{\delta}\,(w, x) > d$	[Negative extension]
c.	$[\![\text{empty}]\!]^{w,g,\#}$	=	$\lambda d_{d} \,.\, \lambda x_{e} \,.\, 0 < \mu_{\delta}\,(w, x) \leq d$	[Extension gap]

[92] In what follows, I am assuming Heim and Kratzer's (1998: 37) well-established conventions to read λ-terms. In particular, if γ is a statement, '[λψ . γ]'should be read as 'the function that maps ψ to T if γ, and to F otherwise.'

Furthermore, and as in previous chapters, I will identify the positive extension of a sentence S (which is not world-relative), notated $[\![S]\!]^{g,+}$ (or simply S^{+}), with (the function that characterises) the set of worlds in which the sentence is true. Similarly, I will identify the negative extension of S, notated $[\![S]\!]^{g,-}$ (or simply S^{-}), with (the function that characterises) the set of worlds in which the sentence is false; and the extension gap of S, notated $[\![S]\!]^{g,\#}$ (or simply $S^{\#}$), with (the function that characterises) the set of worlds in which the sentence is undefined.

Hence:

(13)
a. $[\![S]\!]^{g,+} = \lambda w\,.\,[\![S]\!]^{w,g} = T$ [Positive extension of S]
b. $[\![S]\!]^{g,-} = \lambda w\,.\,[\![S]\!]^{w,g} = F$ [Negative extension of S]
c. $[\![S]\!]^{g,\#} = \lambda w\,.\,[\![S]\!]^{w,g} = \#$ [Extension gap of S]

Let's now go back to (11). μ_{count} is the measure function that counts objects and/or measures quantities:[93] for example, $\mu_{\text{count}}(w, \iota x\,.\,\text{piggy-bank}'(w, x))$ stands for the number of coins inside the piggy bank in question in w. g(ρ) is the value of the granularity parameter at a context: what g(ρ) does is to shift upwards the falsity threshold of the adjective. If g(ρ) = 0, then 'empty$_\rho$' has a positive and a negative extension (i.e. $[\![\text{empty}_\rho]\!]^{w,\,g[0/\rho],\,+} = \lambda x\,.\,\mu_{\text{count}}(w, x) = 0$ and $[\![\text{empty}_\rho]\!]^{w,\,g[0/\rho],\,-} = \lambda x\,.\,\mu_{\text{count}}(w, x) > 0$, respectively), but no extension gap (i.e. $[\![\text{empty}_\rho]\!]^{w,\,g[10/\rho],\,\#} = \lambda x\,.\,0 < \mu_{\text{count}}(w, x) \leq 0$); if g(ρ) > 0, 'empty$_\rho$' has a positive and a negative extension, as well as an extension gap: for example, if g(ρ) = 10, then $[\![\text{empty}_\rho]\!]^{w,\,g[10/\rho],\,+} = \lambda x\,.\,\mu_{\text{count}}(w, x) = 0$, $[\![\text{empty}_\rho]\!]^{w,\,g[10/\rho],\,-} = \lambda x\,.\,\mu_{\text{count}}(w, x) > 10$, and $[\![\text{empty}_\rho]\!]^{w,\,g[10/\rho],\,\#} = \lambda x\,.\,0 < \mu_{\text{count}}(w, x) \leq 10$. Notice that, in the proposed semantics, context—represented by the assignment function g—intrudes upon the negative extension and the extension gap of the total adjective, but not upon its positive extension: indeed, $[\![\text{empty}_\rho]\!]^{w,\,g,\,+} = \lambda x\,.\,\mu_{\text{count}}(w, x) = 0$.

---

[93] Here I am simplifying matters for the sake of exposition. First, what μ_{count} should do, I think, is to measure both the amount of stuff and the capacity of the thing that contains the stuff, and calculate the ratio of the former value to the latter one; hence 'the piggy bank is empty' is true if the number of coins in the piggy bank divided by the total capacity of the piggy bank is 0; conversely, 'the piggy bank is full' is true if the number of coins in the piggy bank divided by the total capacity of the piggy bank is 1. Such a measure function enables us to treat both *empty* and *full* as total adjectives: whereas $[\![\text{empty}]\!]^{w,\,g,\,+} = \lambda d\,.\,\lambda x\,.\,\mu_{\text{count}}(w, x) = 0$, $[\![\text{full}]\!]^{w,\,g,\,+} = \lambda d\,.\,\lambda x\,.\,\mu_{\text{count}}(w, x) = 1$. Second, *empty* and *full*, unlike all the other absolute adjectives, appear to take an extra argument: for example, 'the theatre is empty' can be true of a situation in which there is no spectators in the theatre but there are nonetheless actors on the stage; however, it is also possible to think of a context in which that very same sentence is be false of the said situation (imagine that there's a fire in the theatre: in such a context, 'the theatre is empty' (if uttered by a policeman after inspecting the building) is false if, despite there being no spectators, there are actors on stage). Thus, an extra ('domain') argument should be incorporated somewhere in the semantics of *full* and *empty* (i.e. *empty* means *empty* <u>of something</u>, and *full* means *full* <u>of something</u>).

The value of ρ depends on the context; more specifically, on the I (the issue) and P (the common ground), which I will refer to as the I/P. But how is the value of ρ determined, exactly? The idea I have is the following: ρ takes the unique value that enforces a perfect alignment between the meaning of the sentence and the I/P. In APPENDIX I, I put forward an alignment condition, where the notion of 'perfect alignment' becomes transparent; here I shall do without it and keep the discussion at an intuitive level. Let's imagine that the I/P in operation is such that it does not matter whether there are no coins in the piggy bank or whether there are between 1 and 5 coins; technically, this means that the propositions *there are no coins in the piggy bank* (call it p) and *the number of coins in the piggy bank is greater than 0 and less than or equal to 5* (call it q) are equivalent relative to the I/P being considered (i.e. $p \approx_{I,\mathrm{P}} q$)—see Def. I below (= Def. (V), Chapter V).

(I) Equivalence between propositions rel. to a partition and a probability distribution (v. 1)

Notation:

P is a probability distribution over Ω
I is a partition on Ω: $I = \{i_{\mathrm{i}},\ldots,i_{\mathrm{n}}\}$
i is partition cell: $i \in I$
p and q are propositions: $p, q \in \wp(\Omega)$
$\approx_{I,\mathrm{P}}$ is an equivalence relation between propositions parametrised by I and P

Definition:

$p \approx_{I,\mathrm{P}} q$ if, $\forall i \in I$,

$$(\mathrm{P}(i \mid p) - \mathrm{P}(i)) \times (\mathrm{P}(i \mid q) - \mathrm{P}(i)) > 0 \quad \vee \; \mathrm{P}(i) = \mathrm{P}(i \mid p) = \mathrm{P}(i \mid q),$$

and for any two cells i and $ii \in I$,

$$\mathrm{P}(i \mid p) > \mathrm{P}(ii \mid q) \;\leftrightarrow\; \mathrm{P}(i \mid q) > \mathrm{P}(ii \mid q).$$

Let's also assume that the I/P is such that it does matter whether there are between 1 and 5 coins in the piggy bank or whether there are more than 5 coins; that is, the propositions *the number of coins in the piggy bank is greater than 0 and less than or equal to 5* and *there are more than 5 coins in the piggy bank* (call it r) are not equivalent relative to the I/P (i.e. $q \not\approx_{I,\mathrm{P}} r$). The I/P being considered, thus, has the following two properties:

(14) a. *there are no coins in the piggy bank* (p) $\boldsymbol{\approx_{I,\mathrm{P}}}$ *the number of coins in the piggy bank is greater than 0 and less than or equal to 5* (q)
b. *the number of coins in the piggy bank is greater than 0 and less than or equal to 5* (q) $\boldsymbol{\not\approx_{I,\mathrm{P}}}$ *there're more than 5 coins in the piggy bank* (r)

If, against such an I/P, the sentence 'the piggy bank is empty' were to be uttered: what value should ρ be assigned by g? The answer is 5, because 5 is the unique value that enforces a perfect alignment between the (trivalent) meaning of 'the piggy bank is empty' and the I/P stipulated in (14) as illustrated in (15) below.

(15) $[\![\text{the piggy bank is empty}_\rho]\!]^{+,\, g[5/\rho]} =_* \{w : \text{the p.bank has 0 coins in } w\} = p$
$[\![\text{the piggy bank is empty}_\rho]\!]^{-,\, g[5/\rho]} =_* \{w : \text{the p.bank has more than 5 coins in } w\} = r$
$[\![\text{the piggy bank is empty}_\rho]\!]^{\#,\, g[5/\rho]} =_* \{w : \text{the p.bank has (0, 5] coins in } w\} = q$

*Identifying sets with their characteristic functions.

Indeed, if g(ρ) = 5, 'the piggy bank is empty' will be false as soon as there are more than 5 coins in the piggy bank; and, given the stipulated I/P, non-emptiness *de facto* begins as soon as there are more than 5 coins in the piggy bank (i.e. between 0 and 5 coins, as far as I/P is concerned, the piggy bank is empty: $p \approx_{I,\mathrm{P}} q$).

As illustrated in (15), 'the piggy bank is empty' is predicted to have a gap (against the stipulated I/P); this gap, however, will not be seen. Why so? The explanation goes as follows. By assumption, $p \approx_{I,\mathrm{P}} q$ holds—cf. (14); as a result of this, $S^+ \approx_{I,\mathrm{P}} S^\#$ holds too (because $p = S^+ = \{w$: the piggy bank has 0 coins in $w\}$ and $q = S^\# = \{w$: the piggy bank has (0, 5] coins in $w\}$). When $S^+ \approx_{I,\mathrm{P}} S^\#$ holds, as discussed in the previous chapter, PNM's update mechanism outputs a non-maximal interpretation—that is, for any $p \in \wp(\Omega)$, $\mathrm{P}_3(p) = \sum_{i \in I} \mathrm{P}_1(p|i) \times \mathrm{P}_2(i)$—see Def. II below (= Def. VI, Chapter V). But if 'the piggy bank is empty' is interpreted non-maximally when g(ρ) = 5, then $\{w$: the piggy bank has (0, 5] coins in $w\}$, the extension gap of the sentence, will be invisible: the worlds in the extension gap will all be 'alive' in the final distribution (due to non-maximality) and, as a result, 'the piggy bank is empty' will be perceived to be true enough (as opposed to undefined) relative to such worlds.

(II) PNM's update algorithm

Let I be a partition on Ω, i a partition cell, S a natural language sentence (S^+ = S's positive extension; S^- = S's negative extension; $S^\#$ = S's extension gap), P_0 a prior probability distribution over Ω, and $\approx_{I,\mathrm{P}_0}$ an equivalence relation between propositions. Update with S proceeds as follows. For any $p \in \wp(\Omega)$:

(i) $\mathbf{P_1}(\boldsymbol{p}) = \mathbf{P_0}(\boldsymbol{p} \mid \boldsymbol{\Omega} \setminus \mathbf{S^-})$

(ii) $\mathbf{P_2}(\boldsymbol{p}) = \mathbf{P_1}(\boldsymbol{p} \mid \mathbf{S^+})$

(iii) If $S^+ \approx_{I,\mathrm{P}_0} S^\#$, then $\mathbf{P_3}(\boldsymbol{p}) = \sum_{i \in I} \mathbf{P_1}(\boldsymbol{p} \mid \boldsymbol{i}) \times \mathbf{P_2}(\boldsymbol{i})$. Otherwise, $\mathbf{P_3}(\boldsymbol{p}) = \mathbf{P_2}(\boldsymbol{p})$.

P_3 is the final interpretation.

3.1.2 Negation

Consider (16) below (cf. Chapter III, § 3.4):

(16) Context: *Charlotte and Gerard rented a flat for the weekend in the outskirts of London. The flat was advertised as 'having everything one could wish for' from a full-packed fridge to kitchenware, from TV to board games, etc. However, much to their surprise, the couple found that there was almost nothing to eat in the fridge, no board games, one plate only (and no cutlery). In this context, Gerard points out*:

(Fortunately,) the bookcase in the living room isn't empty.

If there is only one book in the bookcase, (16) is likely to strike Charlotte as being false; indeed, in the stipulated context, (16) appears to mean something like *there is a reasonable number of books in the living room bookcase*. Let's first derive the extensional meaning of (16) on the basis of (11) and the obvious trivalent analogue of propositional negation.

(17)

a. $$[\![\text{the bookcase is empty}_\rho]\!]^{w,\,g} = \begin{cases} \text{T iff } \mu_{\text{count}}\,(w, \iota x \,.\, \text{bookcase}'(w, x)) = 0 \\ \text{F iff } \mu_{\text{count}}\,(w, \iota x \,.\, \text{bookcase}'(w, x)) > \text{g}(\rho) \\ \#\text{ iff } 0 < \mu_{\text{count}}\,(w, \iota x \,.\, \text{bookcase}'(w, x)) \leq \text{g}(\rho \end{cases}$$

b. $$[\![\text{not S}]\!]^{w,\,g} = \begin{cases} \text{T iff } [\![\text{S}]\!]^{w,\,g} = \text{F} \\ \text{F iff } [\![\text{S}]\!]^{w,\,g} = \text{T} \\ \#\text{ iff } [\![\text{S}]\!]^{w,\,g} = \# \end{cases}$$

c. $$[\![\text{the bookcase is}\mathbf{n't}\text{ empty}_\rho]\!]^{w,g} = \begin{cases} \text{T iff } \mu_{\text{count}}\,(w, \iota x \,.\, \text{bookcase}'(w, x)) > \text{g}(\rho) \\ \text{F iff } \mu_{\text{count}}\,(w, \iota x \,.\, \text{bookcase}'(w, x)) = 0 \\ \#\text{ iff } \mu_{\text{count}}\, 0 < (w, \iota x \,.\, \text{bookcase}'(w, x)) \leq \text{g}(\rho) \end{cases}$$

Given the context stipulated in (16), it seems reasonable to assume that the *I*/P in operation is one that cares whether there are, say, more than 30 books in the bookcase, but does not care whether there are 0 or between 1 and 30 books; these assumptions are laid out in (18).

(18) a. *there are no books in bookcase* (*s*) $\approx_{\boldsymbol{I},\mathbf{P}}$ *the number of books in the bookcase is greater than 0 and less than or equal to 30* (*t*)
b. *the number of books in the bookcase is greater than 0 and less than or equal to 30* (*t*) $\not\approx_{\boldsymbol{I},\mathbf{P}}$ *there're more than 30 books in the bookcase* (*u*)

If (16) is uttered against such an *I*/P, what value should ρ take? The answer is obvious: g(ρ) is going to equal 30. Indeed, g(ρ) = 30 secures a perfect alignment between the *I*/P and the meaning of (16): the *I*/P starts to 'visualise' books as soon as there are 31 books or more, and, if g(ρ) = 30, (16) will be predicted to be true as soon as there are 31 books or more. Thus, (16), against the stipulated *I*/P, will have the meaning in (19).

(19) $[\![$the bookcase isn't empty$_\rho]\!]^{+,\, g[30/\rho]}$ $=_*$ $\{w$: the b.case has more than 30 books in $w\}$ $= u$
$[\![$the bookcase isn't empty$_\rho]\!]^{-,\, g[30/\rho]}$ $=_*$ $\{w$: the b.case has 0 books in $w\}$ $= s$
$[\![$the bookcase isn't empty$_\rho]\!]^{\#,\, g[30/\rho]}$ $=_*$ $\{w$: the b.case has (0, 30] books in $w\}$ $= t$

*Identifying sets with their characteristic functions.

What (pragmatic) interpretation does PNM's update algorithm predict given (19)? It predicts P_3, the final interpretation, to be equal to P_2. Let's see why. By assumption, the *I*/P in operation is one relative to which the propositions *there are more than 30 books in the bookcase* (or *u*) and *the number of books in the bookcase is greater than 0 and less than or equal to 30* (or *t*) are not equivalent ($u \not\approx_{I,\mathrm{P}} t$). If $u \not\approx_{I,\mathrm{P}} t$, then $S^+ \not\approx_{I,\mathrm{P}_0} S^\#$: this is because $u = S^+ = \{w$: the bookcase has more than 30 books in $w\}$ and $t = S^\# = \{w$: the bookcase has (0, 30] books in $w\}$. When $S^+ \approx_{I,\mathrm{P}} S^\#$ fails to hold, as in this case, P_3, the final interpretation that PNM's update mechanism outputs, equals P_2 (that is, PNM outputs a maximal interpretation)—cf. Def. II. P_2 is obtained by updating the common ground with S^+: thus, to obtain the final interpretation, all that needs to be done is to update the common ground with $\{w$: the bookcase has more than 30 books in $w\}$. This seems to be the right result: (16), in the stipulated context, appears to be interpreted as meaning *there is a reasonable number of books* (more than 30, by assumption) *in the bookcase.*

It should be noted straight away that there is something quite not right here: (16), in the stipulated context, is interpreted maximally —that is, $P_3 = P_2$—, yet it is not interpreted precisely (i.e. the precise interpretation of 'the bookcase isn't empty' is, of course, *there is at least 1 book in the bookcase*). But this makes no sense because 'maximal interpretation' and 'precise interpretation' are supposed to mean the same thing. This is a terminological accident, which follows from having identified the two possible interpretations of PNM's update rule with the effect that those interpretations have on plural sentences. Indeed, when the final interpretation is equal to $\sum_{i \in I} \mathrm{P}_1(p|i) \times \mathrm{P}_2(i)$, for any $p \in \wp(\Omega)$, the interpretation of an unquantified plural sentence such as 'the professors smiled' or 'the professors didn't smile' is non-maximal or imprecise and, when the final interpretation is equal to P_2, it is maximal. However, in sentences such as 'X is empty' or 'X isn't empty', this no longer holds. The facts are as follows: the final interpretation of a sentence of the form 'X is TA' will be equal to P_2 if $g(\rho) = 0$ and equal to $\sum_{i \in I} \mathrm{P}_1(p|i) \times \mathrm{P}_2(i)$, for any $p \in \wp(\Omega)$, if $g(\rho) > 0$; this interpretation will be precise/maximal if $g(\rho) = 0$ and non-maximal/imprecise if $g(\rho) > 0$; the final interpretation of a sentence of the form 'X is *not* TA' will always be equal to P_2;[94] this interpretation will be maximal/precise if $g(\rho) = 0$ and non-maximal/imprecise if $g(\rho) > 0$.

[94] If $g(\rho) = 0$, then $S^+ \approx_{I,\mathrm{P}} S^\#$ will not hold (because $S^\# = \emptyset$) and, as a result, $P_3 = P_2$; if $g(\rho) > 0$, $S^+ \approx_{I,\mathrm{P}_0} S^\#$ will not hold either: this is because an *I*/P that takes the value of ρ to, for example, 30 is an *I*/P that cares whether 1-30, or whether 31 or more. See APPENDIX I for an explicit account of how the value of ρ is determined.

Though a sentence of the form 'X is *not* TA' is predicted to have a gap when g(p) > 0, such a gap will not be visible. Indeed, the final interpretation (P_3) of 'X is *not* TA', as just explained, will be equal to P_2; this means that, in P_3, the extension gap of 'X is *not* TA' is going to have probability zero. 'X is *not* TA', therefore, if evaluated relative to a world in its extension gap, is going to be perceived as false (and not as undefined). The result appears to generalise: any *I*/P such that g(ρ) > 0 is an *I*/P that induces a *semantic* gap: but PNM, given any such *I*/P, pulls down a pragmatic veil, namely, (non-)maximality, that makes the gap invisible.

3.1.3 Partial adjectives

Partial adjectives (PAs) are the complementary antonyms of total adjectives (TAs): in the trivalent semantics proposed here, what this means is that the positive and negative extensions of a total adjective α are the negative and positive extensions of α's partial counterpart, respectively—cf. (10) with (20) below.

(20)
$$[\![\mathrm{PA}]\!]^{w,\,g} \quad = \quad \lambda d_d\,.\,\lambda x_e\,. \begin{cases} \text{T iff } \mu_\delta\,(w, x) > d \\ \text{F iff } \mu_\delta\,(w, x) = 0 \\ \#\text{ iff } 0 < \mu_\delta\,(w, x) \leq d \end{cases}$$

Let's illustrate this with the total/partial pair *straight*/*bent*, where μ_{cur} is the measure function that measures the degree of curvature of an object.

(21)

a.
$$[\![\text{straight}]\!]^{w,\,g} \quad = \quad \lambda d\,.\,\lambda x\,. \begin{cases} \text{T iff } \mu_{\text{cur}}\,(w, x) = 0 \\ \text{F iff } \mu_{\text{cur}}\,(w, x) > d \\ \#\text{ iff } 0 < \mu_{\text{cur}}\,(w, x) \leq d \end{cases}$$

b.
$$[\![\text{bent}]\!]^{w,\,g} \quad = \quad \lambda d\,.\,\lambda x\,. \begin{cases} \text{T iff } \mu_{\text{cur}}\,(w, x) > d \\ \text{F iff } \mu_{\text{cur}}\,(w, x) = 0 \\ \#\text{ iff } 0 < \mu_{\text{cur}}\,(w, x) \leq d \end{cases}$$

Given the semantics in (20), sentences of the form 'X is PA' are predicted to be interpreted just like sentences of the form 'X is *not* TA' (see the discussion on 'the bookcase isn't empty' in § 3.1.2).

3.2 Vagueness again

In § 2, it was argued that absolute adjectives, unlike relative ones, are not vague. This is perhaps not entirely accurate. Let's consider, once again, example (16), given as (22) below:

(22) Context: Charlotte and Gerard rented a flat for the weekend in the outskirts of London. The flat was advertised as 'having everything one could wish for' from a full-packed fridge to

kitchenware, from TV to board games, etc. However, much to their surprise, the couple found that there was almost nothing to eat in the fridge, no board games, one plate only (and no cutlery). In this context, Gerard points out:

(Fortunately,) the bookcase in the living room isn't empty.

In § 3.1.2, I stipulated that in (16)/(22), the *I*/P in operation is such that it cares whether there are more than 30 books in the bookcase, but does not care whether there are 0 or between 1 and 30 books; given this *I*/P, the trivalent meaning of (22) takes the following shape:

(23) $[\![\text{the bookcase isn't empty}_\rho]\!]^{+,\, g[30/\rho]} =^{*} \{w : \text{the b.case has more than 30 books in } w\}$
$[\![\text{the bookcase isn't empty}_\rho]\!]^{-,\, g[30/\rho]} =^{*} \{w : \text{the b.case has 0 books in } w\}$
$[\![\text{the bookcase isn't empty}_\rho]\!]^{\#,\, g[30/\rho]} =^{*} \{w : \text{the b.case has (0, 30] books in in } w\}$

*Identifying sets with their characteristic functions.

As discussed in § 3.1.2, updating the common ground with 'the bookcase in the living room isn't empty' amounts to updating the common ground with its truth-conditions ($P_3 = P_2$); given the stipulated *I*/P, it amounts to updating the common ground with $\{w$: the bookcase has more than 30 books in $w\}$. Now, it is clear that (22), if uttered in the context stipulated, is never interpreted in such a way: if Charlotte were to find out there are in fact 29 books in the bookcase, she would surely not accuse Gerard of having said something false. Indeed, there seems to be no fact of the matter as to what precise number of books there must be in the bookcase for Gerard's utterance to be accepted as a true description of the situation at hand; in other words, there seems to be vagueness about just how much (non-zero) imprecision is intended.[95]

In the case of relative adjectives, the facts are not quite the same; in degree semantics, the simplest semantic treatment that can be given to a relative adjective (for example, *tall*) has the following form:

(24)
$$[\![\text{tall}_\theta]\!]^{w,g} = \lambda x\,.\begin{cases} \text{T iff } \mu_{\text{height}}(w, x) > g(\theta) \\ \text{F iff } \mu_{\text{height}}(w, x) \leq g(\theta) \\ \#\text{ otherwise (i.e. never)} \end{cases}$$

According to (24), and in plain terms, someone is tall if s/he is above a context-dependent threshold θ. This threshold θ is, just like the granularity parameter ρ, an object-language variable that ranges over degrees.

[95] Thanks to Manuel Križ for discussion on this point.

As discussed in Chapter III, relative adjectives and, derivatively, sentences that contain them (e.g. 'John is tall'), lack sharp boundaries: there does not seem to be a precise point at which John stops being tall and begins to be non-tall; if one accepts the semantics in (24), a sharper diagnosis becomes available: what there is, as a matter of fact, is uncertainty about the value of θ (see, for example, Lassiter and Goodman 2013, and Qing and Franke 2014); as a result of this, the [+]/[−] boundary of relative adjectives—which, according to (24), is the only semantic boundary that these adjectives have—is flooded with fuzziness.

Let's now compare (24) with (25) below:

(25)
$$[\![\text{empty}_{\rho}]\!]^{w,g} \quad = \quad \lambda x_e\,.\begin{cases} \text{T iff } \mu_{\text{count}}(w, x) = 0 \\ \text{F iff } \mu_{\text{count}}(w, x) > g(\rho) \\ \text{\# iff } 0 < \mu_{\text{count}}(w, x) \leq g(\rho) \end{cases}$$

In (25), if there is uncertainty about the precise value of ρ, then the [+ or #]/[−] boundary is expected to be fuzzy (just like the [+]/[−] boundary is expected to be fuzzy in relative adjectives). However, in (25), there is a semantic boundary, namely, the [+]/[# or −] boundary, which does not depend on the value that ρ takes and hence remains sharp; indeed, given (25), 'the bookcase is empty' is true if there are no books in the bookcase, and non-true (undefined or false) otherwise; conversely, 'the bookcase isn't empty' is false if there are no books in the bookcase, and non-false (undefined or true) otherwise.

Thus, at this level of analysis, it does not make much sense to talk about an adjective being or not being vague: there is vagueness (or uncertainty) about the precise value that a given contextual variable takes. In relative adjectives, this vagueness blurs the [+]/[−] boundary (this accounts for why these adjectives induce soritical reasoning); by contrast, in total adjectives, it blurs the [+ or #]/[−] boundary but leaves the [+]/[# or −] untouched, whereas in partial adjectives, it blurs the [+]/[# or −] boundary but leaves the [+ or #]/[−] untouched (having one sharp boundary appears to be enough to prevent soritical reasoning from being triggered). The upshot of this is that, in a sentence of the form 'John is (not) tall', there is going to be vagueness at the truth-falsity boundary; by contrast, in a sentence of the form 'X is (not) empty', as soon as g(ρ) > 0, there is going to be vagueness about how just big the homogeneity gap is and, as a result, about how much deviation from a precise interpretation is pragmatically admissible.

Now, why is there uncertainty about the value or ρ (or about the value of θ)? This is a hard question. Presumably, there is always uncertainty about what the common ground is; in addition, there is often uncertainty about what the issue is. Even if, for each conceivable *I*/P pair, there is a unique value that ρ takes (as I have proposed), in the real world—for example, in the context stipulated in (29)—, a large number of *I*/P pairs is likely to be under consideration.

Everything that has been said in Chapter III, § 3.3, about imprecision and how it differs from vagueness, stands, as far as I can tell; if vagueness did not exist, then, the semantic interpretation of (22) would be exactly the one in (23); PNM's update algorithm would then generate an imprecise (yet non-vague) interpretation, because $g(\rho) > 0$; this interpretation would be the same interpretation that would be obtained if the common ground were to be updated with the sentence 'there are more than 30 books in the bookcase'. Once vagueness is incorporated into the system (and I want to remain silent on the question of how vagueness should be incorporated into the system), then imprecision itself becomes vague.

Now, there is a puzzle here; according to the account just given, imprecision is intertwined with vagueness because there is always uncertainty about what the value of ρ is; but, if this is the case, why is it that one often has complete certitude that $g(\rho) = 0$? In other words, why is it that when a sentence is interpreted precisely, there is no uncertainty about whether the sentence was intended to be interpreted precisely? The answer cannot be: 'because, at $g(\rho) = 0$, the (homogeneity) gap effectively closes': that would be conflating vagueness with homogeneity (cf. Chapter IV, § 2.2.2). Homogeneity is a pre-requisite for imprecision, but not for vagueness: after all, relative adjectives are not homogenous and are nonetheless vague.

This is a difficult puzzle and, unfortunately, I can hardly do more than describing what the facts themselves suggest. It appears to be the case that the difference between *none* (i.e. 0) and *some* (i.e. > 0) is cognitively much more salient than the difference between n and n+1 (where $n \neq 0$). Thus, in contexts in which it matters whether, for example, no one or someone is in the theatre (for example, if there is a fire in the theatre), no uncertainty arises as to what the value of ρ should be: according to all *I*/P pairs under consideration, $g(\rho) = 0$; as a result, in such contexts, 'the theatre is empty' is interpreted precisely. However, in contexts in which it does not matter whether there are no people or whether there are n amount of people, but it does matter whether there are n or more than n people, uncertainty would appear to kick in: in such contexts, not all *I*/P pairs under consideration agree on what the value of ρ should be.

If the sort of picture I have sketched above is remotely accurate, then another puzzle needs explaining: why is it that relative adjectives never have non-vague interpretations (in natural contexts; cf. Chapter III, 3.3). Indeed, given a (natural) context in which the difference between *none* and *some* matters, one would expect the adjective 'hairy', for example, to get to mean *has at least 1 hair*, or the adjective 'rich', *possesses at least 1 thing*. In other words, it should be possible for (positive) relative adjectives, in contexts in which the difference between *none* and *some* matters, to effectively function as partial adjectives.

For this second puzzle, I do have a tentative answer: relative adjectives' threshold-value space is restricted; in particular, 0 is not a possible threshold value. This may seem like an ad-hoc stipulation, but I do not think that it is. There are independent reasons to think that the threshold associated with relative adjectives is sensitive to semantic restrictions: the threshold of *hot*, for example, cannot take any value, only values that are greater than those values that the threshold associated with *warm* can take. If I am right in thinking that relative adjectives' threshold-value space excludes 0, then the puzzle dissolves: relative adjectives, because their threshold argument can never be 0, are insensitive to contexts in which the difference between *none* and *some* matters and, as a result, are indifferent to those *I*/P configurations in which non-vagueness can be seized.

3.3 Strict behaviour in contradiction test

As discussed, absolute adjectives, just like definite plurals, exhibit *strict behaviour in contradiction test* (or SBCT); consider, for example, (26):

(26) # The theatre was empty, and only two people were sitting in the front.

(26) feels very much like a contradiction, and the semantics just proposed predict (26) to be one: more precisely, it predicts the positive extension of (26) to be empty.[96] Indeed, according to (10), 'X is empty' is true if there are no things in X, irrespective of which value ρ takes. The upshot of this is that that (26) can never be true: as soon as the first conjunct is true (as soon as there are no people in the theatre), the second conjunct is false. Because (26) can never be true, it is bound not to be p-relevant (cf. Chapter V, § 2.3); furthermore, since $[\![(26)]\!]^{+} = \emptyset$, if (26) were to be used to update the common ground, PNM's update algorithm will crash (because step (ii) of the algorithm would be undefined; cf. § Def. II).

(10), it should be noted, could be re-worked in a bivalent format: indeed, instead of shifting the [+ or #]/[−] boundary, ρ could shift, as illustrated in (27), the [+]/[−] boundary, and thus semantically generate imprecise *readings* without the need of having a global pragmatic mechanism such as PNM.

(27) $$[\![\text{empty}_{\rho}]\!]^{w,g} \quad = \quad \lambda x \,.\, \begin{cases} \text{T iff } \mu_{\text{count}}\,(w, x) \leq \text{g}(\rho) \\ \text{F iff } \mu_{\text{count}}\,(w, x) > \text{g}(\rho) \\ \text{\# otherwise (i.e. never)} \end{cases}$$

[96] In a trivalent system, there are two obvious definitions of 'contradiction': either a proposition that is false in every interpretation, or a proposition that is either false or undefined in every interpretation. If the latter definition is adopted, then (26) is a contradiction.

Let's assume that 'the theatre was empty' is uttered against an *I*/P that does not care whether there are no people in the theatre or whether there are between 1 and 10. Against such an *I*/P, given (27), 'the theatre was empty' should have the following meaning:

(28)
$$[\![\text{the theatre was empty}_\rho]\!]^{w,\, g[10/\rho]} = \begin{cases} \text{T iff } \mu_{\text{count}}\,(w, \iota x \,.\, \text{theatre}'(w, x)) \leq 10 \\ \text{F iff } \mu_{\text{count}}\,(w, \iota x \,.\, \text{theatre}'(w, x)) > 10 \\ \#\text{ otherwise (i.e. never)} \end{cases}$$

Prima facie, (28) looks good: if there happened to be 11 or more people in the theatre, 'the theatre is empty' would be false, whereas if there happened to be, say, only 5 people in there, then the sentence would be true. Now, if imprecision were to originate at the lexical level as in (27), and then percolate up to the semantics of the complex expression as in (28), then the contradictoriness of (26), repeated below as (29)a, becomes a complete mystery.

(29) a. # The theatre was empty, and only 2 people were sitting in the front.
b. There were between 0 and 10 people in the theatre, and only 2 people were sitting in the front.

Indeed, if the meaning of 'empty$_\rho$' was the one in (27), (29)a would be predicted, if uttered against an *I*/P that does not care whether there were no people in the theatre or whether there were between 1 and 10, to have the same reading as (29)b (because g, given such an *I*/P, would be expected to assign ρ the value 10). The fact that (29)a is a contradiction (and does not get to mean (29)b irrespective of context) suggests that the local/lexical account of imprecision proposed in (27) is not on the right track.

To conclude: the contradictoriness of (26)/(29)a indicates that (im)precision (or (non-)maximality), as stipulated in PNM, must be a global mechanism: that is, a mechanism that kicks in after the compositional semantics has computed a semantic value for the sentence.

4 COMPLETELY AND SLIGHTLY

4.1 Completely, slightly, and the total/partial distinction

One of the greatest difficulties when it comes to formalising the meaning of gradable adjectives is to account for the (seemingly idiosyncratic) distribution of degree modifiers such as *completely* (or *perfectly*, or *totally*) and *slightly*. The general observation, as illustrated in (30), is that *completely* can modify total adjectives (but not relative or partial adjectives), whereas *slightly* can modify partial adjectives (but not relative or total adjectives).

(30) a. *Completely* straight / *bent / *tall
b. *Slightly* bent / *straight / *tall

In what follows, I will discuss what has come to become the standard approach to explaining the distribution facts of these degree modifiers: the scale structure approach.

4.1.1 The scale structure approach

In the tradition of degree semantics, gradable adjectives are analysed as having a measure function somewhere in their meaning. Measure functions, as discussed, associate entities with ordered values on a scale. Classically, scales are formalised as triples of the form $\langle \mathcal{D}, R, \delta \rangle$, where δ is the dimension of measurements (height, price, likelihood, age, etc.), $\mathcal{D}$ is a set of degrees, and R either a total or partial order on $\mathcal{D}$ (Bartsch and Vennemann 1972; Bierwisch 1989).[97] For short, I will refer to the set of degrees associated with the dimension δ as $\mathcal{D}_{\delta}$, and to the scale proper as $\mathcal{D}_{\delta/R}$.

Kennedy and McNally (2005) and Kennedy (2007) have made the claim that adjectival scales come in four types, i.e.: open, upper-closed, lower-closed, and fully closed (details to come); furthermore, these authors have proposed that that relative adjectives are associated with open scales, total adjectives with either upper-closed or fully-closed scales, and partial adjectives with lower-closed scales. This account, it should be noted, is at odds with the analysis that I have proposed in § 3.1, in which the scale, except for the dimension of measurement, is always the same.

For the sake of exposition, let's assume that degrees are values that are isomorphic to the real numbers between 0 and 1; depending on whether 0 and/or 1 are members of $\mathcal{D}$, then four scale types can be derived. If neither 0 nor 1 are elements of $\mathcal{D}$, then the scale is is said to be *open* (or *unbounded*); if 0 or 1, but not both, are elements of $\mathcal{D}$, then the scale can be either *upper-closed* or *lower-closed*, depending on whether the closed endpoint stands for the maximal degree of the measured property or for the absence of it; finally, if both 0 and 1 are elements of $\mathcal{D}$, then the scale is *fully-closed*.

Figure 1. 4-scale typology.[98]

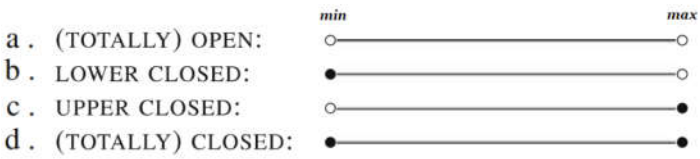

[97] In § 3.1, I do not specify an ordering relationship (and assume that the relevant order is the natural order of the real numbers).
[98] From Kennedy (2007).

Let's first consider the case of a relative adjective such as *tall*; according to the scale structure account, $\mathcal{D}_{\text{height}/R}$, the scale on which height is measured, lacks endpoints: in other words, 0 and 1 are not members of $\mathcal{D}_{\text{height}}$. The relevant intuition appears to be this: the height scale has no upper-limit (a thing can always be taller than it is) nor does it have a lower limit corresponding to the absence of height (if a thing had no height, it would not be a thing). By contrast, *straight* (a total adjective) and *bent* (a partial adjective) are stipulated to map objects onto a scale which is closed at one end: the intuition here is that something is straight when it possesses complete straightness (or zero bentness) and bent as soon as it has a bentness score greater than zero. In this framework, the scale that *straight* and *bent* are associated with is the same except for the ordering relation: $\mathcal{D}_{\text{cur}/>}$, for *straight*, and $\mathcal{D}_{\text{cur}/<}$, for *bent*, where the former is an upper-closed scale (e.g. 0 is the maximal point of straightness) and the latter a lower-closed scale (e.g. 0 is the absence of bentness).[99] Finally, there is a subset of total adjectives (*empty* and *full*, for example) that are stipulated to lexicalise scales that are closed at both ends. Why so? One of the relevant datapoints is that these adjectives can be combined with proportional modifiers such as *half* (whereas other total adjectives such as *straight*, partial adjectives such as *bent*, and relative adjectives such as *tall*, cannot).

Kennedy (2007), among others, has claimed that the (alleged) differences in scale structure among gradable adjectives are responsible for the distribution facts of degree modifiers such as of *completely* and *slightly*. If one assumes, as Kennedy does, that that the semantics of *slightly* requires there to be a lower limit in the adjectival scale, and that the semantics of *completely* requires there to be an upper limit, then *slightly* is predicted to felicitously combine with partial adjectives (but not with total adjectives) because the former, unlike the latter, lexicalise lower-closed scales; conversely, *completely* can combine with total adjectives (and not with partial ones), because the former, unlike the latter, lexicalise upper-closed scales. On this account, *completely* and *slightly* are both predicted to be bad with relative adjectives, because these adjectives are, by hypothesis, associated with open scales (that is, scales which are neither lower- nor upper-closed).

4.1.1.1 Problems

From the outset, it should be noted that it is not even clear that the stipulation that different adjectives are associated with different scales is able to provide a principled account for the distribution of *completely* and *slightly*. For example, and as noted in Solt (2012), *slightly*, according to the scale structure approach, is predicted to felicitously combine with adjectives that are associated with scales

[99] This technique comes out handy when the meaning of the gradable adjectives is analysed as denoting a measure function of type $\langle e,d \rangle$, as in Kennedy (2007). On Kennedy's account, if the ordering relation was not reversed from one member of the antonym pair to the other, then *straight* and *bent* (or *tall* and *short*) would have the exact same denotation. In the semantics proposed in §3.1, there is no need to ever reverse the scale's ordering relation.

that possess a lower limit; *empty* is, on this account, a fully-closed adjective and, as a result, possesses a lower limit: why is it, then, that 'slightly empty' is infelicitous (and does not get to mean *nearly full*)? I do not know of a good answer to this question.

But even if the scale structure approach could make sense of the distribution facts of *completely* and *slightly*, it would do so at a very high price: making incorrect predictions about the semantic behaviour of relative adjectives. It is clear that one can say 'this rod isn't *bent*' to describe a rod that has zero bentness; this is predicted on the scale structure account: *bent* is a partial adjective and, therefore, the scale that it lexicalises has a point that corresponds to zero bentness. This account, however, also predicts that it should not be possible to use a relative adjective α to describe an object that has a zero degree of the α-property; this is because relative adjectives are supposed to be associated with open scales, and these scales lack a point that stands for the absence of the relevant property.[100]

This prediction is not hard to test; consider, for example, the case of relative adjectives such as *rich* and *poor*. Intuition dictates that the scale associated with *rich* and *poor* is closed at one end: after all, there is such a thing as possessing nothing, and someone who possesses nothing (no clothes, no money, no properties, etc.) should be in the negative extension of *rich* (or in the positive extension of *poor*). If this scale, let's call it $\mathcal{D}_{\text{poss/R}}$ (where 'poss' stands for possessions), were to be open at both ends, then it would not be possible to describe someone who possesses nothing (i.e. someone who has literally 0 possessions) with the sentences 's/he isn't rich' or 's/he is poor'—however, as far as I can tell, it is possible. Thus, if one holds to the view that $\mathcal{D}_{\text{poss/R}}$ is open, then one cannot account for this straightforward linguistic fact; if, instead, one accepts that $\mathcal{D}_{\text{poss/R}}$ must be closed at one end, then, one no longer can appeal to considerations of scale structure to explain why 'completely rich/poor' and 'slightly rich/poor' are deviant.

The same argument can be made in connection with the (relative) adjective *hairy*.[101] The scale associated with this adjective, intuitively, is closed at one end: after all, there is such a thing as having no hair, and a person or animal who is hairless (a Sphynx cat, for example) should be in the negative extension of *hairy*. If the scale associated with *hairy* were to be open, then the sentence 's/he isn't hairy (at all)' could not be used to describe a person or animal without hair—however, as far as I can tell, it is possible.[102] This indicates that the scale that *hairy* lexicalises has a point that stands either for the

---

[100] It should be noticed that, the claim that relative adjectives lexicalise open scales (scales that have neither a 0 nor a 1 endpoint) is conceptually distinct from the claim that I have made in § 3.2: the claim that relative adjectives' threshold-value space excludes 0. The former is a claim about scale structure, the latter is not. It should also be noticed that the scale structure claim entails the threshold-value space claim (but not vice versa).

[101] Thanks to Benjamin Spector for thinking of this adjective.

[102] Of course, there is a slight oddness in describing, for example, a hairless cat with the sentence 'the cat isn't hairy'; if the cat really has no hair, there is a sense in which the speaker should have gone for the more informative 'the car is hairless'. More technically, 'the cat isn't hairy' appears to trigger the implicature that the cat is not totally hairless, hence the slight oddity of using 'the cat isn't hairy' to describe a hairless cat. However, this is an implicature and, as such, it can be cancelled:

absence of hair (or, equivalently, for the maximal degree of hairlessness); but, as soon as one accepts this (and one has to), the scale structure account crumbles: indeed, the account then predicts that *hairy* should felicitously combine with either *slightly* or *completely*—however both 'completely hairy' and 'slightly hairy' are semantically deviant.

Then there are (relative) adjectives such as *expensive* and *inexpensive* (or *cheap*); these adjectives, just like with *rich*/*poor*, or *hairy*, also seem to be associated with a scale that is closed at one end: the price/cost scale, at least intuitively, should be sensitive to the fact that certain things have zero price or cost. However, in this particular case, the linguistic facts appear (only appear) to suggest otherwise; for example, one cannot describe a bunch of items that are known to be free (e.g. a bunch of paper napkins on a bar table) with the sentence 'these are inexpensive' or 'these aren't expensive'. Can this datapoint be taken as evidence that the scale associated with (*in*)*expensive* is open? I do not think it can; there is a more plausible (or at least equally reasonable) explanation for the above datapoint: (*in*)*expensive* triggers the presupposition that the object that it applies to has a price.

(31) The 'Hey, wait a minute!' test (von Fintel 2004)
A: The napkins are inexpensive…
B: Hey, wait a minute! I didn't know that the napkins were up for sale.

If this is correct—and (31) appears to suggest that it is—, the oddness of 'these are(n't)inexpensive' (if uttered to describe a bunch of items that have no price) is due to presupposition failure: nothing needs to be stipulated about the nature of the scale that (*in*)*expensive* lexicalises.

Finally, there are (relative) adjectives such as *tall*; people seem to have the strong intuition that the height scale, as opposed to the price/cost scale, is open. I suspect that this is because it is easy to think of entities that do not have a price (e.g. a napkin in a bar) but very hard (if possible at all) to think of entities that have no height; having no height amounts to being immaterial and, being immaterial, at least intuitively, amounts to not being an entity. Hence, people appear to reason, why would the height scale have a 0 (an endpoint corresponding to the absence of height) if no entity will ever be sent to 0? This sort of intuitions, as far as I can tell, are of little value when it comes to shedding light on aspects of semantic structure.

Given these considerations, I remain thoroughly unconvinced that a distinction in terms of scale structure can be made between relative and absolute adjectives. First, it remains unclear whether such a distinction accounts for the distribution facts of *completely* and *slightly* (why is it that *slightly* doesn't

---

'the cat isn't hairy; in fact, it is totally hairless'. If *hairy* was associated with an open scale, 'the cat isn't hairy' would entail that the cat has some hair and, as a result, 'the cat isn't hairy; in fact, it is totally hairless' would be a contradiction.

go with *empty*, for example?); in addition, the claim that relative adjectives lexicalise open scales is arbitrary (why is it that the scale associated with *bent* has a 0 while the one associated with *hairy* doesn't?) but also empirically inadequate (it makes incorrect predictions, as discussed above). As things stand, I see no reason not to model gradable adjectives, including the so-called 'fully-closed' adjectives (see ftn. 93), using a measure function that maps entities onto a scale of the form $\langle \mathbb{R}_{+}, \delta \rangle$: a dense scale going from 0 to ∞.[103]

The distribution facts of *completely* and *slightly*, however, are clear, and need to be accounted for somehow: *slightly* felicitously combines with partial adjectives (*slightly* bent, *slightly* open, *slightly* impure), but not with relative adjectives (**slightly* poor/rich, **slightly* tall/short, **slightly* expensive/inexpensive); *completely*, in turn, felicitously combines with total adjectives (*completely* empty/full, *completely* straight, *completely* pure), but not with relative adjectives (**completely* poor/rich, **completely* tall/short, **completely* expensive/inexpensive). In the next section, I will provide a novel account of these facts.

4.1.2 The semantics of 'completely' *and* 'slightly'[104]

My proposal can be succinctly stated as follows: 'completely' and 'slightly' are homogeneity removers: 'completely' removes homogeneity in total adjectives, whereas 'slightly' removes homogeneity in partial adjectives. Let's get started then. Gradable adjectives (total, partial and relative adjectives) are of type $\langle d, \langle e,t \rangle \rangle$ (see (32)); 'completely' and 'slightly', in turn, are of type $\langle \langle d, \langle e,t \rangle \rangle, \langle e,t \rangle \rangle$. As before, μ_δ is a measure function along the dimension δ, x is a variable ranging over entities, and d is a variable ranging over degrees.

(32) a. <u>Total Adjectives</u>

$$[\![\text{TA}]\!]^{w,g} = \lambda d_d \,.\, \lambda x_e \,. \begin{cases} \text{T iff } \mu_\delta(w, x) = 0 \\ \text{F iff } \mu_\delta(w, x) > d \\ \text{\# iff } 0 < \mu_\delta(w, x) \leq d \end{cases}$$

b. <u>Partial Adjectives</u>

$$[\![\text{PA}]\!]^{w,g} = \lambda d_d \,.\, \lambda x_e \,. \begin{cases} \text{T iff } \mu_\delta(w, x) > d \\ \text{F iff } \mu_\delta(w, x) = 0 \\ \text{\# iff } 0 < \mu_\delta(w, x) \leq d \end{cases}$$

---

103 For arguments in support of the view that measurement scales are dense, see Fox and Hackl (2006) and Bale (2008). Though I am proposing that scales have this form $\langle \mathbb{R}_{+}, \delta \rangle$, as opposed to this form $\langle \mathbb{Q}_{+}, \delta \rangle$, I have no reason to prefer the former over the latter structure: in principle, and as far as I can tell, the rational (as opposed to the real) numbers could be used to build the scale.

104 Thanks to Benjamin Spector for invaluable help writing this section!

c. <u>(Positive) Relative Adjectives</u>

$$[\![\text{RA}]\!]^{w,g} \quad = \quad \lambda d_d \,.\, \lambda x_e \,.\, \begin{cases} \text{T iff } \mu_\delta(w, x) > d \\ \text{F iff } \mu_\delta(w, x) \leq d \\ \text{\# otherwise (i.e. never)} \end{cases}$$

The semantics of 'completely' and 'slightly' are as follows (extra notation: A stands for the denotation of a gradable adjective in w, either $[\![\text{TA}]\!]^{w,g}$, $[\![\text{PA}]\!]^{w,g}$, or $[\![\text{RA}]\!]^{w,g}$):

(33)

$$[\![\text{completely}]\!]^{w,g} = \lambda \mathbb{A}_{\langle d,\langle e,t\rangle\rangle} \,.\, \lambda x_e \,.\, \begin{cases} \text{T iff } \forall d\ (\mathbb{A}(d))\,(x) = \text{T} \\ \text{F iff } \exists d\ (\mathbb{A}(d))\,(x) = \text{F} \\ \text{\# otherwise} \end{cases}$$

(34)

$$[\![\text{slightly}]\!]^{w,g} \quad = \quad \lambda \mathbb{A}_{\langle d,\langle e,t\rangle\rangle} \,.\, \lambda x_e \,.\, \begin{cases} \text{T iff } \exists d\ (\mathbb{A}(d))\,(x) = \text{T} \\ \text{F iff } \forall d\ (\mathbb{A}(d))\,(x) = \text{F} \\ \text{\# otherwise} \end{cases}$$

The function in (33) takes two arguments, a gradable adjective meaning, notated A, and an entity x. For any A, and for any x, $[\![\text{completely}]\!]^{w,g}(\mathbb{A})(x)$ = T iff, for every d, $\mathbb{A}(d)$ maps x to T; and $[\![\text{completely}]\!]^{w,g}(\mathbb{A})(x)$ = F iff, for some d, $\mathbb{A}(d)$ maps x to F. The function in (34) does the reverse: for any A, and for any x, $[\![\text{slightly}]\!]^{w,g}(\mathbb{A})(x)$ = T iff, for some d, $\mathbb{A}(d)$ maps x to T; and $[\![\text{slightly}]\!]^{w,g}(\mathbb{A})(x)$ = F iff, for every d, $\mathbb{A}(d)$ maps x to F.

Let's see how the entries in (33) and (34) work if composed with a total and a partial adjective, respectively.

(35)

a. $$[\![\text{completely}]\!]^{w,g}\,([\![\text{straight}]\!]^{w,g}) \quad = \quad \lambda x \,.\, \begin{cases} \text{T iff } \forall d\ ([\![\text{straight}]\!]^{w,g}(d))\,(x) = \text{T} \\ \text{F iff } \exists d\ ([\![\text{straight}]\!]^{w,g}(d))\,(x) = \text{F} \\ \text{\# otherwise (i.e. never)} \end{cases}$$

b. $$[\![\text{slightly}]\!]^{w,g}\,([\![\text{bent}]\!]^{w,g}) \quad = \quad \lambda x \,.\, \begin{cases} \text{T iff } \exists d\ ([\![\text{bent}]\!]^{w,g}(d))\,(x) = \text{T} \\ \text{F iff } \forall d\ ([\![\text{bent}]\!]^{w,g}(d))\,(x) = \text{F} \\ \text{\# otherwise (i.e. never)} \end{cases}$$

Let's first consider (35)a. For any x, $[\![\text{completely straight}]\!]^{w,g}(x)$ = T iff, for every d, $[\![\text{straight}]\!]^{w,g}(d)$ maps x to T. As established in (32)a, $[\![\text{straight}]\!]^{w,g}(d)$ maps x to T iff x has 0 degree of curvature in w (irrespective of which value d ends up taking); thus, on the proposed semantics, $[\![\text{completely straight}]\!]^{w,g,+}$ and $[\![\text{straight}]\!]^{w,g,+}(d)$ are the same function. For any x, $[\![\text{completely straight}]\!]^{w,g}(x)$ = F iff, for some d, $[\![\text{straight}]\!]^{w,g}(d)$ maps x to F. 0 is a degree, and $[\![\text{straight}]\!]^{w,g}(0)$ maps x to F iff x has a degree of curvature greater than 0 in w; hence, $[\![\text{completely straight}]\!]^{w,g}$ also maps x to F iff x has a degree of curvature greater than 0 in w. The net effect of 'completely', therefore, is that of removing homogeneity in total adjectives: 'x is completely straight' is true if x has 0 degree of curvature and false otherwise.

Let's now take look at (35)b. For any x, $[\![\text{slightly bent}]\!]^{w,g}(x)$ = T iff, for some d, $[\![\text{bent}]\!]^{w,g}(d)$ maps x to T. The upshot of this is that, if x has a degree of curvature greater than 0 in w, $[\![\text{slightly bent}]\!]^{w,g}$ will map x to T (because, if x has a degree of curvature greater than 0 in w, $[\![\text{bent}]\!]^{w,g}(0)$ will map x to T). For any x, $[\![\text{slightly bent}]\!]^{w,g}(x)$ = F iff, for some d, $[\![\text{bent}]\!]^{w,g}(d)$ maps x to F. $[\![\text{bent}]\!]^{w,g}(d)$, as can be seen in (32)b, maps x to F iff x has 0 degree of curvature in w (irrespective of the value that d ends up taking); thus, on the proposed semantics, $[\![\text{slightly bent}]\!]^{w,g,-}$ and $[\![\text{bent}]\!]^{w,g,-}(d)$ are the same function. The net effect of *slightly*, therefore, is that of removing homogeneity in partial adjectives: 'x is slightly bent' is true if x has a degree of curvature greater than 0 and false otherwise.

Next, I will show how the entries in (33) and (34) predict the deviance of 'slightly straight' and 'completely bent'.

(36)

$$\text{a. } \#\ [\![\text{slightly}]\!]^{w,g}([\![\text{straight}]\!]^{w,g}) = \lambda x\,.\begin{cases} \text{T iff } \exists d\,([\![\text{straight}]\!]^{w,g}(d))(x) = \text{T} \\ \text{F iff } \forall d\,([\![\text{straight}]\!]^{w,g}(d))(x) = \text{F (i.e. never)} \\ \#\ \text{otherwise} \end{cases}$$

$$\text{b. } \#\ [\![\text{completely}]\!]^{w,g}([\![\text{bent}]\!]^{w,g}) = \lambda x\,.\begin{cases} \forall d\,([\![\text{bent}]\!]^{w,g}(d))(x) = \text{T (i.e. never)} \\ \exists d\,([\![\text{bent}]\!]^{w,g}(d))(x) = \text{F} \\ \#\ \text{otherwise} \end{cases}$$

Let's first consider (36)a. For any x, $[\![\text{slightly straight}]\!]^{w,g}(x)$ = F iff, for every d, $[\![\text{straight}]\!]^{w,g}(d)$ maps x to F. $[\![\text{straight}]\!]^{w,g}(d)$ maps x to F iff x's degree of curvature in w is greater than d; hence, $[\![\text{slightly straight}]\!]^{w,g}$ will map x to F iff, for every d, x's degree of curvature in w is greater than d. However, there is no x such that its degree of curvature is greater than every possible degree of curvature: indeed, there is always going to be a value that d can take that is greater than the degree of curvature of x. The upshot of this is that 'x is slightly straight' can never be false. Let's now consider (36)b. For any x, $[\![\text{completely bent}]\!]^{w,g}(x)$ = T iff, for every d, $[\![\text{bent}]\!]^{w,g}(d)$ maps x to T. $[\![\text{bent}]\!]^{w,g}(d)$ maps x to T iff x's degree of curvature in w is greater than d; hence, $[\![\text{completely bent}]\!]^{w,g}$ will map x to T iff, for every d, x's degree of curvature in w is greater than d. However, as just noted, there is no x such that its degree of curvature is greater than every possible degree of curvature. The upshot of this is that 'x is completely bent' can never be true. The semantic oddness of 'slightly bent' and 'completely bent' is therefore accounted for: $[\![\text{slightly}]\!]^{w,g}$ makes the negative extension of 'straight' collapse, whereas $[\![\text{completely}]\!]^{w,g}$ makes the positive extension of 'bent' collapse.

It remains to be explained why 'slightly' and 'completely' are deviant if combined with relative adjectives.

(37)

a. $\# \, [\![\text{slightly}]\!]^{w,g} ([\![\text{tall}]\!]^{w,g}) = \lambda x \,.\, \begin{cases} \text{T iff } \exists d \, ([\![\text{tall}]\!]^{w,g} (d)) \, (x) = \text{T (i.e. always)} \\ \text{F iff } \forall d \, ([\![\text{tall}]\!]^{w,g} (d)) \, (x) = \text{F} \\ \# \text{ otherwise} \end{cases}$

b. $\# \, [\![\text{completely}]\!]^{w,g} ([\![\text{tall}]\!]^{w,g}) = \lambda x \,.\, \begin{cases} \text{T iff } \forall d \, ([\![\text{tall}]\!]^{w,g} (d)) \, (x) = \text{T (i.e. never)} \\ \text{F iff } \exists d \, ([\![\text{tall}]\!]^{w,g} (d)) \, (x) = \text{F} \\ \# \text{ otherwise} \end{cases}$

Let's have a look at 'slightly tall' first; according (37)a, for any x, $[\![\text{slightly tall}]\!]^{w,g}(x)$ = T iff, for some d, $[\![\text{tall}]\!]^{w,g}(d)$ maps x to T. $[\![\text{tall}]\!]^{w,g}(d)$ maps x to T iff x's height is greater than d; hence, $[\![\text{slightly tall}]\!]^{w,g}$ maps x to T iff, for some d, x's height is greater than d. The upshot of this is that $[\![\text{slightly}]\!]^{w,g}$ weakens the meaning of 'tall' to triviality: indeed, no matter how tall or how short x is, there will always be a degree d such that $[\![\text{tall}]\!]^{w,g}(d)$ maps x to T. Hence, on the proposed semantics, 'x is slightly tall' is bound to always be true. Let's now consider the case of 'completely tall'; according to (37)b, for every x, $[\![\text{completely tall}]\!]^{w,g} (x)$ = T iff, for every d, $[\![\text{tall}]\!]^{w,g}(d)$ maps x to T. $[\![\text{tall}]\!]^{w,g}(d)$ maps x to T iff x's height is greater than d; hence, $[\![\text{completely tall}]\!]^{w,g}$ maps x to T iff, for every d, x's height is greater than d. The upshot of this is that 'x is completely tall' will never be true: indeed, there is always a value that d can take that is greater than the height of x.

Before closing up this section, I want to make two observations. First, the meaning of *slightly* given in (34) predicts that (38)a should be truth-conditionally equivalent to (38)b. *Prima facie*, this seems like a bad prediction: if the rod in question is extremely bent, then (38)b is true whereas (38)a seems false.

(38) a. The rod is slightly bent.
b. The rod has some degree of curvature greater than 0.

To explain the contrast in meaning between (38)a and (38)b, pragmatic considerations need to be invoked; and, it should be stressed, this is not a last resort: there is in fact independent evidence that suggests that the inference that (38)a triggers to the effect that the rod is not very bent comes about via a scalar implicature. I cannot review all the relevant data here; but consider, for example, the contrast between (39)a and (39)b.

(39) Context: There are 10 rods. You and your friend are not allowed to keep those which are very bent. You pick one and tell your friend…[105]

a. We can keep this one because it's only slightly bent.
b. ? We can keep this one because it's slightly bent.

If 'slightly' meant *only slightly* (≅ *slightly but not very much*), then the contrast between (39)a and (39)b would be a complete mystery. Conversely, if the upper-bounded aspect part of the meaning of *slightly* comes about via pragmatic reasoning, then the said contrast is expected: it is just a fact that scalar implicatures are hard (if not impossible) to compute in the scope of *because*; compare, for example, (40)a with (40)b.

(40) a. I am sad because only some of my friends came to my wedding.
b. ? I am sad because some of my friends came to my wedding.

(40)b, quite clearly, does not receive the same interpretation as (40)a, which shows that the negative implication associated with *some* is not triggered in the scope of *because*. I take the parallelism between (39) and (40) to indicate that the upper-bounded aspect of the meaning of *slightly* (just like the upper-bounded aspect of the meaning of *some*) must come via pragmatic reasoning (see Spector 2014 for a related discussion on *almost*). Thus, that the meaning that I have proposed for *slightly* is not upper-bounded is not only not a problem: it is in fact a good feature of the proposal.

The second issue concerns the distribution of *slightly* and *completely*. Though the proposed analysis makes sense of a great deal of data, it cannot explain the strange behaviour of *open* (and, presumably, of some other adjectives that I have not yet identified). The adjective *open* is peculiar because it appears to be good with both *slightly* and *completely* (as opposed to *closed*, which is only good with *completely*). This poses a problem: if *open* is treated as a total adjective, the proposed analysis predicts that it should be good with *completely* and bad with *slightly*; if, alternatively, it *open* is treated as a partial adjective, then the reverse should be the case. I do not yet have a solution to this puzzle. One could argue that *open* is ambiguous between a total and a partial adjective; such a solution, however, seems rather ad-hoc. It could also be argued that *open* is neither a total nor a partial adjective but something like hybrid of the two; I am attracted to this idea: whether it is feasible, I do not know.

[105] Thanks to Benjamin Spector for the example.

4.1.3 Imprecision, once again

Analysing *completely* and *slightly* as homogeneity removers is not a free ticket: it makes specific predictions regarding the pragmatic behaviour of certain sentences. Specifically, such an account predicts that (41)a should support imprecise interpretation, while (41)b should not; likewise, it predicts that (42)a should be able to be interpreted imprecisely, while (42)b should not.

(41) a. The bookcase is empty.
b. The bookcase is completely empty.

(42) a. The rod is bent.
b. The rod is slightly bent.

The former prediction we know to be correct: when combined with an absolute adjective, *completely* prevents the sentence that contains the said adjective from being interpreted imprecisely (cf. Chapter III). The latter prediction, by contrast, has not yet been tested. Does *slightly* in fact block imprecision? Insofar as it removes homogeneity, as explained in Chapter IV, it should.

Testing whether *slightly* blocks imprecise interpretation is not straightforward, however. First, *slightly* appears to be a positive polarity item (an item that cannot scope below negation; see, for example, Ladusaw 1979, Krifka 1991, van der Wouden 1997, Szabolcsi 2004, and Spector 2014, among others); this, unfortunately, limits the sort of linguistic environments in which the prediction can be tested. In addition, a sentence such as (42)b, as discussed, triggers the implicature that the rod was not very bent, an inference which is absent in (42)a, the *slightly*-free sentence. Due to this, (42)a and (42)b are not adequate minimal pairs, at least not for our purposes.

I can think of two examples, however, that circumvent these issues and provide support for the case that *slightly* has an imprecision-blocking effect. Consider, first, (43) below.

(43) <u>Context:</u> Frank bought an old fishing rod online; like all old finishing rods, it wasn't perfectly straight, but, for an old fishing rod, it was in great shape and could almost pass as new. However, Frank eventually decided to return the fishing rod. When his son demanded an explanation, Frank replied…

a. I returned the fishing rod because it was bent.
b. I returned the fishing rod because it was slightly bent.

In the stipulated context, (43)a feels false; technically, what seems to be going on here is this: the granularity parameter associated with *bent* is set at value (much) higher than 0 and, as a result of this, (43)a ends up presupposing that the fishing rod was *significantly* bent; this inference contradicts common ground knowledge: the old fishing rod was, as a matter of fact, not significantly bent. (43)b, conversely, does not contradict common ground knowledge: it is just an insane statement, because the fact that the rod was slightly bent is clearly not a good reason to send it back (after all, no old fishing rod is perfectly straight). The contrast between (43)a and (43)b, therefore, suggests that the presence of *slightly*, as predicted, prevents (43)b from being interpreted imprecisely.

(44) Context: Frank's daughter comes back from work and is informed of the fact that his father decided to send back the fishing rod. Rather confused, she utters…

a. Really? But the rod was is such good shape… *It wasn't even bent!*
b. Really? But the rod was is such good shape… *It wasn't even slightly bent!*

(44)a, in the context stipulated in (43)+(44), may not be true (strictly speaking) but it is undoubtedly true enough. (44)b, by contrast, comes across as a false statement: '[the rod] wasn't even slightly bent' appears to entail, as predicted in the semantics that I have proposed in § 4.1.2, that the rod was perfectly straight, which contradicts contextual knowledge. The contrast between (44)a and (44)b, therefore, provides evidence which is at least compatible with the account presented here: *slightly* removes homogeneity and, by so doing, disallows imprecise interpretation.

4.1.3.1 The comparative form

If *slightly* is a homogeneity remover, then a more general prediction can be stated: sentences that have as constituents constructions that can be modified by *slightly* should allow for imprecise interpretation when the said constructions are not modified by *slightly*. In addition to partial adjectives, what other constructions does *slightly* modify? *Slightly* modifies comparative predicates, as shown below:

(45) a. Mary is *slightly* taller than Paula.
b. Mary is *slightly* less than Paula.
c. This rod is *slightly* more bent than that rod.
d. This rod is *slightly* less bent than that rod.
e. This glass *slightly* emptier than that glass.
f. This glass *slightly* emptier than that glass.

The question is, can then the sentences in (46) be interpreted imprecisely?

(46) a. Mary is (not) taller than Paula.
b. Mary is (not) less tall than Paula.
c. This rod is (not) more bent than that rod.
d. This rod is (not) less bent than that rod.
e. This glass is (not) emptier than that glass.
f. This glass is (not) emptier than that glass.

The answer is: yes, they can. Consider, for example, (47) below.

(47) <u>Context:</u> Mary broke up with John; she didn't like him: he wasn't tall enough for her. Paula wants to find a new boyfriend for Mary; she asks her: 'Would you be interested in having a date with Mike?' Mary replies…

a. Mike? He isn't taller than John.
b. Mike? He isn't even slightly taller than John.

(47)a supports imprecise interpretation: indeed, (47)a can be used by Mary if Mike is as a matter of fact a bit taller than John but not significantly so. (47)b, by contrast, cannot: (47)b is false if Mike happens to be a bit taller than John.

Stating a compositional semantics of the comparison form goes beyond the scope of this thesis.

5 SUMMARY

In this chapter, I have used PNM to derive imprecision in sentences containing absolute adjectives, sentences such as 'the piggy bank is(n't) *empty*'. PNM's underlying assumption, assumption that inherits from Križ (2015), is that (im)precision (or (non-)maximality) has homogeneity as a pre-requisite: hence, in order for PNM to do the required job, 'the piggy bank is(n't)empty' must be posited to be homogenous. There is a problem with doing so, however: 'the piggy bank is(n't)empty', unlike unquantified plurals sentences, is not visibly homogenous; so, unless one can make sense of the idea of invisible homogeneity, this approach does not look very promising. In the first part of this chapter, therefore, I have been engaged with the task of building a homogenous semantics for absolute adjectives which, in conjunction with (non-)maximality and an alignment condition (see § 6), predicts gap invisibility.

In the second part of this chapter, I have criticised the account that seeks to explain the distribution facts of *completely* and *slightly* by appealing to considerations of scale structure. This account, I have argued, does not succeed in explaining the relevant facts and, furthermore, makes incorrect predictions about the semantic behaviour of relative adjectives. I have proposed a different account of the distribution facts of *completely* and *slightly*, which is tied up to the semantics for absolute adjectives developed in § 3.1: on this account, *completely* and *slightly* are both homogeneity removers: *completely* removes homogeneity in total adjectives (but messes up with the meaning of partial and relative adjectives) whereas *slightly* removes homogeneity in partial adjectives (but messes up with the meaning of total and relative adjectives).

6 APPENDIX I

I will provide an alignment condition within the classical framework (cf. Chapter V, § 2.1): partition semantics without probabilities. Ideally, the alignment condition should be stated within the more sophisticated probabilistic framework developed in Chapter V, § 2.2; to my regret, I have not yet managed to do so.

6.1 Definitions

To begin with, let's define the following (very simple) notion of equivalence between propositions:

(III) Equivalence between propositions relative to an issue

I is a partition on C, the Context Set: $I = \{i_i,\dots,i_n\}$
i is partition cell: $i \in I$
p and q are propositions: $p, q \in \wp(\Omega)$

Two propositions p and q are equivalent relative to an issue I (written $p \approx_I q$) if they intersect the same i-cells.

The alignment condition is as follows:

(IV) Alignment condition

If there is a maximal ρ-value such that $S^+ \approx_I S^\#$, then ρ should be assigned that value. If there is a maximal ρ-value such that $S^- \approx_I S^\#$, then ρ should be assigned that value. If neither of these values exists, then g(ρ) = 0.

Note that (IV) presupposes that there should be a unique maximal ρ-value that meets either one of the two sub-conditions.

The reason why (IV) asks for the *maximal* ρ-value is because we want there to be as much homogeneity (and hence non-maximality) as the context allows for.

6.2 Test case: the piggy bank is(n't) empty

6.2.1 The issue

In the classical framework, the issue is a partition of the Context Set, the set of worlds compatible with the common ground. Let's assume that the Context Set entails that: (a) to buy a cake one needs more than £5; (b) the speaker keeps her money in the piggy bank (i.e. she doesn't have money elsewhere).

The issue that I will assume to be in operation can be paraphrased as *do you have enough money to buy a cake?* This issue, let's call it I, has two cells, i.e.: i_1 (NO cell) and i_2 (YES cell).

6.2.2 The positive sentence

(48) A: Do you have enough money to buy a cake?
B: The piggy bank is empty.

The meaning of (48)B is as follows:

(49) $[\![\text{the piggy bank is empty}_\rho]\!]^{g,+} =^* \{w : \text{the piggy bank has £0 in } w\}$
$[\![\text{the piggy bank is empty}_\rho]\!]^{g,-} =^* \{w : \text{the piggy bank has more than £ g}(\rho) \text{ in } w\}$
$[\![\text{the piggy bank is empty}_\rho]\!]^{g,\#} =^* \{w : \text{the piggy bank has £ } (0, \text{g}(\rho)] \text{ in } w\}$

*Identifying sets with their characterising functions.

6.2.2.1 How to find the value of ρ given I

Figure 2. A (partial) representation of issue *I*.

i_1: Not enough money to buy cake (≤ £5 in piggy bank)	w_0: There's no money in the piggy bank. w_1: There's £1 in the piggy bank. w_2: There's £2 in the piggy bank. w_3: There's £3 in the piggy bank. w_4: There's £4 in the piggy bank. w_5: There's £5 in the piggy bank.
i_2: Enough money to buy cake (> £5 in piggy bank)	w_6: There's £5.05 in the piggy bank. w_7: There's £6 in the piggy bank. etc.

There is no ρ-value (neither maximal nor non-maximal) such that $S^- \approx_I S^\#$: if g(ρ) < 5, then S^- intersects both i_1 and i_2, while $S^\#$ only intersects i_1; if g(ρ) = 5, then S^- intersects only i_2, while $S^\#$ intersects only i_1; if g(ρ) > 5, then S^- intersects only i_2, while $S^\#$ intersects i_1 and i_2. However, there is a maximal ρ-value such that $S^+ \approx_I S^\#$, namely, 5: hence, as per (IV), g(ρ) = 5. The value of ρ cannot be smaller than 5 (for example, 4): the condition instructs us to pick the *maximal* ρ-value such that $S^+ \approx_I S^\#$; neither can it be greater than 5 (for example, 6): as soon as g(ρ) > 5, $S^+ \not\approx_I S^\#$: indeed, when g(ρ) > 5, S^+ intersects i_1, and $S^\#$ intersects both i_1 and i_2.

In a non-probabilistic framework, to generate the final interpretation of the sentence (the sentence's communicated content), one just takes the union of all the cells that the positive extension intersects (cf. Križ's algorithm; Chapter IV, § 3): (48)B's communicated content, therefore, will be i_1. The worlds in the extension gap—which, if g(ρ) = 5, is $\{w : \text{the piggy bank has £ } (0, 5] \text{ in } w\}$—are all bound to be in i_1 (because 5 is the maximal ρ-value at which $S^+ \approx_I S^\#$ holds and S^+ only intersects i_1). The upshot of this is that the extension gap of (48)B will not be visible: indeed, if evaluated relative to the worlds in the extension gap, the sentence, though semantically undefined in those worlds, will be judged to be true enough (i.e. i_1 is true in those worlds).

6.2.3 The negative sentence

(50) A: Do you have enough money to buy a cake?
B: The piggy bank isn't empty.

The meaning of (50)B is as follows.

(51) $[\![\text{the piggy bank is}\mathbf{n't}\text{ empty}_\rho]\!]^{g,+} =_* \{w : \text{the piggy bank has more than £ g}(\rho) \text{ in } w\}$
$[\![\text{the piggy bank is}\mathbf{n't}\text{ empty}_\rho]\!]^{g,-} =_* \{w : \text{the piggy bank has £0 in } w\}$
$[\![\text{the piggy bank is}\mathbf{n't}\text{ empty}_\rho]\!]^{g,\#} =_* \{w : \text{the piggy bank has £ } (0, g(\rho)] \text{ in } w\}$

*Identifying sets with their characterising functions.

6.2.3.1 How to find the value of ρ *given I*

There is no ρ-value (neither maximal nor non-maximal) such that $S^+ \approx_I S^\#$: if $g(\rho) < 5$, then S^+ intersects both i_1 and i_2, while $S^\#$ only intersects i_1; if $g(\rho) = 5$, then S^+ intersects only i_2, while $S^\#$ intersects only i_1; if $g(\rho) > 5$, then S^+ intersects only i_2, while $S^\#$ intersects both i_1 and i_2. However, there is a maximal ρ-value such that $S^- \approx_I S^\#$, namely, 5: hence, as per (IV), $g(\rho) = 5$.

As before, to derive the final interpretation of the sentence (the sentence's communicated content), the union of all the cells that the sentence's positive extension intersects is taken: (50)B's communicated content, therefore, is i_2. The worlds in the extension gap—which, if $g(\rho) = 5$, is $\{w$: the piggy bank has £ (0, 5] in $w\}$—are all elements of i_1, as noted in § 6.2.2.1: hence, relative to these worlds, (50)B, though semantically undefined in these worlds, is bound to be perceived as false: this is because i_2, the communicated content of the sentence, is false in those worlds. Thus, here again, (non-)maximality has the effect of making the gap invisible.

CLOSING REMARKS

In Chapters II and III, I was concerned with the question of whether content words can be subjected to lexical modulation, a mechanism which is hypothesised to come in at least two varieties, lexical narrowing and lexical broadening. I have found no convincing evidence that any such mechanism is operative in natural language: according to the data examined in this thesis, lexical modulation is nothing else but an illusion.

In Chapter III, the conclusion was reached that the phenomenon of (im)precision cannot be adequately described in terms of lexical broadening; this conclusion, importantly, does not entail that absolute adjectives, and other expressions whose presence is correlated with the emergence of imprecise interpretation, are not context-sensitive. In Chapter VI, absolute adjectives have been formalised as $\langle d,\langle e,t\rangle\rangle$ functions, the degree argument being provided via ρ, a contextual variable. This formalisation makes absolute adjectives context-sensitive: their meaning is tied up to the value of ρ, which is contextually determined. It should be stressed, however, that what ρ does is to induce homogeneity (*not* imprecision) at the lexical level. Imprecision as such, in both Križ (2015) and PNM, is a non-semantic mechanism that operates on whole sentences and has homogeneity as a pre-requisite.

One of the issues that have been raised in this thesis is the issue of gap (in)visibility. According to the account put forward in Chapter VI, the gap in sentences such as ‘X is empty’ is pragmatically induced via ρ: because the same pragmatic considerations that induce the gap also cover it up, the gap cannot be seen. The gap in plural sentences, by contrast, isn’t context-sensitive: ‘the professors smiled’ is true if all the professors smiled, false if none of them did, and undefined otherwise; as a result, this gap is visible. Now, there are reasons to suspect that the gap in plural sentences may, in fact, be context-sensitive too (at least, stipulating that this is the case improves PNM’s predictive power and accuracy; cf. Chapter V, § 3.5). It is an open question, which I will not tackle here, whether the gap in plural sentences can be made context-sensitive (reliant on one—or more—granularity parameters) without (incorrectly) predicting its invisibility.

REFERENCES

Alxatib, Sam, Peter Pagin, and Uli Sauerland. 2013. "Acceptable Contradictions: Pragmatics or Semantics? A Reply to Cobreros et Al." *Journal of Philosophical Logic*, 619–634.

Alxatib, Sam, and Francis Jeffry Pelletier. 2011. "The Psychology of Vagueness: Borderline Cases and Contradictions." *Mind & Language* 26 (3): 287–326. https://doi.org/10.1111/j.1468-0017.2011.01419.x.

Arregui, Ana. 2008. "Some Remarks on Domain Widening." In *Proceedings of the 27th West Coast Conference on Formal Linguistics*, 45–53.

Bale, Alan Clinton. 2008. "A Universal Scale of Comparison." *Linguistics and Philosophy* 31 (1): 1. https://doi.org/10.1007/s10988-008-9028-z.

Bar-Lev, Moshe. 2018. "Free Choice, Homogeneity, and Innocent Inclusion." PhD Thesis, Doctoral Dissertation, Hebrew University of Jerusalem.

Bartsch, Renate, and Theo Vennemann. 1972. "Relative Adjectives and Comparison." *UCLA Papers in Syntax* 2: 80–107.

Beaver, David, and Emiel Krahmer. 2001. "A Partial Account of Presupposition Projection." *Journal of Logic, Language and Information* 10 (2): 147. https://doi.org/10.1023/A:1008371413822.

Bergen, Leon. 2016. "Joint Inference in Pragmatic Reasoning." PhD Thesis, Massachusetts Institute of Technology.

Bergen, Leon, Roger Levy, and Noah Goodman. 2016. "Pragmatic Reasoning through Semantic Inference." *Semantics & Pragmatics* 9: 20.

Bierwisch, Manfred. 1989. "The Semantics of Gradation." In *Dimensional Adjectives: Grammatical Structure and Conceptual Interpretation*, edited by M. Bierwisch and E. Lang, 71–261. Springer-Verlag.

Breheny, Richard. 2003. "A Lexical Account of Implicit (Bound) Contextual Dependence." In *Semantics and Linguistic Theory*, 13:55–72.

———. 2005. "Exhaustivity, Homogeneity, and Definiteness." In *Fifteenth Amsterdam Colloquium*, 59–65.

Brisson, Christine M. 1998. "Distributivity, Maximality and Floating Quantifiers." PhD Thesis, Rutgers, The State University of New Jersey.

Büring, Daniel. 2003. "On D-Trees, Beans, And B-Accents." *Linguistics and Philosophy* 26 (5): 511–45. https://doi.org/10.1023/A:1025887707652.

Burnett, Heather. 2014. "A Delineation Solution to the Puzzles of Absolute Adjectives." *Linguistics & Philosophy* 37 (1): 1–39.

Carlson, Greg N. 1977. "Reference to Kinds in English." PhD Thesis, University of Massachusetts, Amherst.

———. 1982. “Generic Terms and Generic Sentences.” *Journal of Philosophical Logic* 11 (2): 145–181.

Carnap, Rudolf. 1950. *Logical Foundations of Probability*. Chicago: University of Chicago Press.

Carston, Robyn. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. John Wiley & Sons.

———. 2012. “Word Meaning and Concept Expressed.” *The Linguistic Review* 29 (4): 607–623. https://doi.org/10.1515/tlr-2012-0022.

Carston, Robyn, and George Powell. 2008. “Relevance Theory—New Directions and Developments.” In *The Oxford Handbook of Philosophy of Language*, Ernest Lepore and Barry C. Smith (eds.). Oxford: Oxford University Press.

Carter, Sam. 2016. “The Dynamics of Loose Talk.” In *NASSLLI 2016*.

———. 2019. “The Dynamics of Loose Talk.” *Noûs*. https://doi.org/10.1111/nous.12306.

Chierchia, Gennaro. 1995. “Individual Level Predicates as Inherent Generics.” In *The Generic Book*, edited by Greg N. Carlson and Francis J. Pelletier, 176–223. Chicago University Press.

Chierchia, Gennaro, Danny Fox, and Benjamin Spector. 2008. “The Grammatical View of Scalar Implicatures and the Relationship between Semantics and Pragmatics.” In *Handbook of Semantics*, edited by Paul Portner, Claudia Maienborn, and Klaus von Heusinger. New York: Mouton de Gruyter.

Cobreros, Pablo, Paul Egré, David Ripley, and Robert van Rooij. 2012. “Tolerant, Classical, Strict.” *Journal of Philosophical Logic* 41 (2): 347–85. https://doi.org/10.1007/s10992-010-9165-z.

Cohen, L. Jonathan. 1971. “Some Remarks on Grice’s Views about the Logical Particles of Natural Language.” In *Pragmatics of Natural Languages*, edited by Yehoshua Bar-Hillel, 50–68. Dordrecht: Reidel.

———. 1986. “How Is Conceptual Innovation Possible?” *Erkenntnis* 25 (2): 221–238.

Colston, Herbert L., and Shauna B. Keller. 1998. “You’ll Never Believe This: Irony and Hyperbole in Expressing Surprise.” *Journal of Psycholinguistic Research* 27 (4): 499–513. https://doi.org/10.1023/A:1023229304509.

Condoravdi, Cleo A. 1994. “Descriptions in Context.” PhD Thesis, Yale University.

Cooper, Robin. 1996. “The Role of Situations in Generalized Quantifiers.” *The Handbook of Contemporary Semantic Theory* 65: 86.

Cresswell, Max J. 1976. “The Semantics of Degree.” In *Montague Grammar*, edited by Barbara Partee, 261–292. New York: Academic Press.

Del Pinal, Guillermo. 2018. “Meaning, Modulation, and Context: A Multidimensional Semantics for Truth-Conditional Pragmatics.” *Linguistics and Philosophy* 41 (2): 165–207.

Dummett, Michael. 1978. *Truth and Other Enigmas*. London: Duckworth.

Egré, Paul. 2019. “Respects for Contradictions.” In *Graham Priest on Dialetheism and Paraconsistency*, edited by Can Başkent and Thomas Macaulay Ferguson, 39–57.

Outstanding Contributions to Logic. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-25365-3_4.

Egré, Paul, Vincent de Gardelle, and David Ripley. 2013. "Vagueness and Order Effects in Color Categorization." *Journal of Logic, Language and Information* 22 (4): 391–420. https://doi.org/10.1007/s10849-013-9183-7.

Egré, Paul, and Jérémy Zehr. 2018. "Are Gaps Preferred to Gluts? A Closer Look at Borderline Contradictions." In *The Semantics of Gradability, Vagueness, and Scale Structure: Experimental Perspectives*, edited by Elena Castroviejo, Louise McNally, and Galit Weidman Sassoon, 25–58. Language, Cognition, and Mind. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-77791-7_2.

Etxeberria, Urtzi. 2005. "Contextual Restriction and Quantification in Basque." *Belgian Journal of Linguistics* 19 (1): 211–245.

Farkas, Donka F., and Katalin E. Kiss. 2000. "On the Comparative and Absolute Readings of Superlatives." *Natural Language & Linguistic Theory* 18 (3): 417–455.

Fine, Kit. 1975. "Vagueness, Truth and Logic." *Synthese* 30 (3/4): 265–300.

Fintel, Kai von. 1994. "Restrictions on Quantifier Domains." PhD Thesis, University of Massachusetts at Amherst.

———. 2004. "Would You Believe It? The King of France Is Back! Presuppositions and Truth-Value Intuitions." In *Descriptions and Beyond*, edited by Marga Reimer and Anne Bezuidenhout, 315–41. Oxford: OUP.

———. 2014. "Quantifier Domain Restriction." Presented at the Semantics Workshop, Rutgers University.

Fodor, Janet Dean. 1970. "The Linguistic Description of Opaque Contents." PhD Thesis, Massachusetts Institute of Technology.

Fox, Danny. 2018. "Partition by Exhaustification: Comments on Dayal 1996." *Proceedings of Sinn Und Bedeutung* 22 (1): 403–34.

Fox, Danny, and Martin Hackl. 2006. "The Universal Density of Measurement." *Linguistics and Philosophy* 29 (5): 537–86. https://doi.org/10.1007/s10988-006-9004-4.

Gajewski, Jon Robert. 2005. "Neg-Raising : Polarity and Presupposition." Thesis, Massachusetts Institute of Technology. https://dspace.mit.edu/handle/1721.1/33696.

Giannakidou, Anastasia. 2004. "Domain Restriction and the Arguments of Quantificational Determiners." *Semantics and Linguistic Theory* 14: 110–26.

Gillon, Carrie. 2006. "The Semantics of Determiners: Domain Restriction in Skwxwú7mesh." PhD Thesis, University of British Columbia.

Grice, Herbert Paul. 1975. "Logic and Conversation." In *Syntax and Semantics: Speech Acts*, Peter Cole and Jerry L. Morgan, 3:41–58. New York: Academic Press.

Groenendijk, Jeroen Antonius Gerardus, and Martin Johan Bastiaan Stokhof. 1984. “Studies on the Semantics of Questions and the Pragmatics of Answers.” PhD Thesis, Univ. Amsterdam.

———. 1990. “Partitioning Logical Space.” Annotated handout.

Heim, Irene. 1983. “On the Projection Problem for Presuppositions.” In *Formal Semantics - The Essential Readings*, edited by P. Portner and B. H. Partee, 249–260. Blackwell.

———. 1994. “Interrogative Semantics and Karttunen’s Semantics for Know.” In *Proceedings of IATL*, 1:128–144.

———. 1999. “Notes on Superlatives.” Ms, MIT.

Heim, Irene, and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.

Hoek, Daniel. 2018. “Conversational Exculpature.” *The Philosophical Review* 127 (2): 151–96.

Horn, Laurence. 1972. “The Semantics of Logical Operators in English.” PhD Thesis, UCLA.

Jakobson, Roman. 1960. “Closing Statement: Linguistics and Poetics.” In *Style in Language*, edited by Thomas Albert Sebeok, 350–77. M.I.T.

Kadmon, Nirit, and Fred Landman. 1993. “Any.” *Linguistics and Philosophy* 16 (4): 353–422.

Kao, Justine T., Jean Y. Wu, Leon Bergen, and Noah D. Goodman. 2014. “Nonliteral Understanding of Number Words.” *Proceedings of the National Academy of Sciences* 111 (33): 12002–7. https://doi.org/10.1073/pnas.1407479111.

Keefe, Rosanna. 2000. *Theories of Vagueness*. Cambridge University Press.

———. 2007. “Vagueness Without Context Change.” *Mind* 116 (462): 275–92. https://doi.org/10.1093/mind/fzm275.

Kennedy, Christopher. 2007. “Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives.” *Linguistics and Philosophy* 30 (1): 1–45.

———. 2011. “Vagueness and Comparison.” In *Vagueness and Language Use*, edited by Paul Egré and Nathan Klinedinst, 73–97. Palgrave Studies in Pragmatics, Language and Cognition. London: Palgrave Macmillan UK. https://doi.org/10.1057/9780230299313_4.

Kennedy, Christopher, and Louise McNally. 2005. “Scale Structure, Degree Modification, and the Semantics of Gradable Predicates.” *Language* 81 (2): 345–381.

———. 2010. “Color, Context, and Compositionality.” *Synthese* 174 (1): 79–98. https://doi.org/10.1007/s11229-009-9685-7.

Klein, Ewan. 1980. “A Semantics for Positive and Comparative Adjectives.” *Linguistics and Philosophy* 4 (1): 1–45.

Klinedinst, Nathan, and Daniel Rothschild. 2011. “Exhaustivity in Questions with Non-Factives.” *Semantics and Pragmatics* 4 (2): 1–23.

Kratzer, Angelika. 1981. “The Notional Category of Modality.” In *Words, Worlds, and Contexts: New Approaches in Word Semantics*, edited by Hans-Jürgen Eikmeyer and Hannes Rieser. Walter de Gruyter.

———. 2004. “Covert Quantifier Restrictions in Natural Languages.” Palazzo Feltrenelli, Gargnano, Italy.

———. 2007. “Situations in Natural Language Semantics.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2019. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/sum2019/entries/situations-semantics/.

Krifka, Manfred. 1991. “Some Remarks on Polarity Items.” In *Semantic Universals and Universal Semantics*, edited by Dietmar Zaefferer, 150–189. Berlin: Foris.

Križ, Manuel. 2015. “Aspects of Homogeneity in the Semantics of Natural Language.” University of Vienna.

———. 2017. “Bare Plurals, Multiplicity, and Homogeneity.”

———. 2019. “Homogeneity Effects in Natural Language Semantics.” *Language and Linguistics Compass* 13 (11). https://onlinelibrary.wiley.com/doi/full/10.1111/lnc3.12350.

Kullback, S., and R. A. Leibler. 1951. “On Information and Sufficiency.” *The Annals of Mathematical Statistics* 22 (1): 79–86.

Kyburg, Alice, and Michael Morreau. 2000. “Fitting Words: Vague Language in Context.” *Linguistics and Philosophy* 23 (6): 577–97.

Ladusaw, William A. 1979. “Polarity Sensitivity as Inherent Scope Relations.” University of Texas at Austin.

Lasersohn, Peter. 1999. “Pragmatic Halos.” *Language* 75 (3): 522–51. https://doi.org/10.2307/417059.

Lassiter, Daniel, and Noah D. Goodman. 2013. “Context, Scale Structure, and Statistics in the Interpretation of Positive-Form Adjectives.” In *Semantics and Linguistic Theory*, 23:587–610.

Lewis, David. 1975. “Adverbs of Quantification.” In *Formal Semantics of Natural Language*, edited by Edward L. Keenan, 178–188. Cambridge University Press.

———. 1979. “Scorekeeping in a Language Game.” *Journal of Philosophical Logic* 8 (1): 339–359.

———. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.

Löbner, Sebastian. 2000. “Polarity in Natural Language: Predication, Quantification and Negation in Particular and Characterizing Sentences.” *Linguistics and Philosophy* 23 (3): 213–308. https://doi.org/10.1023/A:1005571202592.

Ludlow, Peter. 2014. *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. Oxford University Press.

Magri, Giorgio. 2014. “An Account for the Homogeneity Effect Triggered by Plural Definites and Conjunction Based on Double Strengthening.” In *Pragmatics, Semantics and the Case of Scalar Implicatures*, edited by Salvatore Pistoia Reda, 99–145. Palgrave Studies in Pragmatics, Language and Cognition. London: Palgrave Macmillan UK. https://doi.org/10.1057/9781137333285_5.

Malamud, Sophia A. 2012. “The Meaning of Plural Definites: A Decision-Theoretic Approach.” *Semantics and Pragmatics* 5 (0): 1–58. https://doi.org/10.3765/sp.5.3.

Martí, Luisa. 2003. “Contextual Variables.” PhD Thesis, University of Connecticut.

McCarthy, Michael, and Ronald Carter. 2004. “‘There’s Millions of Them’: Hyperbole in Everyday Conversation.” *Journal of Pragmatics* 36 (2): 149–84. https://doi.org/10.1016/S0378-2166(03)00116-4.

Milsark, Gary L. 1974. “Existential Sentences in English.” PhD Thesis, Massachusetts Institute of Technology.

Mitchell, Jonathan Edward. 1986. “The Formal Semantics of Point of View.” PhD Thesis, University of Massachusetts Amherst.

Moxey, Linda M., and Anthony J. Sanford. 1993. *Communicating Quantities: A Psychological Perspective.* Lawrence Erlbaum Associates, Inc.

Pagin, Peter, and Francis Jeffry Pelletier. 2007. “Context, Content, and Composition.”

Partee, Barbara. 1989. “Binding Implicit Variables in Quantified Contexts.” In *CLS*, 25:342–365.

Pelletier, Jeff. 2013. “Context, Compositionality, and Brevity.” In *Brevity*, edited by L. Goldstein, 178–97. Oxford: Oxford University Press.

Pinkal, Manfred. 1995. *Logic and Lexicon*. Dordrecht: Kluwer.

———. 1996. “Vagueness, Ambiguity, and Underspecification.” *Semantics and Linguistic Theory* 6: 185–201.

Qing, Ciyang, and Michael Franke. 2014. “Gradable Adjectives, Vagueness, and Optimal Language Use: A Speaker-Oriented Model.” *Semantics and Linguistic Theory* 24 (0): 23–41. https://doi.org/10.3765/salt.v24i0.2412.

Recanati, François. 1993. *Direct Reference: From Language to Thought*. Oxford: Blackwell.

———. 2002. *The Limits of Expressibility*. Edited by B. Smith. Cambridge: Cambridge University Press.

———. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.

———. 2009. “Compositionality, Semantic Flexibility, and Context-Dependence.” In *The Oxford Handbook of Compositionality*, W. Hinzen, E. Machery, and M. Werning. Oxford: Oxford University Press.

———. 2010. *Truth-Conditional Pragmatics*. Vol. 25. Oxford University Press Oxford.

Ripley, David. 2011. “Contradictions at the Borders.” In *Vagueness in Communication*, edited by Rick Nouwen, Robert van Rooij, Uli Sauerland, and Hans-Christian Schmitz, 169–88. Lecture Notes in Computer Science. Springer Berlin Heidelberg.

Roberts, Craige. (1996) 2012. “Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics.” *Semantics and Pragmatics* 5. https://doi.org/10.3765/sp.5.6.

Rothschild, Daniel. 2015. “Lexical Pragmatics and the Weakness of Belief.” UCL Pragmatics reading group.

Rotstein, Carmen, and Yoad Winter. 2004. “Total Adjectives vs. Partial Adjectives: Scale Structure and Higher-Order Modifiers.” *Natural Language Semantics* 12 (3): 259–88. https://doi.org/10.1023/B:NALS.0000034517.56898.9a.

Sainsbury, Mark. 1991. “Is There Higher-Order Vagueness?” *The Philosophical Quarterly (1950-)* 41 (163): 167–82. https://doi.org/10.2307/2219591.

Sauerland, Uli. 2014. “Intermediate Scalar Implicatures.” In *Pragmatics, Semantics and the Case of Scalar Implicatures*, edited by Salvatore Pistoia-Reda, 72–98. Springer.

Sauerland, Uli, Jan Anderssen, and Kazuko Yatsushiro. 2005. “The Plural Is Semantically Unmarked.” In *Linguistic Evidence: Empirical, Theoretical, and Computational Perspectives*, edited by Marga Reis and Stephan Kepser, 413–34. Berlin: Walter de Gruyter.

Scha, Remko. 1981. “Distributive, Collective and Cumulative Quantification.” In *Formal Methods in the Study of Language*, edited by J. Groenendijk, T. Janssen, and M. Stokhof, 2:483–512. Amsterdam: Mathematisch Centrum.

Schwarz, Florian. 2009. “Two Types of Definites in Natural Language.” PhD Thesis, University of Massachusetts at Amherst.

Searle, John R. 1980. “The Background of Meaning.” In *Speech Act Theory and Pragmatics*, edited by F. Kiefer, M. Bierwisch, and John R. Searle, 221–232. Dordrecht: Reidel.

———. 1992. *The Rediscovery of the Mind*. MIT press.

Serchuk, Phil, Ian Hargreaves, and Richard Zach. 2011. “Vagueness, Logic and Use: Four Experimental Studies on Vagueness.” *Mind & Language* 26 (5): 540–73. https://doi.org/10.1111/j.1468-0017.2011.01430.x.

Seuren, Pieter A. M. 1973. “The Comparative.” In *Generative Grammar in Europe*, edited by F. Kiefer and N. Ruwet, 528–564. Dordrecht: Reidel.

Sharvit, Yael, and Penka Stateva. 2002. “Superlative Expressions, Context, and Focus.” *Linguistics and Philosophy* 25 (4): 453–504.

Simons, Mandy. 2005. “Presuppositions and Relevance.” In *Semantics Versus Pragmatics*, edited by Zoltan Gendler Szabo, 329–255. Oxford University Press.

Simons, Mandy, Judith Tonhauser, David Beaver, and Craige Roberts. 2010. “What Projects and Why.” *Semantics and Linguistic Theory* 20 (0): 309–27. https://doi.org/10.3765/salt.v20i0.2584.

Soames, Scott. 1986. “Incomplete Definite Descriptions.” *Notre Dame Journal of Formal Logic* 27 (3): 349–375.

Solt, Stephanie. 2012. “Comparison to Arbitrary Standards.” In *Proceedings of Sinn Und Bedeutung*, 16:557–570.

Spector, Benjamin. 2007a. “Aspects of the Pragmatics of Plural Morphology: On Higher-Order Implicatures.” In *Presupposition and Implicature in Compositional Semantics*, edited by Uli

Sauerland and Penka Stateva, 243–81. Palgrave Studies in Pragmatics, Language and Cognition. London: Palgrave Macmillan UK. https://doi.org/10.1057/9780230210752_9.

———. 2007b. “Scalar Implicatures: Exhaustivity and Gricean Reasoning.” In *Questions in Dynamic Semantics*, edited by Maria Aloni, Butler Alastair, and Dekker Paul, 17:225–49. Current Research in the Semantics / Pragmatics Interface. Brill.

———. 2012. “Vagueness, (Local) Accommodation, Presupposition and Restrictors.” Course Notes on Trivalent Semantics for Vagueness and Presupposition, Vienna.

———. 2013. “Homogeneity and Plurals: From the Strongest Meaning Hypothesis to Supervaluations.” In *Sinn Und Bedeutung*. Vol. 18.

———. 2014. “Global Positive Polarity Items and Obligatory Exhaustivity.” *Semantics and Pragmatics* 7 (0): 11-1–61. https://doi.org/10.3765/sp.7.11.

———. 2016a. “Multivalent Semantics for Vagueness and Presupposition.” *Topoi* 35 (1): 45–55.

———. 2016b. “Presupposition (and Vagueness) Projection at the Propositional Level.” ESSLLI 2016 Course notes on Trivalents Logics and Natural Language Meaning.

Spector, Benjamin, and Paul Egré. 2015. “A Uniform Semantics for Embedded Interrogatives: An Answer, Not Necessarily the Answer.” *Synthese* 192 (6): 1729–84. https://doi.org/10.1007/s11229-015-0722-4.

Sperber, Dan, and Deirdre Wilson. 1986. *Relevance: Communication and Cognition*. Vol. 142. Harvard University Press Cambridge, MA.

———. 2008. “A Deflationary Account of Metaphors.” In *The Cambridge Handbook of Metaphor and Thought*, edited by R. Gibbs, 84–105.

Stanley, Jason. 2000. “Context and Logical Form.” *Linguistics and Philosophy* 23 (4): 391–434. https://doi.org/10.1023/A:1005599312747.

———. 2002. “Making It Articulated.” *Mind & Language* 17 (1–2): 149–68. https://doi.org/10.1111/1468-0017.00193.

Stanley, Jason, and Zoltán Gendler Szabó. 2000. “On Quantifier Domain Restriction.” *Mind & Language* 15 (2–3): 219–261.

Szabolcsi, Anna. 2004. “Positive Polarity – Negative Polarity.” *Natural Language & Linguistic Theory* 22 (2): 409–52. https://doi.org/10.1023/B:NALA.0000015791.00288.43.

Tonhauser, Judith. 2012. “Diagnosing (Not-) at-Issue Content.” In *Proceedings of Semantics of Under-Represented Languages of the Americas (SULA)*, 6:239–254.

Travis, Charles. 1997. “Pragmatics.” In *A Companion to the Philosophy of Language*, edited by Crispin Wright and Bob Hale, 87–107.

———. 2008. *Occasion-Sensitivity: Selected Essays*. Oxford University Press.

Unger, Peter. 1975. *Ignorance: A Case for Scepticism*. OUP Oxford.

Unnsteinsson, Elmar. 2014. “Compositionality and Sandbag Semantics.” *Synthese* 191 (14): 3329–3350.

Westerståhl, Dag. 1984. “Determiners and Context Sets.” In *Generalized Quantifiers in Natural Language*, Johan van Benthem and Alice ter Meulen, 46–71. Walter de Gruyter.

Williamson, Timothy. 1999. “On the Structure of Higher-Order Vagueness.” *Mind* 108 (429): 127–43.

Wilson, Deirdre. 2011. “Parallels and Differences in the Treatment of Metaphor in Relevance Theory and Cognitive Linguistics.” *Intercultural Pragmatics* 8 (2): 196–213.

Wilson, Deirdre, and Robyn Carston. 2007. “A Unitary Approach to Lexical Pragmatics: Relevance, Inference and Ad Hoc Concepts.” In *Pragmatics*, Noel Burton-Roberts, 230–59. London: Palgrave Macmillan.

Wouden, Ton van der. 1997. *Negative Contexts: Collocation, Polarity and Multiple Negation*. London & New York: Routledge.

Yoon, Youngeun. 1996. “Total and Partial Predicates and the Weak and Strong Interpretations.” *Natural Language Semantics* 4 (3): 217–36. https://doi.org/10.1007/BF00372820.

Zweig, Eytan. 2008. “Dependent Plurals and Plural Meaning.” Thesis, New York University. http://lear.unive.it/jspui/handle/11707/4453.

Zwicky, Arnold M., and Jerrold M. Sadock. 1975. “Ambiguity Tests and How to Fail Them.” *Syntax and Semantics* 4: 1–36.