The Logicality of Language: Contextualism vs. Semantic Minimalism *

Guillermo Del Pinal
delpinal@illinois.edu
UIUC
November 20, 2020

forthcoming in *Mind*

Abstract

The Logicality of Language is the hypothesis that the language system has access to a 'natural' logic that can identify and filter out as unacceptable expressions that have trivial meanings—i.e., that are true/false in all possible worlds or situations in which they are defined. This hypothesis helps explain otherwise puzzling patterns concerning the distribution of various functional terms and phrases. Despite its promise, Logicality vastly over-generates unacceptability assignments. Most solutions to this problem rest on specific stipulations about the properties of logical form—roughly, the level of linguistic representation which feeds into the interpretation procedures—and have substantial implications for traditional philosophical disputes about the nature of language. Specifically, Contextualism and Semantic Minimalism, construed as competing hypothesis about the nature and degree of context-sensitivity at the level of logical form, suggest different approaches to the over-generation problem. In this paper, I explore the implications of pairing Logicality with various forms of Contextualism and Semantic Minimalism. I argue that, to adequately solve the over-generation problem, Logicality should be implemented in a constrained Contextualist framework.

Keywords: Logicality, Contextualism, Semantic Minimalism, semantics vs. pragmatics, natural logic, modularity, grammaticality, triviality, quantifiers **Words:** 12,128

1 Introduction

According to the 'generative' tradition in linguistics and philosophy, the human language system consists of a (recursive) structure building device and a compositional interpretation procedure which together determine the class of expressions that belong to a natural language such as English. The 'Logicality of language' is the hypothesis that the language system also includes a kind of 'natural logic' that can perform certain unconscious, automatic inferences (Gajewski 2002, 2008a, Fox 2000, Fox & Hackl 2007, Chierchia 2013, Abrusán 2014, Del Pinal

* For helpful comments and discussions of various drafts of this paper, I'm grateful to Gabe Dupre, Michael Glanzberg, Jeremy Goodman, Eleonore Neufeld, Paul Pietroski, Brian Rabern, Uli Sauerland, and Eric Swanson. This paper was presented at the Mind and Language in LA (MLLA) talk series—I thank the audience for a lively discussion. Thanks also to the editors of *Mind* and two anonymous reviewers for extremely helpful comments and questions. The final paper has been much improved thanks to their suggestions. Finally, I owe special thanks to Gennaro Chierchia for ongoing conversations about Logicality which have served as a constant source of ideas and inspiration for this and related projects.

2019). On this view, the language system can identify and filter-out as strictly unacceptable expressions that, although syntactically well-formed, are uninformative in the sense of being 'trivial', i.e., are either true or false in every world or situation in which they are defined. This hypothesis is motivated by acceptability patterns which capture the distribution of various functional terms and phrases, such as the patterns for quantifiers in (1)-(3), where the sentences in (a)-(b) illustrate an instance of the generalization in (c). The accounts cited in each case show that the target generalization can be derived from (i) reasonable hypotheses about the semantics of functional terms, and (ii) the assumption that expressions which are logically/trivially true or false are marked as strictly unacceptable.

(1) Connected *but*-exceptives: (von Fintel 1993)
  a. *Some student/s but John passed the exam. [trivially false]
  b. No student but John passed the exam.
  c. **Generalization:** Only universal (positive/negative) quantifiers can host *but*-exceptive phrases in their restrictors.

(2) *There*-existentials: (Barwise & Cooper 1981)
  a. *There is every red apple in the garden. [trivially true]
  b. There are some red apples in the garden.
  c. **Generalization:** Only weak determiners can occur in *there*-existential sentences.

(3) Polarity sensitive items: (Chierchia 2013)
  a. *Mary has any marbles. [trivially false]
  b. Mary doesn't have any marbles.
  c. **Generalization:** Negative polarity items such as *any* are only licensed in downward entailing environments.

The reason why each of the marked expressions is trivial is opaque to pre-theoretical reflection. Indeed, the accounts which derive the target generalizations include some of the most elegant and sophisticated analyses in formal semantics. Proponents of Logicality have uncovered many systematic patterns involving expressions which (i) are arguably syntactically well-formed, (ii) can be shown to be trivial, and (iii) are judged as strictly unacceptable. Triviality-based analyses shed light on the distribution of quantifiers, attitude verbs, numerals and exhaustification operators, among other functional terms and phrases.

Despite its considerable empirical payoffs, the Logicality of language hypothesis faces an important challenge, recognized from the outset by its main proponents (Gajewski 2002, Fox & Hackl 2007, Chierchia 2013, Abrusán 2014). If the language system includes a computational system which automatically identifies and filters out as strictly unacceptable expressions which are logically trivial, why are many of the intuitively most obvious examples of tautologies and contradictions, such as those in (4), strictly acceptable (even if sometimes a bit odd)?

(4) Superficial tautologies and contradictions:
  a. If John is a cheater, then John is a cheater.
  b. It is raining and it is not raining.

Why does the 'natural logic' of language identify as trivially false and hence unacceptable expressions like (1a) and (3a) but not the intuitively simpler contradiction in (4b)? Similarly why does the language system filter out as unacceptable trivially true expressions like (2a) but not superficial tautologies like (4a)? Proponents of Logicality have to find a principled way of separating the class of trivial expressions which feel 'ungrammatical' from the class of (superficial) trivialities which are strictly acceptable. Call this the 'over-generation of unacceptability' problem.

The project of finding an implementation of Logicality that addresses the over-generation problem is of considerable theoretical interest. As we will see, solutions to this problem rest on substantive assumptions about the nature of logical form, i.e., about the level of representation that is the input to the interpretation function or procedures. For this reason, the project of finding a viable implementation of Logicality interacts in meaningful ways with traditional philosophical debates between Contextualists and Semantic Minimalists, which are centered on disputes about the nature of logical form. According to Contextualists, most/all terms can be represented as (or can be modified by) characters whose open parameters have to be fixed by context before they can determine an extension given a world/situation (e.g., Carston 2002, Stanley 2007, Recanati 2010, Rothschild & Segal 2009). Minimalists hold, in contrast, that while natural languages have a class of genuinely context-sensitive terms (incl., demonstratives and indexicals), most open-class terms do not have covert context-sensitive parameters (e.g., Borg 2004, Cappelen & Lepore 2005). Logicality can be combined with a Contextualist or a Minimalist conception of logical form—and as we will see, each approach issues in a range of unique yet reasonably promising solutions to the over-generation problem.

To begin to appreciate what is distinctive about Contextualism and Semantic Minimalism, taken as solutions to the over-generation problem, consider how each approach may take advantage of a key difference between the ungrammatical trivialities in (1a)-(3a) and the superficial, acceptable trivialities in (4). This difference depends on distinguishing between 'functional' or 'logical' terms (e.g., *all, few, any, and, but*) and 'content' or 'referential' terms (e.g., *cheater, John, rain, love*). As a first pass (see §6), we can say that functional terms are typically assigned high types, their semantic effect is inference-based, and they make up the 'closed' class vocabulary which shows limited variation within and across languages. Content terms, in contrast, are typically assigned lower types which correspond to entities, events, sets of or relations between members of those basic types, and they make up the 'open' class vocabulary which can change in relatively unconstrained ways within and across languages. Crucially, in cases like (1a)-(3a) the trivialities depend only on the configuration of functional or logical terms (see §2-3 below and Gajewski 2002, 2009, Chierchia 2013, Abrusán 2014, Del Pinal 2019). Yet in cases like (4), their status as trivial also depends on the identity of each token of their content terms. Building on that distinction between strictly unacceptable and acceptable, 'superficial' trivialities, consider a Contextualist and a Minimalist friendly proposal for tackling the over-generation problem. Let us begin with the former:

Logicality + Modulation. The meaning of all content terms (incl., variables which are assigned values of the same types) can be modulated by context-sensitive operators present in logical form. Expressions whose triviality depends on the co-identity of content terms are not seen as trivial because each token can be modulated in slightly different ways in its local context, thereby avoiding triviality. Crucially, modulation over content terms doesn't help rescue expressions whose triviality depends solely on the configuration of their logical/functional terms. For logical terms, unlike content ones, can't be modulated.

For example, (1a) is marked as ungrammatical because modulating the meaning of terms like *student*, *pass* and *exam* doesn't 'rescue' the expression from triviality. In contrast, (4a) is not marked as ungrammatical because modulating each token of *cheater* in slightly different ways rescues the expression from triviality (see §3.1). Crucially, this approach to the over-generation problem is not available to Semantic Minimalists—for it appeals to (semantic) modulation operators over all content terms and variables of any 'referential' types—but other promising approaches are compatible with their core commitments. Consider the following 'syntactic' approach:

Logicality + Syntactic skeletons. There is a level of representation which is sensitive to logical/functional terms, but is blind to the specific semantic value and identity of content terms. Grammatically-relevant triviality is determined at this level. Accordingly, expressions whose triviality depends only on the configuration of logical terms can be proven to be trivial, whereas those whose triviality also depends on seeing the co-identity of their content terms are not seen as trivial. At the (later) stage of processing in which the meaning/identity of content terms is fully represented, there is no rampant (linguistically triggered) context-sensitivity.

From this perspective, (1a) is marked as ungrammatical because we can prove that it is trivial even if we do not know what specific semantic value each of its content terms ultimately receives. In contrast, (4a) is not marked as strictly ungrammatical because, to determine if it is trivial, we need to know whether each token of *cheater* receives the same semantic value—and this is not something that can be determined at the level of syntactic representation in which grammatically relevant trivialities are computed (see §3.2). This syntactic approach was adopted by early proponents of Logicality to tackle the over-generation problem (e.g., Gajewski 2002, Fox & Hackl 2007, Chierchia 2013)

The aim of this paper is to show that a refined version of the Contextualist position of Logicality + Modulation is superior to various implementations of Logicality which are inspired by or compatible with Semantic Minimalism. My argument has two parts. The first argues against popular approaches to the over-generation problem along the lines of Logicality + Syntactic skeletons (§2-§4). The key cases involve acceptable superficial trivialities similar to (4a)-(4b), except that the relevant individual terms or predicates are syntactically co-bound or in some form of anaphoric relation. I will argue that only Logicality + Modulation—according to which logical forms include general modulation operators over content terms and individual/predicate variables—can explain why these variants of superficial trivialities are strictly acceptable. The second part examines three novel, Minimalist-friendly attempts to solve the over-generation problem while avoiding the shortcomings of Logicality + Syntactic skeletons (§5). One proposal is that triviality is checked only within minimal syntactic phases, another is that triviality is determined relative to a specific kind of non-classical natural logic, and a third is that triviality is a result of lexical presuppositions. None of these proposals appeal to semantic modulation operators, or posit any kind of ubiquitous context-sensitivity across the lexicon. While each has advantages, I argue that, ultimately, only Logicality + Modulation can maintain the triviality-based accounts of patterns such as (1)-(3)—and similar generalizations which help capture the distribution of functional terms and phrases—without simultaneously over-generating unacceptability assignments for various kinds of 'superficial', acceptable trivialities. Importantly, it does *not* follow from my argument that *any* version of Contextualism is a suitable partner of Logicality: as already alluded, I will argue that we need a version in which modulation is computed by context-sensitive operators present in logical form and is confined to content terms and variables of the corresponding referential types (cf., Martí 2006, Stanley 2007). Radical Contextualism—roughly, the (popular) view that *all* terms can be modulated to increase the coherence or utility of utterances—has to be rejected, if Logicality is accepted.

2 The Logicality of language: A case study of quantifiers and exceptives

According to Logicality, the language system can identify and filter-out expressions which are trivial, i.e., true or false in all worlds/situations in which they are defined. This hypothesis can be used to derive generalizations, such as those in (1)-(3), which capture the distribution of various functional terms and phrases, yet it should be implemented in a way that avoids the over-generation problem. To evaluate different implementations of Logicality, it will be useful to begin by reviewing one triviality-based analysis in detail. This section presents an influential triviality-based account of exceptive-*but* phrases, due to von Fintel (1993). Additional acceptability patterns and accounts will be discussed in later sections.

The basic contrast concerning which quantifiers can host exceptive-*but* phrases in their restrictors is repeated in (5), and the general acceptability pattern is summarized in (6).

(5) a. *Some student/s but John passed the exam.
b. Every student but John passed the exam.

(6) Generalization:
a. ✓: *every, all, none, no*
b. ✗: the rest

The quantifiers that can and those that can't host exceptive-*but* phrases in their restrictors belong to the same syntactic category—partly for this reason, there is no principled syntactic explanation for the acceptability pattern in (6). In contrast, the class of quantifiers that can host exceptive-*but* phrases share a unique semantic characterization: they are the universal (positive/negative) quantifiers. This characterization provides a clue for deriving the target pattern: formulate a plausible entry for exceptive-*but* and examine how it interacts with universal vs. non-universal quantifiers.

Suppose that expressions like (5a) and (5b) are parsed as in (7). A natural hypothesis is that *but* subtracts the set denoted by its complement from that denoted by the next term it combines with, as captured in (8).

(7) $[[\mathbf{D}\ [A\ [\text{but}\ C]]]\ P]$

(8) $[\![[\mathbf{D}\ [A\ [\text{but}\ C]]]\ P]\!] = 1 \text{ iff } \begin{cases} (i)\ C \neq \emptyset \\ (ii)\ \mathbf{D}(A - C)(P) = 1 \end{cases}$

Applied to (5b), this simple subtraction hypothesis generates the truth-conditions in (9):

(9) $[\![[\text{Every}_{\mathbf{D}}\ [\text{student}_A\ [\text{but John}_C]]]\ \text{passed}_P]\!] = 1$ iff
(i) $\{\text{John}\} \neq \emptyset$
(ii) $\{x : x \text{ is student}\} - \{\text{John}\} \subseteq \{x : x \text{ passed}\}$

Now, consider a word w_1 in which every student including John passed the exam. (5b) is intuitively false w_1. Yet the truth-conditions in (9) predict that it should be true, since in w_1 every student other than John passed the exam. This suggests that a simple subtraction operation, as in (8), can't be the whole semantic contribution of exceptive-*but*. von Fintel (1993) proposes instead an analysis closer to (10), which adds condition (*iii*) to the original subtraction-based entry. This captures the idea that the complement of *but* should be the smallest set one can subtract from the restrictor of **D** while preserving the truth of the quantified statement.

(10) $$[\![[\mathbf{D}\ [A\ [\text{but}\ C]]]\ P]\!] = 1 \text{ iff } \begin{cases} (i)\ C \neq \emptyset \\ (ii)\ \mathbf{D}(A - C)(P) = 1 \\ (iii)\ \underbrace{\forall S[\mathbf{D}(A - S)(P) = 1 \rightarrow C \subseteq S]}_{\text{`the \textit{least} you can take out' condition}} \end{cases}$$

Applied to (5b), the entry in (10) generates the truth-conditions in (11):

(11) $[\![[\text{Every}_{\mathbf{D}}\ [\text{student}_A\ [\text{but John}_C]]]\ \text{passed}_P]\!] = 1$ iff

(*i*) $\{\text{John}\} \neq \emptyset$

(*ii*) $\{x : x \text{ is student}\} - \{\text{John}\} \subseteq \{x : x \text{ passed}\}$

(*iii*) $\forall S[\{x : x \text{ is student}\} - S \subseteq \{x : x \text{ passed}\} \rightarrow \{\text{John}\} \subseteq S]$

Consider again w_1, where every student including John passed. This analysis now correctly predicts that (5b) is false in w_1. For although conditions (*i*)-(*ii*) are obviously satisfied—since every student other than John passed in w_1—(*iii*) isn't. Simply let $S = \emptyset$, then the antecedent of (*iii*) is true while the consequent is false. This analysis also captures cases in which (5b) is intuitively true, such as a world w_2 in which John did *not* pass but every other student did pass. The truth-conditions in (11) are satisfied in w_2. Conditions (*i*)-(*ii*) are satisfied because every student other than John passed, and (*iii*) because any set substituted for S which doesn't include John—such as the empty set or the singleton set of any other student—would make the antecedent false, hence the whole conditional true. It is easy to check that the analysis in (10) assigns appropriate truth-conditions to exceptive-*but* sentences with the other (positive/negative) universal quantifiers, at least in direct instantiations of (7).[1]

1 Two clarifications. First, on this version of von Fintel's (1993) account, the first argument of exceptive-*but* is of type $< e, t >$—i.e., takes characteristic functions of sets of entities. In cases like (11), this requires a type shifting operation from John to {John}. While other compositional routes are explored in von Fintel (1993), all still require that *but* be assigned a high type. Second, (10) is intended to capture the meaning of *but*, not of all exceptive terms/phrases. Indeed, most semanticists think that *except (for)* has a freer/more inclusive distribution than *but*. This is partly explained by assuming that the former doesn't include the 'least you can take out' condition (*iii*). To be sure, Gajewski (2008b), Hirsch (2016) and Crnič (2018) have explored the hypothesis that the 'least you can take out' condition is not directly contributed by *but*; rather, it arises from the interaction between *but* (taken as just a set subtraction operation) and an exhaustification operator. On these views, the difference in the distribution between *but* and *except (for)* is captured by stipulating that while *but* phrases obligatorily trigger exhaustification, *except (for)* phrases trigger it only optionally. For simple sentences like those in (5), these accounts also predict the acceptability pattern in (6), for reasons parallel to those we present below. Accordingly—and because details of the compositional source of (*iii*) don't affect broader issues about how best to implement Logicality—I focus here on von Fintel's original (1993) account.

In addition to capturing the intuitive truth-conditions of acceptable exceptive-*but* sentences, the analysis in (10) is also crucial to derive the acceptability patterns summarized in (6). The key step is to recognize that the universal quantifiers in (6a) are all left-downward-entailing—this is what guarantees that there can be minimal exceptions to the corresponding universal generalizations, and that sentences like (5a) are predicted to have contingent truth-conditions, as we just saw. In contrast, left-upward-entailing quantifiers—e.g., *some,* (*at least*) *three,* (*at least*) *four*, etc.—hosting an exceptive-*but* phrase in their restrictors, always fail to simultaneously satisfy (*i*)-(*iii*), thereby generating truth-conditions that are trivially false.

(12) $\mathbf{D}$ is a left upward entailing quantifier iff $\forall A, A^+, P$ s.t.
$[\![\mathbf{D}]\!](A)(P) = 1$ & $A \subseteq A^+$, $[\![\mathbf{D}]\!](A^+)(P) = 1$

The reason for this is simple. Suppose $[\![\mathbf{D}]\!](A)(P) = 1$ and that the restrictor $A = A^+ - C$, where $C \neq \emptyset$. If $\mathbf{D}$ is left-upward-entailing, it follows that $[\![\mathbf{D}]\!](A^+)(P) = 1$, since $A \subseteq A^+$. That is, one could always have subtracted from A^+ a smaller set than C—including the empty set—and still get a true statement. Accordingly, expressions with left-upward-entailing quantifiers with exceptive-*but* phrases in their restrictors can't satisfy the 'least you can take out' condition (*iii*). Given Logicality, such trivially false expressions are marked as strictly unacceptable.

To illustrate this result—i.e., that left-upward-entailing quantifiers, when hosting exceptive-*but* phrases in their restrictors, generate trivial truth-conditions—consider (5a). Given the account of *but* in (10), (5a) is assigned the truth-conditions in (13):

(13) $[\![[\text{Some}_{\mathbf{D}}\ [\text{student}_A\ [\text{but John}_C]]]\ \text{passed}_P]]\!] = 1$ iff
(*i*) $\{\text{John}\} \neq \emptyset$
(*ii*) $(\{x : x \text{ is student}\} - \{\text{John}\}) \cap \{x : x \text{ passed}\} \neq \emptyset$
(*iii*) $\forall S[(\{x : x \text{ is student}\} - S) \cap \{x : x \text{ passed}\} \neq \emptyset \rightarrow \{\text{John}\} \subseteq S]$

Obviously, these conditions are not satisfied in worlds where no student passed. What we need to check, then, is if they are satisfied in any worlds in which at least some students passed. (*i*)-(*ii*) are only satisfied in worlds in which at least some students other than John passed. Amongst those worlds, there are two cases to check for condition (*iii*): worlds in which John also passed, and worlds in which he didn't. In either case, let $S = \emptyset$. The antecedent of (*iii*) is then true—since some students passed in those worlds—while the consequent is obviously false. Hence in any world in which conditions (*i*)-(*ii*) of (13) are satisfied, the 'least you can take out' condition (*iii*) won't be. Since (5a) is assigned trivial truth-conditions, the Logicality hypothesis correctly predicts that it is marked as unacceptable.

The final class we need to consider is that of left non-monotonic quantifiers such as *exactly 3*, *most*, and *few*. The standard view is that these quantifiers can't host exceptive-*but* phrases in their restrictors, as captured in (14):

(14) a. *Exactly three students but John passed the exam.
b. *Most students but John passed the exam.
c. *Few students but John passed the exam.

Despite some complications, von Fintel's (1993) account also arguably predicts this result. Let us focus on (14a). Given the account of *but* in (10), (14a) is assigned the truth-conditions in (15):

(15) $[\![[\text{Exactly three}_{\mathbf{D}}\ [\text{students}_A\ [\text{but John}_C]]]\ \text{passed}_P]\!] = 1$ iff
(i) $\{\text{John}\} \neq \emptyset$ &
(ii) $\mathbf{card}((\{x : x \text{ is student}\} - \{\text{John}\}) \cap \{x : x \text{ passed}\}) = 3$ &
(iii) $\forall S[\mathbf{card}((\{x : x \text{ is student}\} - S) \cap \{x : x \text{ passed}\}) = 3 \rightarrow \{\text{John}\} \subseteq S]$

The truth-conditions in (15) are clearly not satisfied in worlds where no student passed, as well as in worlds where exactly one, two, or at least four students (excluding John) passed. To determine if they are trivially false, we have to check if there are any worlds which satisfy them. There are two relevant remaining cases to consider. The first consists of worlds where exactly three students passed and John is not in that set, i.e., he did not pass. This would satisfy conditions (i)-(ii), but not (iii). For let $S = \emptyset$, then the antecedent of (iii) is true while the consequent is false. The second consists of worlds where exactly three students other than John passed and John also passed. This would again satisfy conditions (i)-(ii): the set of students subtracting {John} includes exactly three that passed. But it again fails condition (iii). For let S equal any singleton set containing any student other than John who passed, then the antecedent of (iii) is true but the consequent is false, since {John} is not a subset of any of those singleton sets. It follows that the truth-conditions of (14a) in (15) are trivially false, so by Logicality (14a) is correctly predicted to be marked as unacceptable.[2]

---

2 *Most* and *few*-quantified sentences with *but*-phrases in their restrictors, such as (14b) and (14c), present additional complications. For brevity, I focus on the case of *most* (the case of *few* is quite similar). Given the entry for *but* in (10)—and assuming *most* means 'more than half'—(14b) is assigned the truth-conditions in (A):

(A) $[\![[\text{Most}_{\mathbf{D}}\ [\text{students}_A\ [\text{but John}_C]]]\ \text{passed}_P]\!] = 1$ iff
(i) $\{\text{John}\} \neq \emptyset$
(ii) $\mathbf{card}((\{x : x \text{ is student}\} - \{\text{John}\}) \cap \{x : x \text{ passed}\}) > 1/2\mathbf{card}(\{x : x \text{ is student}\} - \{\text{John}\})$
(iii) $\forall S[\mathbf{card}((\{x : x \text{ is student}\} - S) \cap \{x : x \text{ passed}\}) > 1/2\mathbf{card}(\{x : x \text{ is student}\} - S) \rightarrow \{\text{John}\} \subseteq S]$

It is easy to check that most situations don't satisfy conditions (i)-(iii). So just like the corresponding *exactly n* sentences, (14b) is predicted to come out as false in general, a desirable result insofar as we are trying to show that (14b) is unacceptable because it has trivial truth-conditions. However, there is a type of situation in which the conditions in (A) are satisfied. Suppose there are just two students, incl. John, and that only John failed. (i)-(ii) are satisfied because the cardinality of the set of students excluding John who passed is greater than that of half the set of students excluding John. (iii) is satisfied because if $S = \emptyset$, the antecedent of (iii) is false (since one student failed and one passed), and if S is the singleton set of the other student, the antecedent of (iii) is again false (since John, the only other student, did not pass). Either way, the conditional in (iii) comes out true. It follows that, on this account, sentences like (14b) may be true but only when their restrictor is a singleton set.

However, building on Heim (1991), Hirsch (2016), a.o., has noted that *most*-sentences seem to be infelicitous when interpreters know or presuppose that they have a singleton set as a restrictor (either of individuals or pluralities). This observation is motivated—independently of our target sentences—by examples like *#most tallest student/s in the class passed* and *#most of my parents came to visit*. One way of accounting for this observation is to argue that (under certain conditions) *most*-sentences

Summing up, using independently justified entries for the relevant functional terms, we have identified a semantically definable class of quantificational determiners which can host exceptive-*but* phrases in their restrictors. Specifically, we have shown that the (positive/negative) universal quantifiers, which are left downward entailing, can host *but* phrases in their restrictors without generating trivial readings. We also showed that, in contrast, left-upward entailing quantifiers, and arguably also the left non-monotonic ones, generate trivially false readings when hosting *but* phrases in their restrictors. Based on those results, we can derive the distributional generalization in (6) concerning the interaction between quantifiers and exceptive-*but* phrases if we adopt Logicality, i.e., the hypothesis that sentences with trivial truth-conditions are identified and marked as unacceptable by the language system. Following Fox & Hackl (2007), let us call the computational system that can identify and filter out such grammatically relevant trivial expressions the 'Deductive System' (DS).

3 The over-generation problem and Contextualist vs. Minimalist conceptions of logical form

Logicality supports elegant accounts of the distribution of quantifiers and many other functional terms and phrases. The problem for any triviality-based account, however, is that many superficial tautologies and contradictions, such as those in (4), are strictly acceptable. This is unexpected if the language system includes a DS that automatically filters out trivial expressions. Can we implement Logicality so that the DS doesn't over-generate assignments of triviality, hence of strict unacceptability? Call 'L-trivial' the set of expressions that is predicted to be strictly unacceptable relative to each solution of the over-generation problem. One way of approaching the over-generation problem is by examining different assumptions about logical form, specifically, about the properties of the linguistic representations 'seen' by the DS. This is where Contextualism and Semantic Minimalism enter the discussion, since they issue in distinctive hypotheses about the nature of logical form. In this section, I present what I believe are the most promising ways of pairing each of Contextualism and Semantic Minimalism with Logicality. Although both proposals help with the over-generation problem, I will argue in §4 that only the Contextualist approach provides a fully general solution.

trigger an obligatory 'not all' implicature, which would generate a contradiction whenever the target restrictor is a singleton set. Alternatively, we could add a singleton set ban as a presupposition on the restrictor of *most*. Either way, we would block the one type of situation in which the truth-conditions in (A) can be satisfied, and predict that (14b) comes out (whenever defined) as trivially false, hence is marked as unacceptable. Parallel issues (and solutions) apply to *few*-quantified sentences like (14c).

One final concern. Although the usual judgment amongst linguists working on connected exceptives is that (14b)-(14c) are indeed unacceptable (see e.g. von Fintel 1993, Gajewski 2008b, Hirsch 2016, Crnič 2018), a reviewer reports that (14b)-(14c) feel kind of acceptable, even if a bit odd, and to have the truth-conditions that would be assigned if we use the bare subtraction entry for *but* in (8). Assuming von Fintel (1993)'s account, this is not an expected pattern of judgments—yet it is also not entirely surprising. According to von Fintel (1993), a bare set subtraction operator is part of the functional (fixed) repertoire of natural languages. Although in English it is usually lexicalized by 'free' exceptives such as *except (for)*, it is possible that in some idiolects, or stages of grammaticalization, it is also lexicalized with *but* (while triggering only 'optional' exhaustification).

3.1 Contextualism as Logicality + Modulated Logical Forms

Consider first the version of Contextualism which I propose to adopt as an implementation of Logicality. The goal here is to begin to illustrate how it addresses the over-generation problem—the full justification for all the components of this account will emerge gradually as we discuss, in later sections, additional acceptability patterns. On this version of Contextualism modulation is performed by an operator, $\mathcal{R}$, present in logical form. $\mathcal{R}$ is a polymorphic type operator that is generated as a sister to all and only content terms and variables that can be assigned any 'referential' types (i.e., individual and predicate variables) (cf. Del Pinal 2019, Chierchia 2019). The resulting hypothesis is schematically captured in (16):

(16) **Logicality + Modulated logical forms**

a. Language and its DS 'see' modulated LFs: representations like standard LFs except that all non-logical terms are arguments of $\mathcal{R}$ operators. If an expression can't be 'rescued' from triviality by possible modulations of $\mathcal{R}$ operators it is marked as unacceptable.

b. To obtain a modulated LF for α

(i) Identify the minimal projections of any content terms and (individual and predicate) variables of α (any 'referential' points).

(ii) Add $\mathcal{R}$ as a sister.

On this view, the DS interacts with modulated logical forms. These representations involve a covert $\mathcal{R}$ operator—a character interpreted in its *local context*—which attaches to all content terms and variables and can modulate their meaning.[3] The class of content terms and variables consists of open class terms such as *John* and *red* and individual and predicate-type variables (for refinements, see §6). Although $\mathcal{R}$ is obligatorily inserted in its licensed positions, it can be lazy: i.e., it can compute the identity function, which results in a kind of vacuous modulation. The modulated logical forms of some basic examples of unacceptable vs. (superficial) acceptable trivialities can be represented roughly as follows:

(17) *Some students but the lazy ones passed the exam.

a. Modulated LF:
[[[Some [$\mathcal{R}_{c'}$(students) [but the $\mathcal{R}_{c''}$(lazy ones)]]] [$\mathcal{R}_{c'''}$(passed)]]

(18) It is raining and it is not raining.

a. Modulated LF:
[[It is $\mathcal{R}_{c'}$(raining)] [and [it is not $\mathcal{R}_{c''}$(raining)]]]

3 Logicality + Modulated LFs builds on constrained Contextualist accounts in which modulation operators are present in logical form and operate only on non-logical terms (e.g., Szabó & Stanley 2000, Stanley 2007, Martí 2006, Sauerland 2014). Unlike radical Contextualism, Logicality + Modulated LFs is compatible with the hypothesis that the language system is relatively modular, and also with the standard compositional explanations of systematicity and productivity. Indeed, we can stipulate that the expressive power of $\mathcal{R}$ is rather constrained, although this approach is compatible with various implementations of content class modulation, incl., versions somewhat similar to those recently explored by Abrusán et al. (2019, 2018), except that on my view modulation over functional/logical terms should be categorically ruled out.

Given modulated logical forms, the subset of the L-trivial sentences—the trivial sentences that are marked as strictly unacceptable—can be defined as follows:

(19) **L-triviality with modulated logical forms**:

a. A sentence is L-trivial iff (whenever defined) it comes out as uniformly true/false for every modulation available to each instance of $\mathcal{R}$. L-trivial sentences are marked as 'ungrammatical'

b. A sentence is trivial iff (whenever defined) it comes out as uniformly true/false for the default value (the identity map) of modulations. Trivial but not L-trivial sentences are not marked as 'ungrammatical' by the DS.

To see why Logicality + Modulated LFs helps with the over-generation problem, let us examine how it derives that observed acceptability patterns for sentences with exceptive-*but* phrases and for our basic examples of superficial, acceptable trivialities. Let us begin with the latter, simpler case. It is easy to see that, on this account, superficial trivialities like (18) do not come out as L-trivial. In this specific case, each token of *rain* can be modulated in a slightly different way given its local context, generating readings like 'it is raining but it is not raining hard'. In general, modulated logical forms rescue from L-triviality many 'superficial' (and intuitively acceptable) tautologies and contradictions, for in all these cases their triviality depends on computing just the identity function over each token of their non-logical terms.

The next step is to show that Logicality + Modulated LFs makes the right predictions for the acceptability patterns with exceptive-*but* phrases. The basic contrast is repeated in (20) and (21). It is easy to see that (20a) comes out as contingent. For one possible modulation of each token of $\mathcal{R}$ is the identity function, and in that case (20a) is contingent: e.g., it is true in worlds in which all the non-lazy students passed and the lazy ones did not pass, and false in worlds in which all the students failed. Consider next (21), given its modulated LF in (21a). The aim is to show that we can't 'over-rescue' in this kind of case. Applying the entry for exceptive *but* in (10), we get the interpretation in (21b). From this we can see that (21a) is false whenever $\mathcal{R}_{c''}(\llbracket \textit{lazy students} \rrbracket) = \emptyset$ (or undefined, if condition (i) is treated as a presupposition). Accordingly, pick any modulation for $\mathcal{R}_{c''}$ that restricts the relevant set of lazy students, so long as the resulting set is not empty, the aim being to see if we can find a modulation under which it comes out as true. Since *some* is left-upward entailing, we can always subtract less than $\mathcal{R}_{c''}(\llbracket \textit{lazy students} \rrbracket)$, whatever that is, since we can simply subtract the empty set. As a result, the 'least you can take out condition' (= condition (iii)) of exceptive-*but* is necessarily violated, and (21a) can't come out as true.[4]

(20) All students but the lazy ones passed the exam.

a. Modulated LF:
[[All [$\mathcal{R}_{c'}$(students) [but the $\mathcal{R}_{c''}$(lazy ones)]]] [$\mathcal{R}_{c'''}$(passed)]]

4 We said earlier that $\mathcal{R}$ can 'restrict' (move to a subset of the set denoted by its argument) but also 'loosen' (move to a superset of the set denoted by its argument) interpretations. It is easy to see that cases like (21) cannot be rescued when $\mathcal{R}$ 'loosens' interpretations: if the set of lazy students is *not* the *smallest* set one can subtract from the set of students while maintaining truth, then, a fortiori, no superset of that set will be smallest set one can subtract from the set of students while maintaining truth.

b. Interpretation: [Contingent]

$$= 1 \text{ iff } \begin{cases} (i)\ [\![\mathcal{R}_{c''}(lazy\ ones)]\!] \neq \emptyset\ \wedge \\ (ii)\ [\![All]\!]([\![\mathcal{R}_{c'}(students)]\!] - [\![\mathcal{R}_{c''}(lazy\ students)]\!])([\![passed]\!]) = 1\ \wedge \\ (iii)\ \forall S([\![All]\!]([\![\mathcal{R}_{c'}(students)]\!] - S)([\![\mathcal{R}_{c'''}(passed)]\!]) = 1 \\ \rightarrow [\![\mathcal{R}_{c''}(lazy\ students)]\!] \subseteq S) \end{cases}$$

(21) *Some students but the lazy ones passed the exam.

a. Modulated LF:
[[Some [$\mathcal{R}_{c'}$(students) [but the $\mathcal{R}_{c''}$(lazy ones)]]] $\mathcal{R}_{c'''}$(passed)]

b. Interpretation: [Trivially false]

$$= 1 \text{ iff } \begin{cases} (i)\ [\![\mathcal{R}_{c''}(lazy\ ones)]\!] \neq \emptyset\ \wedge \\ (ii)\ [\![some]\!]([\![\mathcal{R}_{c'}(students)]\!] - [\![\mathcal{R}_{c''}(lazy\ students)]\!])([\![passed]\!]) = 1\ \wedge \\ (iii)\ \forall S([\![some]\!]([\![\mathcal{R}_{c'}(students)]\!] - S)([\![\mathcal{R}_{c'''}(passed)]\!]) = 1 \\ \rightarrow [\![\mathcal{R}_{c''}(lazy\ students)]\!] \subseteq S) \end{cases}$$

Summing up, Logicality + Modulated LFs issues in a promising solution to the over-generation problem: while standard examples of superficial trivialities don't come out as L-trivial, the unacceptable examples with exceptive-*but* do come out as L-trivial. This approach also preserves the L-triviality-based accounts of the other acceptability patterns in (1)-(3) (see Del Pinal 2019, Chierchia 2019), and supports various additional applications of Logicality that we discuss in §4-§5. Before introducing competing approaches to the over-generation problem compatible with Semantic Minimalism, let me clarify how Logicality + Modulated LFs relates to broader Contextualist approaches to logical form.

At this point, it is easy to see why not just any version of Contextualism will work as a suitable partner of Logicality. Specifically, radical versions of Contextualism in which *all* terms are subject to modulation (cf. Carston 2002, Recanati 2004, 2010) systematically under-generate assignments of unacceptability in the kinds of cases considered here. Suppose that the meaning of any term, including functional/logical ones, could be modulated so as to increase the utility of assertions (where rescuing an assertion from strict unacceptability would be a special case of this function). On this view, we could parse an exceptive sentence like (22) as in (22a), i.e., with a modulation operator over exceptive *but* (I omit other possible modulations for simplicity). We have seen that what makes left-upward entailing quantifiers such as *some* generate trivial readings in these sentences is the 'least you can take out' condition (=(*iii*)) of *but*. Accordingly, we could rescue assertions of (22) and the like from triviality via a modulation operation that simply drops that condition, and outputs a bare set subtraction meaning, as captured in (22b).

(22) *Some students but the lazy ones passed the exam.

a. Modulated LF:
[[Some [students [$\mathcal{R}_{c'}$(but) the lazy ones]]] passed]

b. Interpretation:

$$= 1 \text{ iff } \begin{cases} (i)\ [\![lazy\ ones]\!] \neq \emptyset\ \wedge \\ (ii)\ [\![some]\!]([\![students]\!] - [\![lazy\ students]\!])([\![passed]\!]) = 1 \end{cases}$$

In this case, (22) comes out as contingent—e.g., it is false in worlds in which no students passed, and true in worlds in which at least one student who is not amongst the lazy ones passed (including worlds in which all the lazy students also passed). Accordingly, allowing modulation over functional terms would result in the incorrect prediction that (22) has an acceptable reading, paraphrasable as 'at least some students who are not amongst the lazy ones passed the exam'. In general, positions that allow for modulation to operate over functional terms make systematically incorrect 'over-rescuing' predictions. This result is important because radical Contextualism remains an attractive position amongst philosophers of language—yet it is simply not a viable position for those who also accept Logicality.

Logicality + Modulated LFs qualifies as a constrained version of Contextualism due to two core features of the modulation operator $\mathcal{R}$. First, $\mathcal{R}$ is attached exclusively to non-logical terms, where the target class includes content terms like *John* and *red* and individual and predicate variables. Second, $\mathcal{R}$ may result in non-trivial modulations—including those that 'rescue' expressions which would otherwise be informationally useless—but can also simply compute the identity function. Insofar as we go for a moderate form of Contextualism, these constrains seem natural, and as we will see in §4, both have desirable empirical consequences. Radical Contextualist may hold that there is no strict distinction between logical and non-logical terms relevant to modulation, and/or that interpreters are required to non-trivially modulate all token uses of (non-logical) terms. Still, we can explore a constrained Contextualist approach in which there is a distinction between logical and non-logical terms which affects the domain on modulation, and in which terms that in principle may be modulated are only non-trivially modulated under certain conditions. On this approach, functional/logical terms such as quantifiers, coordinators and modals form a relatively closed class system such that token uses of them can't be synchronically modulated, while content terms such as nouns, verbs and variables of the same semantic types form a relatively open system such that uses of them can be synchronically modulated to increase the coherence/informativeness of utterances. How to precisely separate the logical and non-logical terms is of course a difficult issue; yet it is one that, as we will see, any viable approach to the over-generation problem needs to face (we revisit this issue in §6).

3.2 Semantic Minimalism as Logicality + Skeletons

Consider next a Minimalist-friendly notion of logical form that can be paired with Logicality to tackle the over-generation problem. The key stipulation, due to Gajewski (2002), is that the DS operates over a level of representation—called 'logical skeletons'—that is 'blind' to the identity and specific content of all non-logical terms (see also Fox & Hackl 2007, Gajewski 2009, Chierchia 2006, 2013, Abrusán 2011). The resulting package—quite popular amongst proponents of Logicality as an approach to the over-generation problem—is captured in (23):

(23) **Logicality + Logical Skeletons**

a. Language and its DS 'see' only 'logical skeletons': representations that are underspecified with respect to the meaning/identity of their non-logical terms. Expressions whose skeletons can be proven trivial are marked as unacceptable.

b. To obtain the logical skeleton of an LF α

(i) Identify the maximal constituents of α containing no logical terms

(ii) Replace each such constituent with a fresh constant of the same type.

On this view, the DS is radically 'blind' to all non-logical terms—crucially, it does not even see when two content tokens are tokens of the same content/non-logical term. Accordingly, the logical skeletons of our basic target examples would look roughly as in (24a) and (25a). Note that in (25a) each token of *rain* is replaced with a new variable of the same type.

(24) *Some students but John passed the exam

a. Skeleton:
[[Some [$\mathrm{P}_{<e,t>}$s [but $\mathrm{S}_{<e>}$]]] [$\mathrm{V}_{<e,<e,t>}$-ed the $\mathrm{E}_{<e,t>}$]]

(25) It is raining and it is not raining.

a. Skeleton:
[[It is $\mathrm{P}_{<e,t>}$-ing] [and [it is not $\mathrm{R}_{<e,t>}$-ing]]]

To complete this account, we have to specify which subset of the trivial sentences are marked as unacceptable, i.e., we have to define the set of 'L-trivial' sentences:

(26) **L-triviality with Skeletons**

(i) A sentence is L-trivial iff its logical skeleton = 1 (or 0) for *all* interpretations in which it is defined.
(ii) A sentence is strictly unacceptable if it is or contains an L-trivial sentence.

Logical skeletons correspond (roughly) to a level of syntactic representation advanced on independent grounds by work in distributional morphology, according to which content/open-class terms are inserted 'late' in the derivation process (Marantz 1994, Harley 2014). From this perspective, it is not ad hoc to stipulate that the DS applies at a stage of processing in which only the functional skeleton of expressions is explicitly represented.

To see why Logicality + Skeletons helps with the over-generation problem, consider why it predicts that acceptable, superficial trivialities are not L-trivial while still supporting the triviality-based account of acceptability patterns with exceptive-*but* sentences. It is easy to check that, based on their skeletons, superficial trivialities such as (25) do not come out as L-trivial, hence are not marked as strictly unacceptable. Simply consider an interpretation of P and R in which they are not equivalent. Logicality + Skeletons also makes the right predictions for the target patterns with exceptive-*but* sentences. The basic contrast is repeated in (27) and (28). The skeleton in (27a) can come out as true or false depending on the values assigned to P_1, P_2 and P_3. In contrast, the skeleton in (28a) can never come out as true. Given the entry for *but* in (10), $I(P_2) \neq \emptyset$ (= condition (i)), otherwise (28) is false (or a presupposition failure). Against that constraint, take any interpretation of $P_1 \ldots P_3$ which satisfies (ii). To check if the 'least you can take out condition' (= condition (iii)) can be satisfied, let $S = \emptyset$. Since *some* is left-upward-entailing, the antecedent of (iii) will be satisfied while the consequent is false for any non-empty interpretation of P_2. Accordingly, (28) comes out as trivially false (whenever defined), and is correctly marked as unacceptable.

(27) Every student but John passed the exam

a. Logical skeleton: [[*every* [P_1 [*but* P_2]]] P_3]
b. Interpretation: [Contingent]

$$= 1 \text{ iff } \begin{cases} (i)\ I(P_2) \neq \emptyset\ \wedge \\ (ii)\ [\![every]\!](I(P_1) - I(P_2))(I(P_3)) = 1\ \wedge \\ (iii)\ \forall S[[\![every]\!](I(P_1) - S)(I(P_3)) = 1 \rightarrow I(P_2) \subseteq S] \end{cases}$$

(28) *Some students but John passed the exam.

a. Logical skeleton: $[[some\ [P_1\ [but\ P_2]]]\ P_3]$

b. Interpretation: [Trivially false]

$$= 1 \text{ iff } \begin{cases} (i)\ I(P_2) \neq \emptyset \wedge \\ (ii)\ [\![some]\!](I(P_1) - I(P_2))(I(P_3)) = 1 \wedge \\ (iii)\ \forall S[[\![some]\!](I(P_1) - S)(I(P_3)) = 1 \rightarrow I(P_2) \subseteq S] \end{cases}$$

Summing up, we have seen that if we assume that the system that searches for 'trivialities' runs on skeletons, we capture (roughly) the correct distinction between superficial trivialities and strictly unacceptable L-trivialities. Although again we have derived this result for only one case—connected *but*-exceptives—the acceptability patterns for the other cases in (1)-(3) can also arguably be derived from the triviality vs. contingency of the corresponding logical skeletons (see Gajewski 2002, 2008a, 2009, Chierchia 2013).[5] In addition, Logicality + Skeletons need not posit that all or most content terms include genuine context-sensitive parameters. In other words, holding that there is a level of processing in which the identity of content/open-class terms is ignored is compatible with holding that, once the identity and semantic values of these terms are recovered, most of them don't have any context-sensitive parameters. Logicality + Skeletons, then, is a reasonable approach to over-generation which postulates a grammatically-relevant level of representation that is fully compatible with the commitments of Semantic Minimalism.

4 Modulated LFs vs. Skeletons: Superficial trivialities with bound variables

We have seen that pairing Logicality with either Modulated Logical Forms (the Contextualist-friendly option) or Logical Skeletons (the Minimalist-friendly option) helps with the over-generation problem. Specifically, each package can explain why some basic cases of superficially trivial sentences are not marked as strictly unacceptable, while supporting the derivation of L-triviality for the target expressions in acceptability patterns, such as (1)-(3), which capture the distribution of various kinds of functional terms and phrases. However, there is a class of cases, to which we now turn, that can discriminate between those proposals. The key examples are similar to superficial contradictions and tautologies, except that the content terms that generate the trivialities are either syntactically co-bound or in some kind of anaphoric dependency relation. I will argue that, for these kinds of superficial trivialities, Logicality + Skeletons, but not Logicality + Modulated LFs, systematically over-generates unacceptability assignments.

4.1 Predicate co-binding in superficial trivialities

According to Skeletons, the identity of content terms is not encoded at the level of representation accessible to the DS: e.g., superficial contradictions like *It is raining and not raining* are 'seen' roughly as *It is P and it is not Q*. This helps explain why some superficial trivialities

5 Yet note that Del Pinal (2019) argues that Logicality + Skeletons has difficulties supporting L-triviality-based accounts of negative polarity items, such as the one defended in Chierchia (2013) (see also §5.2 below). Abrusán (2014) also discusses various acceptability patterns—which arguably call for triviality-based explanation—that are hard to capture based on the skeletons of the relevant expressions.

are acceptable. However, we can construct acceptable superficial trivialities which induce various kinds of syntactic co-dependencies between the target content terms. Crucially, it is hard to deny that these kinds of co-dependencies—especially binding relations—are encoded in the functional skeletons of expressions. In these cases, L-triviality can be proven from their corresponding skeletons, but not, I will argue, from their modulated LFs.

Consider first superficially trivial expressions with co-bound predicate variables, such as (29). Suppose the level of representation where grammaticality is determined is blind to the identity of non-logical terms like *smart*. Still, the structure with co-binding ensures that, whichever specific predicate is ultimately selected, it must co-occur in both conjuncts, as captured in (29a). Since (29a) is L-trivial, (29) is incorrectly predicted to feel ungrammatical. In contrast, consider the modulated LF of (29) in (29b). Since each instance of the co-bound predicate can be modulated in slightly different ways, the DS can't derive L-triviality, and (29) is correctly predicted to feel strictly acceptable.[6]

(29) Smart is what John is and isn't.

a. P is [what$_1$ John is t$_1$ and is not t$_1$]

b. Smart is [what$_1$ John is $\mathcal{R}_{c'}(\mathrm{t}_1)$ and is not $\mathcal{R}_{c''}(\mathrm{t}_1)$]

To make the same point with a slightly different example, consider the embedded question in (30). Since its syntactic skeleton has to encode binding information, as captured in (30a), L-triviality can be easily derived. So (30) is incorrectly predicted to be strictly unacceptable. In contrast, L-triviality is not derivable from the modulated LF for (30), captured in (30c), which can support modulated meanings such as (30d):

(30) I wonder what John is and isn't...

a. what$_1$ John is t$_1$ and is not t$_1$

b. $[\![$(30a)$]\!] \approx$ {p: $\exists$Q[p = John is Q and John is not Q]}

c. what$_1$ John is $\mathcal{R}_{c'}(\mathrm{t}_1)$ and is not $\mathcal{R}_{c''}(\mathrm{t}_1)$

d. $[\![$(30c)$]\!] \approx$ {John is a typical cousin and not a good cousin, John is a typical friend and not a good friend, John is a typical partner and not a good partner, ...}

These kinds of superficial trivialities with predicate co-binding are not just strictly acceptable—in some cases, they are easy to produce and interpret. Imagine that Mary and Peter are perplexed by John's recent selfish behavior:

(31) a. Peter: I wonder what John is and is not ...

b. Mary: A friend ...

In this case, Mary's assertion can be naturally interpreted as saying that John is a friend in one sense, but also not a friend in some other, perhaps deeper sense.

6 A defender of Skeletons can respond that the meaning of *smart* is a character whose parameters have to be saturated in its local context. This kind of response might work for some examples of superficial trivialities with co-binding, such as (29), but it is not a generalizable strategy for Semantic Minimalists. For we can easily construct examples that are structurally like (29) except that the co-bound predicates are not, given basic Minimalist commitments, context-sensitive characters/terms.

Interestingly, given standard accounts of ellipsis and other forms of de-accentuation,[7] even quite simple variants of our original superficial trivialities arguably present a challenge to Skeletons, as pointed out by Sauerland (2017), but not to Modulated LFs. To see why, assume a structural account of ellipsis, according to which elided material is subject to some kind of (anaphoric) syntactic and/or semantic identity constraint, as captured in (32a)-(32b). Note that our original examples of superficial trivialities, such as (33a), don't involve ellipsis or de-accentuation. It is thus reasonable to assume that no identity condition is explicitly imposed over the tokens of the predicates which generate the superficial contradiction. This means that we can generate a skeleton as in (33b), which is not L-trivial.

(32) a. Jasmine is smart but John isn't.
b. Jasmine is smart$_1$ but John isn't ~~smart$_1$~~

(33) a. John is smart and he isn't smart.
b. John is P and he isn't Q

Yet consider simple variants of (33a) with ellipsis, such as the superficial trivialities in (34a) and (35a). The problem for Skeletons is that the syntactic licensing condition has to encode the information that the elided predicate is anaphoric or copied from the non-elided one, as captured in (34b) and (35b). That is, logical/functional skeletons must encode that co-identity information, even if the specific interpretation of the predicate is ignored at this level. The problem, of course, is that structures like (34b) and (35b) are L-trivial. In contrast, the corresponding modulated LFs in (34c) and (35c) meet the syntactic/semantic identity condition on the elided predicate, do not come out as L-trivial, and can be used to explain why these expressions support a reading like 'John is smart in some sense and isn't smart in some other sense'.[8]

(34) a. John is and isn't smart.
b. John is ~~smart$_1$~~ and isn't smart$_1$
c. John is $\mathcal{R}_{c'}$(~~smart$_1$~~) and isn't $\mathcal{R}_{c''}$(smart$_1$)

(35) a. John is smart, but he also isn't.

7 See, e.g., Rooth (1992); for a survey of recent accounts of ellipsis, see Merchant (2019).

8 This analysis seems to assume that, in cases like (34a) and (35a), $\mathcal{R}$ is not part of the copied material, which some might find questionable. Yet even if we assume that the syntactic and/or semantic identity condition applies also to $\mathcal{R}$, we would still get the same result. This is because, in general, elided context-sensitive terms are interpreted in their local context as determined by their LF position. To see this, consider examples like (ia)-(ib):

(i) a. Serena Williams is a great tennis player, and you are one as well.
b. The dutch basketball team is very tall, and so is their football team.

A coach can assert (ia), to motivate a junior player, and use different standards for what counts as a 'great player' for Serena Williams vs. for junior players. Similarly, a fan can assert (ib) and use different standards for what it is to count as a 'tall team' for a basketball vs. a football team. This kind of flexibility is systematically exhibited by elided (context-sensitive) material. While this observation is compatible with a syntactic identity condition, it suggests that a semantic identity condition should in the first instance apply to characters. Either way, we can maintain our basic account of why the modulated LFs of the target superficial trivialities with ellipsis are not L-trivial (see also §5.2).

b. John is smart_1 but he also isn't $\cancel{\text{smart}_1}$
c. John is $\mathcal{R}_{c'}(\text{smart}_1)$ but he also isn't $\mathcal{R}_{c''}(\cancel{\text{smart}_1})$

4.2 Reflexives in superficial trivialities

Logicality + Skeletons, but not Logicality + Modulated LFs, also over-generates unacceptability assignments for superficial trivialities with reflexives (Chierchia 2019).[9] Consider the deceptively simple example in (36). Assuming a bound variable account of reflexives—according to which reflexives have to be bound in their local syntactic environment (Chomsky 1981, Heim & Kratzer 1998)—(36) has an LF as in (36a). Due to the presence of binding, (36a) can be easily shown to be L-trivial, even given its skeleton. As a result, Skeletons incorrectly predicts that (36) is unacceptable.[10] In contrast, Modulated LFs says that $\mathcal{R}$ is triggered as a sister of any referential type, including variables of type e (or $< s, e >$). Accordingly, the modulated LF for (36) is roughly as in (36b), which is not L-trivial because the modulation can be different at each local context for $\mathcal{R}$.

(36) John is not himself
a. John $\lambda x_i[x_i$ is not $\text{himself}_i]$
b. John $\lambda x_i[\mathcal{R}_{c'}(x_i)$ is not $\mathcal{R}_{c''}(\text{himself}_i)]$
c. John is not behaving (today) the way he usually behaves.

The hypothesis that (36) has the modulated LF in (36b) helps explain why it can get the reading paraphrased in (36c). If we assume, for simplicity, that proper names and variables over individuals are of type $< e >$, this means that $\mathcal{R}$ can map individuals of type $< e >$ into individual concepts of type $< s, e >$. In the case at hand, $\mathcal{R}_{c''}$ maps John to an individual level concept like 'the individual that behaves (in the current situation) most similarly to how John usually behaves' (while $\mathcal{R}_{c'}$ is 'lazy', i.e., is resolved to the identity function).[11]

Like superficial trivialities with predicate co-binding, acceptable superficial trivialities with reflexives occur in many kinds of constructions. Consider reflexives in comparatives

9 The problem of acceptable, superficial trivialities with reflexives is briefly discussed by Gajewski (2009), focusing on Skeletons. Del Pinal (2019) tried to deal with these cases without assuming that modulation applies to variables over individuals. The account I present below builds on the recent proposal by Chierchia (2019) to extend the domain of modulation to variables over individuals.

10 Could proponents of Skeletons reply that the DS is also blind to the English copula *is* and treats it as one amongst various other possible relations? That is, can Skeletons treat the copula as an open-class term? This is implausible (see Gajewski 2002, 2009, Abrusán 2014, Del Pinal 2019, Chierchia 2019). First, the copula is syntactically a prototypical functional item. Second, semantic criteria such as identity under domain permutations classify identity as an unambiguous logical constant. Third, treating identity as a non-logical term doesn't help with variants of the basic cases which don't involve the use of the copula, such as superficial trivialities with reflexives in comparatives (discussed below).

11 Chierchia (2019) argues that this application of modulation over individuals is independently supported by an influential approach to de re (and de se) belief. Briefly, the challenge in the de re case is to explain why (ia) can be used to express a non-contradictory belief of John towards his actual brother, appropriate to scenarios like (ib).

(i) a. John believes that his brother is not his brother.
b. John believes that his actual brother is in fact an impostor trying to steal John's inheritance.

such as (37). For the analysis, assume a degree semantics for comparatives where adjectives correspond to relations between individuals and degrees (e.g., Kennedy 2007). The standard LF in (37a) is L-trivial, and generating its skeleton doesn't help due to the presence of the reflexive. In contrast, the modulated LF in (37b) is not L-trivial, which is the desired result.

(37) John was more eloquent than himself.

a. John$_i$ $\lambda x_i[x_i$ was MORE(eloquent) than himself$_i]$
$= \lambda x_i[\text{MORE(eloquent)}(x_i)(x_i)]$ (John$_i$)
where for any u, MORE(eloquent)(u) is the property of being more eloquent than u defined as follows:
u' has the property of being more eloquent than u iff there is some degree d such that u' is at least d-eloquent and u is not d-eloquent.

b. John$_i$ $\lambda x_i[\mathcal{R}_{c'}(x_i)$ was MORE(eloquent) than $\mathcal{R}_{c''}$(himself$_i$)$]$
$= \lambda x_i[\text{MORE(eloquent)}(\mathcal{R}_{c'}(x_i))(\mathcal{R}_{c''}(x_i))]$ (John$_i$)

Superficial trivialities with reflexives such as (37) are not only strictly acceptable but even quite easy to interpret. Suppose that it is common ground between Mary and Peter that John's speeches are usually quite bad. One odd Monday, however, John's speech was amazing, but only Mary was present:

(38) a. Peter: How did John do today?
b. Mary: It was unreal! I mean, he was more eloquent than himself.

Given a modulated LF roughly analogous to (37b), we predict that Mary's assertion is strictly acceptable and can convey something like that John's degree of eloquence (on that odd Monday) was higher than the degree of eloquence that he usually or normally displays.

4.3 Too much modulation?

Logicality + Modulated LFs says that the modulation operator, $\mathcal{R}$, appears as a sister of all content terms and variables. The account of superficial trivialities with bound variables in §4.1-4.2 builds on that assumption. One might worry, however, that while that assumption helps with the over-generation problem, it gives too much expressive power to $\mathcal{R}$, thereby forcing 'informative' readings for superficial tautologies and contradictions. Yet there are contexts in which the intended readings are precisely the trivial ones. Consider example (39), where the context as updated by the first assertion suggests that the speaker intends that the complement of the belief attribution should be assigned its trivial, contradictory reading.

(39) John is totally irrational. He believes that he will both win and not win the race.

One influential approach to these cases, going back to Quine (1956), Kaplan (1968) and Cresswell & Von Stechow (1982), appeals to concepts through which the relevant individual is accessed by the attitude holder, where a belief is de re about an individual u whenever u reliably induces a concept in the belief holder a which identifies u for a in a's belief state. For (ia), such concept might be 'the man who wants to share John's inheritance'. Charlow & Sharvit (2014) propose an implementation of this approach in which the LFs for de re beliefs include 'concept generators', which are inserted in the syntactic spot of the *res* and drive pragmatically the propositional content of the belief. According to Chierchia (2019), the use of modulation over individual terms and variables can arguably be viewed as an extension of Charlow and Sharvit's proposal for the semantics de re belief.

Suppose the embedded clause has a modulated LF as in (40a). This seems to predict that the embedded clause gets the reading in (40b), but in (39) the default reading is closer to (40c).

(40) a. he_1 will $\mathcal{R}_{c'}(\text{win}) \wedge \text{he}_1$ will not $\mathcal{R}_{c''}(\text{win})$
b. John believes that he will win (in one sense of winning) and also that he won't win (in another sense of winning).
c. John believes, in exactly the same sense of winning, that he will win and not win.

Yet Logicality + Modulated LFs, as presented in §3.1, entails that superficially trivial expressions can be assigned trivial readings. On this view, an expressions counts as L-trivial, and is thus filtered out, only if it is trivial on every possible modulation (i.e., resolution of $\mathcal{R}$), which is obviously not the case for (40a). Still, even in such cases, $\mathcal{R}$ can ultimately (i.e., once the context is taken into account) be assigned the laziest modulation, i.e., the identity function. In the case of (39), this choice would generate the intended reading. From this perspective, the second sentence in (39) is trivial but not L-trivial, and is thus correctly predicted to be strictly acceptable. Generalizing, Logicality + Modulated LFs entails that some (acceptable) expressions which are not L-trivial—since they are not trivial on every possible modulation of each token of $\mathcal{R}$—can still be assigned a trivial reading in particular contexts. Indeed, Logicality + Modulated LFs is compatible with the view that lazy modulation is the default, such that $\mathcal{R}$ is only assigned substantial modulation operations when supported by specific patterns of focus/intonation, questions under discussion, and similar factors.

5 Other approaches to Logicality compatible with Semantic Minimalism

The Contextualist package of Logicality + Modulated LFs, I have argued, is descriptively superior to the Semantic Minimalist-friendly package of Logicality + Skeletons. Specifically, Logicality + Modulated LFs issues in a more general solution to the over-generation problem while preserving L-triviality-based accounts of acceptability patterns, such as those in (1)-(3), which help capture the distribution of various functional terms and phrases. For those sympathetic to Logicality, this result amounts to a novel argument for Contextualism over Semantic Minimalism—but only if there are no other viable implementations of Logicality compatible with Minimalism. In this section, I present three additional Minimalist-friendly implementations of Logicality, and argue that each option is descriptively inferior, given the over-generation problem, to Logicality + Modulated LFs. Unlike Skeletons, these proposals have not been explored in the literature; yet each has some prima facie plausibility. Examining why they fail will deepen our understanding of the conditions that should be satisfied by any viable implementation of Logicality.

5.1 L-triviality within Phases

Suppose that the DS sees 'standard' (Semantic Minimalist-friendly) logical forms—i.e., textbook syntactic representations, different from both logical skeletons and modulated logical forms, where only a special class of terms exhibits linguistically-driven context sensitivity. Assume, however, that the DS only checks for trivialities within (and not across) 'minimal syntactic phases'. As a first pass, we can say that a syntactic structure counts as a minimal phase if it can be assigned a propositional type interpretation and has no proper constituents that can also be assigned a propositional type interpretation.

(41) **Logicality + Phases**. The DS sees standard logical forms and filters out all expressions which can be shown to be logically trivial. However, the DS operates only within minimal syntactic phases. Expressions whose triviality depends on comparing information across minimal phases are not seen as L-trivial by the DS, hence are not marked as strictly unacceptable.

The hypothesis that syntactic structures are computed in phases has some independent motivation (Chomsky 1995, Radford 2004). To see why Logicality + Phases has some promise as a solution to the over-generation problem, consider again two basic examples of the kinds of superficial trivialities that implementations of Logicality should *not* classify as L-trivial:

(42) a. If John_1 is wrong, then he_1 is wrong.
b. It is raining and it is not raining.

(42a) and (42b) share the feature that, to identify their triviality, the DS would have to look across more than one minimal propositional structure, i.e., it would have to compare material across distinct syntactic phases. Specifically, to determine if (42a) is a tautology, the DS would have to compare information across two phases, as informally captured in (43a). Similarly, to determine if (42b) is a contradiction, it would need to look across two phases, as informally captured in (43b):

(43) a. [If $\overbrace{[\text{ j}_1 \text{ is W}]}^{\text{Min. Phase}}$, then $\overbrace{[\text{he}_1 \text{ is W}]}^{\text{Min. Phase}}$]
b. [$\underbrace{[\text{It is R-ing}]}_{\text{Min. Phase}}$ and $\underbrace{[\text{it is (not) R-ing}]}_{\text{Min. Phase}}$]

Suppose that the DS only checks for trivialities within minimal syntactic phases. It follows that superficial trivialities like (42a)-(42b) will not be identified and filtered-out by the DS. This holds even if (within each minimal phase) the DS sees otherwise standard logical forms, as in (43a)-(43b). In addition, many of the cases of trivialities that do result in 'ungrammaticality' can be proven from minimal propositional clauses. For example, it is easy to check, for our account of exceptive-*but* phrases in §2, that proving the target cases of L-triviality at no point depends on comparing material across minimal phases.

Despite its advantages, Logicality + Phases both over and under-generates unacceptability assignments. Starting with over-generation, consider again acceptable superficial trivialities with reflexives, such as (44a) and (45a), which as we saw undermine Logicality + Skeletons but not Logicality + Modulated LFs. In these cases, each contradiction or tautology can be proven within a minimal phase (e.g., no connectives or proposition taking operators are essentially involved, and the target reflexives must be bound in their local syntactic environment), as can be seen from their partial LFs in (44b) and (45b):

(44) a. John is (not) himself.
b. John $\underbrace{[t_i \text{ is (not) himself}_i]}_{\text{Min. Phase}}$

(45) a. John is more eloquent than himself.

b. John $\underbrace{[\mathrm{t}_i \text{ is more eloquent than himself}_i]}_{\text{Min. Phase}}$

It follows that many simple superficial trivialities with reflexives would, on this proposal, come out as L-trivial, and so would be incorrectly predicted to feel ungrammatical.

Logicality + Phases also under-generates assignments of unacceptability. The Logicality program includes triviality-based accounts of the distribution of propositional operators, such as attitude verbs. To derive the target trivialities in these cases, the DS would need access to the interaction between propositional operators and the content of their complements. As examples, consider the two patterns in (46) and (47), both of which have a triviality-based explanation. The key point is easy to see (even without getting into details): the relevant trivialities in (46a) and (47a) can only be derived if we can compare material within a minimal propositional phase (the embedded clauses) with operators outside of it (the attitude verbs).[12] Accordingly, if the DS could only prove trivialities within minimal phases, it would *not* filter out as L-trivial (unacceptable) expression such as (46a) and (47a).[13]

(46) Attitude verbs and interrogative embedding: (Mayr 2019)

a. *John believes whether Mary smokes.

b. John knows whether Mary smokes.

(47) Weak presuppositional islands: (Abrusán 2011)

a. *How do you regret that Mary fixed the roof?

b. How do you hope that Mary fixed the roof?

The problems of over and under-generation of unacceptability assignments, taken together, amount to a serious dilemma for Logicality + Phases—and it is hard to imagine a reasonable modification of the notion of minimal phases that can avoid it. On the one hand, to block the incorrect assignment of L-triviality for simple sentences with reflexives such as (44a) and (44b), minimal phases would have to involve 'small' syntactic structures with arguably sub-propositional type interpretations. On the other hand, to prove L-triviality for cases that require access to the interaction between propositional operators and their complements, such as (46a) and (47a), minimal phases would have to involve rather inclusive syntactic structures which may have structures with propositional type interpretations as proper sub-constituents. It is hard to see how a coherent and independently motivated notion of phases might satisfy both of these constraints, since they pull in opposite directions with respect to the size-complexity of the kinds of structures that are evaluated for triviality by the DS.

5.2 Exotic Deductive Systems

Another way of pairing Semantic Minimalism with Logicality is to assume that the DS implements a non-classical 'natural' logic. Many acceptable superficial trivialities, given

12 For demonstration that those accounts of the distribution of attitude verbs can be implemented in Logicality + Modulated LFs, see Del Pinal (2019).

13 Another prominent triviality-based account which also depends on the interaction between propositional operators and their complements is Chierchia's account of the distribution of negative polarity items, already mentioned in §1 and discussed in more detail in §5.2 below.

their standard logical forms, correspond to simple cases of classically trivial formulas: e.g., violations of the law of non-contradiction. By adopting a non-standard logic for the DS—e.g., a relevant logic which allows for $p \wedge \neg p$ to be contingent (i.e., to have some true and some false instantiations)—we can restrict the over-generation of unacceptability for superficial trivialities. And we can do this without holding that skeletons are a level of syntactic representation—for the non-classical DS can run directly on standard LFs.

(48) **Logicality + Exotic DS**. The DS interfaces with standard (Semantic Minimalist-friendly) logical forms. However, the kind of 'natural logic' implemented by the DS is closer to relevant logics (or to even more exotic systems) than to classical logics. All expressions which are trivial relative to the exotic DS are classified as L-trivial, and hence filtered out as strictly unacceptable.

What is the advantage of modeling the DS as a relevant logic (or an even weaker system)? The proposal to run the DS on skeletons in which each content term token is replaced with a new variable of the appropriate type basically mimics some results of such non-classical logics: e.g., *it is raining and not raining* comes out as contingent because it is 'seen' by the DS as 'it is P and it is not Q'. The main objection we raised in §4 against Skeletons concerns cases in which, due to binding or some syntactic/semantic identity constraint on ellipsis, the identity of content term tokens is explicitly encoded by the Grammar. By *directly* modeling the DS as a relevant logic, one may avoid that objection. For on this implementation, formulas like 'it is P and not P'—*seen as such by the DS*—come out as strictly contingent, hence are not filtered out by the DS, even if the tokens of 'P' are co-bound in any way.

Unfortunately, Logicality + Exotic DS faces a dilemma. To fully solve the over-generation problem, not only superficial contradictions, but even tautologies like *if it is raining, then it is raining* would have to come out as contingent. Yet 'if P then P' and various other basic examples of superficial tautologies are valid in relevant logics; hence would come out as L-trivial if the DS is modeled as a relevant logic which runs on standard LFs. The problem, of course, is that such superficial tautologies are in general as acceptable as superficial contradictions. This suggests that to fully deal with the over-generation problem, this direct approach would have to adopt an extremely weak logic for the DS.[14] The problem with adopting a weak logic for the DS, however, is that many of the triviality-based accounts which make Logicality such a powerful hypothesis depend on the validity of various classical formulas and inference rules, such as the LNC, MP and MT.

To illustrate this, consider a simplified version of Chierchia's (2013) triviality-based account of the distribution of negative polarity items (NPIs), focusing on the case of *any*. The basic contrast, presented in (3), is repeated below in (50) and (51). Chierchia argues that *any* is an indefinite with existential force which, unlike its plain counterpart *a/an*, triggers obligatory exhaustification of domain alternatives, O_{DA}, defined as in (49):

14 Indeed, proponents of this package would arguably be forced to hold that the DS is as weak as, say, Korner's (1955) logic for vagueness/inexact concepts (cf. Gajewski 2009). As discussed in Williamson (1994), this system provides truth-tables for the connectives that basically treat each token of a propositional variable as independent (even in formulas like $p \wedge p$), so that the resulting 'logic' is extremely weak. In itself, this might not be a totally unattractive position for a Semantic Minimalist who holds that, while most open-class terms aren't context-sensitive, all or most of them are vague.

(49) a. $[\![O_{DA}\ \phi]\!]^{g,w} = [\![\phi]\!]^{g,w} \wedge \overbrace{\forall p \in [\![\phi]\!]^{DA}[p \rightarrow \lambda w'[\![\phi]\!]^{g,w'} \subseteq p]}^{\text{negate d-alternatives not entailed by } \phi}$

b. $[\![\phi]\!]^{DA} = \{[\![\phi]\!] : D' \subseteq g(D)\}$

In (50), *any* occurs in a downward-entailing environment. Suppose for simplicity that the relevant domain is John's house, which has just a living room and a kitchen. Since the prejacent of O_{DA} entails each of its domain alternatives in (50b), exhaustification is in this case vacuous, as captured in (50c). The result is obviously a contingent statement which can be true or false depending on whether John has any eggs in the world of evaluation.

(50) John doesn't have any eggs.

a. O_{DA} [¬John has an egg $\in D_{house}$]

b. $DA = \{$¬John has an egg $\in D_{house}$, ¬John has an egg $\in D_{kitchen}$, ¬John has an egg $\in D_{living\_room}\}$

c. $[\![$(50a)$]\!]$ = ¬John has an egg $\in D_{house}$

In contrast, in (51) *any* occurs in an upward-entailing environment. As a result, this account now generates trivial truth-conditions. To see why, notice that the prejacent of O_{DA}—namely, that John has an egg in the house—entails neither its alternative that John has an egg in the kitchen, nor its alternative that John has an egg in the living room. Given the definition of O_{DA} in (49), this means that each of these alternatives has to be negated, as captured in (51c). This generates a contradiction, since by assumption the domain of John's house consists just of the subdomains of the kitchen and living room.

(51) *John has any eggs.

a. O_{DA} [John has an egg $\in D_{house}$]

b. $DA = \{$John has an egg $\in D_{house}$, John has an egg $\in D_{kitchen}$, John has an egg $\in D_{living\_room}\}$

c. $[\![$(51a)$]\!]$ = John has an egg $\in D_{house} \wedge$ ¬John has an egg $\in D_{kitchen}$ $\wedge$ ¬John has an egg $\in D_{living\_room}$

What is crucial to note, for us, is that this account depends on the assumption that the DS can identify and filter out violations of the LNC, such as (51c). Yet this is precisely what we would have to reject if we assume that the DS directly implements a non-classical logic in which the LNC is not valid. Like Chierchia's account of NPIs, many other triviality-based accounts that make up the Logicality program depend on the stipulation that the DS is a rather powerful inferential system.

At this point, it is important to understand why, given Chierchia's account of NPIs, Logicality + Modulated LFs filters out expressions like (51) but not superficial contradictions. Consider the modulated LF of (51) in (52a). The modulation function $\mathcal{R}$ can apply to any open-class term in the prejacent of O_{DA}. Once any modulations are inserted into the LF for the prejacent, the formal alternatives are determined from the subdomains of D_{house}, as illustrated in (52b).

(52) a. O_{DA} [John has an $\mathcal{R}_{c'}$(egg) $\in D_{house}$]

b. $DA = \{$John has an $\mathcal{R}_{c'}$(egg) $\in D'$: $D' \subseteq D_{house}\}$

Suppose that *egg* is modulated to 'expensive egg'; we would still derive a contradiction when we exhaustify as in (52a) over the domain alternatives in (52b). The key assumption here is that the interpretation of non-focused terms, even if they are context-sensitive characters, remains constant across formal alternatives. This assumption is independently justified. To see why, consider the exhaustified (scalar) reading of (53), focusing on the behavior of *tall*, a paradigmatic context-sensitive term. Although its context-sensitive parameters can be saturated in different ways in its local context—to capture different thresholds for counting as 'tall'—that interpretation has to be held constant across the formal alternatives (in this case scalar alternatives = SA) used by the exhaustification operator, as captured in (53b).

(53) Some$_{\mathrm{F}}$ students are tall.

a. O_{SA}(Some$_{\mathrm{F}}$ students are tall$_{c'}$)

b. $SA = \{$Some students are tall$_{c'}$, All students are tall$_{c'}\}$

c. Some students are tall$_{c'}$ $\wedge$ $\neg$All students are tall$_{c'}$

d. Some students are tall$_{c'}$ $\wedge$ $\neg$All students are tall$_{c''}$

This explains why (53) can have the enriched reading in (53c) but not the one in (53d), i.e., why (53) cannot be enriched to mean something like 'some students are tall given threshold A, but not all students are tall given higher threshold B'. In contrast, it is clearly possible to switch standards when *tall* occurs in two different local contexts at LF, such as in (54):

(54) My students are tall for US standards, but they aren't tall for Dutch standards.

In short, the principles which guarantee that paradigmatic context-sensitive terms like *tall* are assigned uniform interpretations across formal (scalar) alternatives in structures like (53a), but not in (54), also guarantee that $\mathcal{R}$, which is also a context-sensitive operator, must be assigned a uniform interpretation across domain alternatives in examples like (52), but not when it occurs in different sites at LF such as in typical superficial trivialities.

5.3 Anti-triviality clauses

The third attempt to square Semantic Minimalism with Logicality—to tackle the over-generation problem—is based on a technical trick. As pointed out by Chierchia (2013), we can eliminate, from our theory of the language system, the notion of a DS or natural logic that identifies and filters out L-triviality by introducing specific anti-triviality clauses into the semantic entries for certain functional terms. Using this technique, we can try to reduce L-triviality to presupposition failure.

(55) **Logicality as anti-triviality presuppositions.** The language system doesn't include a DS that identifies and filters out L-trivial expressions. Instead, many functional/logical terms include, as part of their meaning, anti-triviality presuppositions. The class of L-trivial expressions can be reduced to that of expressions which violate such anti-triviality clauses.

Schematic examples of lexical entries with anti-triviality presuppositions for (domain alternatives-based) exhaustification and exceptive-*but* are presented in (56b) and (57b). Given (56b), trivial sentences with NPIs like (56a) come out as presupposition failures; and given (57b), trivial sentences with exceptive phrases like (57a) also come out as presupposition failures.

(56) a. *Sam has any philosophy books

b. $O^*_{DA}(\phi) = \begin{cases} \#\text{ if } O_{DA}(\phi) \text{ is trivial;} \\ O_{DA}(\phi) \text{ otherwise} \end{cases}$

(57) a. *[[Three$_{\mathbf{D}}$ [athletes$_A$ [but John$_C$]]] smoke$_P$]

b. $\textsc{but}^*(\text{C})(\text{A})(\mathbf{D})(\text{P}) = \begin{cases} \#\text{ if } \textsc{but}(C)(A)(\mathbf{D})(P) \text{ is trivial;} \\ \textsc{but}(C)(A)(\mathbf{D})(P) \text{ otherwise} \end{cases}$

This strategy can be generalized: i.e., we can re-write the semantic entries for certain functional terms so that what we originally classified as L-triviality-based cases of unacceptability result instead from violations of explicit anti-triviality clauses. Since the trivialities that result in unacceptability are encoded in specific lexical entries, we avoid the over-generation problem, at least in its original form. As a result, this version of Semantic Minimalism need not appeal to logical skeletons, and is thus not directly undermined by the problems raised against Logicality + Skeletons in §4.

This use of anti-triviality clauses in the entries for functional terms, however, faces serious obstacles. According to Logicality, there is a subset of the trivial sentences, the 'L-trivial' ones, which are unacceptable. According to Logicality + Modulated LFs, we can derive the empirically correct set of L-trivial sentences—hence address the over-generation problem—on the basis of independently justified assumptions about functional terms and the kind of context-sensitivity characteristic of the content-based lexicon. This suggests a rationale for why L-trivial sentences are unacceptable, while merely trivial ones are strictly acceptable: merely trivial sentences can convey (useful) information, depending on the selected modulations, whereas L-trivial ones are not even potentially useful, i.e., they are unrecoverable under all possible modulations. Contrast that picture with the one suggested by the anti-triviality account. The problem is not just that it seems pointless to write a specific anti-triviality presupposition clause into the semantic entry of each functional/logical term involved in triviality-driven acceptability patterns. It is rather that this account doesn't come with an independent rationale for deciding when to include such anti-triviality clauses. As a result, we end up with an ad hoc procedure that faces its own version of the over-generation problem.

For if natural languages can encode anti-triviality clauses, why don't they do so for all functional/logical terms? For example, why don't the entries for *and* and *or* include anti-triviality clauses that filter out trivial conjunctions and disjunctions? Obviously, these entries would over-generate unacceptability for many superficial tautologies and contradictions, given standard logical forms without modulation operators (i.e., given Minimalist-friendly logical forms). For example, given anti-triviality conjunction, $\textsc{and}^*$, defined as in (58), a superficial contradiction like (59) would be incorrectly predicted to be unacceptable, given its standard LF in (59a). This prediction is blocked by adopting the modulated LF in (59b), but this option is not in general available to theorists opting for a Semantic Minimalist-friendly implementation of anti-triviality clauses.

(58) $\textsc{and}^*(\text{p})(\text{q}) = \begin{cases} \#\text{ if } p \wedge q \text{ is trivial;} \\ p \wedge q \text{ otherwise} \end{cases}$

(59) It is raining and it is not raining

a. Standard LF: [It is P [and not it is P]]

b. Modulated LF: [It is $\mathcal{R}_{c'}$(P) [and not it is $\mathcal{R}_{c''}$(P)]]

In short, proponents of this view need to explain why only some functional terms encode anti-triviality clauses. The rationale cannot be that, relative to their standard logical forms, such clauses filter out logically trivial and hence informationally useless expressions, for this wouldn't explain why connectives like *and* and *or* don't also incorporate anti-triviality clauses. In addition, they would also have to specify which kinds of presupposition failures generate judgements of strict unacceptability. According to most extant theories, the observational signature of presupposition failures is something like 'intuitive' oddness (cf. *It is raining, but John knows it isn't raining*), or uncertainty concerning truth-value assignments given all the relevant facts and controlling for vagueness (cf. *The current King of France is bald*). These observational signatures should be distinguished from strict unacceptability, which is closer to the feeling of ungrammaticality. Accordingly, and as pointed out in Chierchia (2013), proponents of Logicality as anti-triviality would have to explain why some but not all presupposition failures give rise to judgements of strict unacceptability. The challenge can be seen more directly in (60a)-(60c). All these expressions involve, given the anti-triviality view, some kind of presupposition failure, but only (60a) feels strictly unacceptable:

(60) a. *Sue broke any cups.

b. ?I met an Italian that turned out not to be Italian.

c. ?Mary knows a pilot who is not a pilot.

The project of specifying which subset of presupposition failures gives rise to strict unacceptability is as hard as that of specifying which subset of trivial sentences counts as L-trivial, i.e., gives rise to strict unacceptability. The problem, of course, is that the anti-triviality proposal was presented, at this point in our dialectic, as a general solution to the latter project.

6 Logical vs. non-logical words and the domain of modulation

In §4-§5, I argued that the Contextualist-friendly package of Logicality + Modulated LFs constitutes a more satisfactory approach to the over-generation problem than various implementations of Logicality which are compatible with Semantic Minimalism. To conclude my argument, I want to clarify and justify a key assumption of my approach. According to Logicality + Modulated LFs, the modulation operator $\mathcal{R}$ is inserted as a sister of all non-logical terminal nodes. Although there is an intuitive difference between logical terms like determiners, connectives, and modals, and content terms like nouns, adjectives and verbs which pick out entities, events, or functions of entities or events, this approach ultimately depends on the availability of a more systematic procedure for separating the logical and non-logical terms. Indeed, this also applies to other implementations of Logicality: e.g., logical skeletons can only be derived from standard logical forms if there is a way of identifying their non-logical points. The goal of this section is to explain why I am optimistic that we will be able to find a computationally tractable procedure for separating the fixed, logical terms of natural languages from the non-logical terms that are open to modulation. My approach builds on previous work on the identification of logical constants, esp., on related observations by Chierchia (2019).

Most of the lexical terms of natural languages that are commonly classified as paradigmatically logical share a cluster of syntactic and semantic properties (von Fintel 1995, Gajewski 2009, MacFarlane 2017, Chierchia 2019). Syntactically, logical terms tend to fall on the functional, closed-class side of the lexicon, while content terms—i.e., referential or world-directed terms—fall on the open-class side of the lexicon. In current generative approaches, functional terms appear on the edges of noun and verb phrases, forming the 'extended projections' of the latter, content-based phrases. Semantically, paradigmatic logical terms share two features that are important for our purposes. First, they pass a range of invariance tests. There are various kinds of invariance tests, some more strict than others (see e.g., van Benthem 1989, McGee 1996, Sher 2003, Sagi 2014, MacFarlane 2017). For the purposes of implementing Logicality, we should use relatively inclusive invariance tests, such as tests that involve permutations of the domain of individuals and events which respect to structural differences across domains such as the mass/count and the event/state distinctions. Second, paradigmatic logical terms tend to be assigned high types. The sorts of terms that pass such inclusive permutation invariance tests and are assigned high types includes determiners (*every*, *none*, *most*), connectives/coordinators (*and*, *or*), modals (*must, might*), exceptives (*but, except*) and exhaustifiers (*even, only, O*)—i.e., all the terms that we have thus far treated as part of the fixed natural logic used by the language system (see Gajewski 2009, Sagi 2014, MacFarlane 2017, Chierchia 2019). In contrast, content terms—incl., individual and predicate variables—typically fail such permutation invariance tests, and are usually assign a 'low' semantic type, corresponding to their role of standing for individuals, events, or predicates of individuals or events.

Although there is a significant overlap between the functional, closed-class, permutation invariant, and high-typed terms, on the one hand, and the content, open-class, non-permutation invariant, and low-typed terms, on the other, there are some mismatches predicted by the different criteria within each of these clusters. How we propose to resolve these mismatches matters to the (empirical) project of picking out the appropriate set of L-trivial expressions. Consider two examples. First, predicates like *exists* come out as logical when classified using certain permutation invariance tests (Gajewski 2009, MacFarlane 2017), but as non-logical when classified using its type, namely, that of a one-place predicate akin to *made of plastic*. If we hold that any terms which pass such permutation invariance tests are treated as logical constants by the language system, hence not in the domain of $\mathcal{R}$ (i.e., not subject to modulation), then sentences like *Pete exists* would come out as L-trivial and incorrectly predicted to feel strictly unacceptable. Second, pronouns—including reflexives—are arguably part of the functional, closed-class vocabulary, and yet are not permutation invariant and their semantic type is, on most accounts, simply that of (variables of) entities (or individual level concepts), or of pluralities of entities. If we hold that any terms which are part of the closed-class vocabulary are treated as logical constants by the language system, they would not be in the domain of $\mathcal{R}$. As a result, superficial trivialities with reflexives such as *John is not himself today* would come out as L-trivial and incorrectly predicted to be unacceptable.

When considering such mismatches across different criteria for separating the logical from the non-logical terms, it is important to appreciate that, given the project of implementing Logicality, our goal is not to select a procedure that picks out the 'true' logical constants. Our goal is the empirical and pragmatic one of selecting a procedure that, when combined with our implementation of Logicality, results in an overall theory that determines the correct set of L-trivial expressions, i.e., that assigns triviality just to those expressions that, while syntactically well-formed, are judged by competent speakers to feel strictly unacceptable. At the same time, it is not appropriate, given that goal, to simply point out that we should use these criteria as reliable diagnostics—and not as necessary/sufficient conditions—for picking out the language system relative logical terms. For any term (or class of terms) that is cross-classified by the criteria within a cluster (e.g., a term that is classified as logical based on an invariance test but as non-logical based on its low semantic type), we still have to decide whether it is in the domain of modulation. And this choice will partly determine whether we derive the correct acceptability patterns for expressions containing that term.

Which criteria, then, should get the highest weight for picking out the language system relative non-logical terms? I propose that the domain of the modulation operator $\mathcal{R}$ should be determined by the semantic types of its possible arguments. Specifically, $\mathcal{R}$ should be treated as a constrained polymorphic type operator, which can take as arguments any terms which have a 'referential' type, relative to the target theory. Given a semantic theory in which the basic domains (excluding the truth values) are those of entities and events, $\mathcal{R}$ will apply to terms and variables of type e, v, and any terms and variables for functions of a type whose first element is of type e or v ($< e, t >$, $< e < e, t >>$, etc.). This proposal excludes any high typed functions from the domain of modulation—i.e., any functions whose first argument is not a (non-truth value) basic type. It follows that determiners, connectives/coordinators, modal auxiliaries, exceptives and exhaustifiers are not in the domain of modulation, the desired result. In addition, this proposal deals nicely with the previous examples of cross-classified terms. First, predicates that apply to the entire domain of entities in all models will be treated by the language system as content terms and subject to modulation—even if, on some tests, they count as permutation invariant (same holds of predicates that are empty in all models). This result might seem problematic for certain projects in philosophical logic, but it helps pick out the correct set of L-trivial expressions. For then sentences like *Pete exists* do not come out as L-trivial, and are thus correctly predicted to be strictly acceptable (*exists* can be modulated to mean something like 'exists relative to some relevant world which need not be the actual one'). Second, on this view reflexives—taken as bound (individual) variables (i.e., of type $< e >$ or $< s, e >$)—are also in the domain of $\mathcal{R}$, even if they are part of the closed-class lexicon. As shown in §4.2, this entails that superficial trivialities with reflexives such as *John is not himself today* do not come out as L-trivial and are correctly predicted to be strictly acceptable.

To be sure, this (preliminary) proposal for specifying the domain of modulation leaves open various important issues. For example, future work should examine acceptability patterns involving mixed or semi logical terms such as prepositions and propositional attitude verbs to determine if those terms are treated by the language system as part of the fixed, logical vocabulary or as part of the non-logical terms that are subject to modulation. [15] Those results will help inform whether or not expressions of the corresponding semantic types in general should be included in the domain of modulation. In addition, relative to semantic theories with a strict correspondence between syntactic categories and semantic types, this

15 For an attempt to reconcile Logicality + Modulated LFs with the view, advocated by Abrusán (2014) and Mayr (2019), that propositional attitude verbs can trigger systematic patterns of L-triviality, as illustrated in (46)-(47), see Del Pinal (2019). Briefly, I argue there that although attitude verbs are subject to modulation, the presuppositions of attitude verbs project from such modifications in the usual way. As a result, the presupposed factivity (or lack thereof) of the attitude verb is preserved across all possible modulations, and this is enough to maintain the triviality-based accounts of Abrusán (2014) and Mayr (2019) of patterns like (46)-(47).

kind of proposal is relatively deterministic and entails that $\mathcal{R}$ will range over nouns, pronouns, verbs, adjectives and adverbs. Yet relative to theories that allow for substantial semantic type variation within each syntactic category, this proposal leaves open various parameters which may be used to explore different ways of fixing the (disputed) boundaries of the domain of $\mathcal{R}$. For those interested in constructing an empirically adequate implementation of Logicality, this proposal can in turn push assumptions—perhaps even revisionary ones—about which semantic types to assign to specific classes of terms.

7 Conclusion

The project of finding an implementation of Logicality that can preserve triviality-based accounts of the distribution of quantifiers, modals, and exhaustifiers, among other logical or semi-logical terms and phrases, without over-generating unacceptability assignments for 'superficial' trivialities opens up a novel way of tackling traditional philosophical disputes about the nature of logical form, including ongoing debates between Contextualists and Semantic Minimalists. This paper explored various implementations of Logicality compatible with these philosophical frameworks. I have argued that each Minimalist-friendly implementation is descriptively inadequate as a general solution to the over-generation problem, while pairing Logicality with a version of Contextualism results in a more promising approach. I also argued that not just any version of Contextualism will work as part of this package: Logicality cannot be paired with radical accounts according to which all terms—including logical terms—can be modulated. Finally, the discussion of various novel Minimalist-friendly proposals revealed some general constraints on any defensible implementation of Logicality: (i) the natural logic used by the language system seems to be quite powerful, and should respect most classical rules of inference,[16] and (ii) triviality-induced unacceptability cannot in general be reduced to violations of explicit and lexically encoded anti-triviality presuppositions.

Semantic Minimalists (and Radical Contextualists) might be tempted to resist these results by rejecting the Logicality of language hypothesis. Although the main goal of this paper is not to directly defend Logicality, I think that the case studies discussed here illustrate the considerable power and elegance of triviality-based explanations of the distribution of functional terms and phrases. It is becoming increasingly clear that rejecting Logicality is a costly move. Any version of Semantic Minimalism or Contextualism—indeed, any hypothesis about the nature of logical form—that depends on that move would have reduced credibility as an empirical hypothesis about a level of representation used by the language system and its interfaces. For this reason, I hope that even philosophers who ultimately reject the specific claims I defend here will be convinced that it is useful to frame traditional debates between Semantic Minimalists and Contextualists as debates that are in part about how to implement Logicality and understand why some syntactically well-formed sentences are automatically filtered out by the language system.

Logicality also interacts in interesting ways with other ongoing debates in Philosophy of Language. First, we have seen that most viable implementations of Logicality (whether via Skeletons or Modulated logical forms) depend on a distinction between functional/logical terms and content/non-logical terms. Although coming up with a principled distinction between

16 Recall, however, that the relevant notion of entailment is close to Strawson-entailment. This is because a (modulated) LF is L-trivial if, *whenever defined*, it is either trivially true or false (or equivalently, is entailed by the empty set whenever it is defined) for all possible modulations (see §3.1).

logical and non-logical terms is difficult, as is well known from related work in philosophical logic (van Benthem 1989, Sher 2003, Sagi 2014, MacFarlane 2015), I have argued, following Chierchia (2019), that there are good reasons to think that such a distinction plays a central role in the architecture of the language system. Still, much works remains to be done to really solidify that hypothesis (see §6). Secondly, some Logicality-style accounts assume that the DS has access to information that goes beyond strictly 'logical' information. For example, accounts of modified numerals (Fox & Hackl 2007), negative islands in comparatives (Gajewski 2008b) and weak presuppositional islands (Abrusán 2014), depend on substantial structural assumptions about the domains of numbers, degrees and manners. In other words, they require (domain-specific) stipulations about natural language metaphysics. A philosophically satisfying implementation of Logicality will have to grapple with these foundational issues at the interface of language, logic and metaphysics.

References

Abrusán, Márta. 2011. Presuppositional and negative islands: A semantic account. *Natural Language Semantics* 19(3). 257–321.

Abrusán, Márta. 2014. *Weak island semantics*. Oxford, UK: Oxford University Press.

Abrusán, Márta, N. Asher & T. Van de Cruys. 2018. Content vs. function words: The view from distributional semantics. In Uli Sauerland & Stephanie Solt (eds.), *Proceedings of* Sinn und Beudeutung *22*, .

Abrusán, Márta, Nicholas Asher & Tim Van de Cruys. 2019. Grammaticality and meaning shift. In Uli Sauerland & Stephanie Solt (eds.), *The semantic conception of logic*, chap. forthcoming. Cambridge, UK: Cambridge University Press.

Barwise, Jon & Robin Cooper. 1981. Generalized quantifiers and natural language. *Linguistics & Philosophy* 4. 159–219.

Borg, Emma. 2004. *Minimal semantics*. Oxford: Oxford University Press.

Cappelen, H. & E. Lepore. 2005. *Insensitive semantics*. Oxford, UK: Basic Blackwell.

Carston, Robyn. 2002. *Thoughts and utterances: The pragmatics of explicit communication*. Malden, MA: Blackwell Publishers Ltd.

Charlow, Simon & Yael Sharvit. 2014. Bound 'de re' pronouns and the lfs of attitude reports. *Semantics and Pragmatics* 7. 1–43. http://dx.doi.org/http://dx.doi.org/10.3765/sp.7.3.

Chierchia, Gennaro. 2006. Broaden your views: Implicatures of domain widening and the "logicality" of language. *Linguistic inquiry* 37(4). 535–590.

Chierchia, Gennaro. 2013. *Logic in grammar: Polarity, free choice, and intervention*, vol. 2. Oxford: Oxford University Press.

Chierchia, Gennaro. 2019. On being trivial: Logic vs grammar. In Gil Sagi & Jack Woods (eds.), *The semantic conception of logic: Essays on consequence, invariance, and meaning*, Cambridge University Press.

Chomsky, Noam. 1981. *Lectures on government and bindinding*. Berlin: Mouton de Gruyter.

Chomsky, Noam. 1995. *The minimalist program*. Cambridge, MA: The MIT Press.

Cresswell, Maxwell J & Arnim Von Stechow. 1982. De re belief generalized. *Linguistics and Philosophy* 5(4). 503–535.

Crnič, Luka. 2018. A note on connected exceptives and approximatives. *Journal of Semantics* 35(4). 741–756.

Del Pinal, Guillermo. 2019. The logicality of language: A new take on triviality, 'ungrammaticality', and logical form. *Nous* 53(4). 785–818. http://dx.doi.org/10.1111/nous.12235.
Fox, Danny. 2000. *Economy and semantic interpretation*. Cambridge, MA: The MIT Press.
Fox, Danny & Martin Hackl. 2007. The universal density of measurement. *Linguistics & Philosophy* 29(5). 537–586. http://dx.doi.org/10.1007/s10988-006-9004-4.
Gajewski, Jon. 2002. L-analyticity and natural language. Manuscript.
Gajewski, Jon. 2008a. Connected exceptives and NPI *any*. *Natural Language Semantics* 16(1). 69–110.
Gajewski, Jon. 2008b. More on quantifiers in comparative clauses. In T. Friedman & S. Ito (eds.), *Semantics and linguistic theory XVIII*, vol. 18, 340–357. Ithaca, NY.
Gajewski, Jon. 2009. L-triviality and grammar. Manuscript.
Harley, Heidi. 2014. On the identity of roots. *Theoretical Linguistics* 40(3-4). 225–275. http://dx.doi.org/https://doi.org/10.1515/tl-2014-0010.
Heim, Irene. 1991. Artikel und defintheit. In Arnim von Stechow & D. Wunderlich (eds.), *Semantik: ein internationales handbuch der zeitgenössischen forschung*, 487–535. De Gruyter.
Heim, Irene & Angelina Kratzer. 1998. *Semantics in Generative Grammar*. Oxford, UK: Blackwell Publishers Ltd.
Hirsch, Aron. 2016. An unexceptional semantics for expressions of exception. *University of Pennsylvania Working Papers in Linguistics* 22(1). 16.
Kaplan, David. 1968. Quantifying in. *Synthese* 19(1-2). 178–214.
Kennedy, Christopher. 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics & Philosophy* 30(1). 1–45.
Körner, S. 1955. *conceptual thinking*. Cambridge, UK: Cambridge University Press.
MacFarlane, John. 2015. *Logical constants*. Stanford Encyclopedia of Philosophy. https://plato.stanford.edu/archives/win2017/entries/logical-constants/.
MacFarlane, John. 2017. Logical constants. In E. N. Zalta (ed.), *Stanford encyclopedia of philosophy*, CSLI Publications winter 2017 edn. ⟨https://plato.stanford.edu/archives/win2017/entries/logical-constants/⟩.
Marantz, Alec. 1994. A late note on late insertion. In *A festschrift for dong-whee yang* Explorations in generative grammar, 396–413. Seoul: Hankuk Publishing Co.
Martí, Luisa. 2006. Unarticulated constituents revisited. *Linguistics and Philosophy* 29(2). 135–166.
Mayr, Clemens. 2019. Triviality and interrogative embedding: context sensitivity, factivity, and neg-raising. *Natural Language Semantics* 27(3). 227–278.
McGee, V. 1996. Logical operations. *Journal of Philosophical Logic* 25. 567–580.
Merchant, Jason. 2019. Ellipsis: A survey of analytical approaches. In Jeroen van Craenenbroeck & Tanja Temmerman (eds.), *The oxford handbook of ellipsis*, chap. 2. Oxford, UK: Oxford University Press.
Quine, Willard V. 1956. Quantifiers and propositional attitudes. *The Journal of Philosophy* 53(5). 177–187.
Radford, Andrew. 2004. *Minimalist syntax: Exploring the structure of english*. Cambridge, UK: Cambridge University Press.
Recanati, François. 2004. *Literal meaning*. Cambridge: Cambridge University Press.
Recanati, François. 2010. *Truth-conditional pragmatics*. Oxford: Oxford University Press.
Rooth, Mats. 1992. A theory of focus interpretation. *Natural language semantics* 1(1). 75–116.

Rothschild, Daniel & Gabriel Segal. 2009. Indexical predicates. *Mind & Language* 24(4). 467–493. http://dx.doi.org/10.1111/j.1468-0017.2009.01371.x.

Sagi, G. 2014. Formality in logic: From logical terms to semantic constraints. *Logique et Analyse* 227. 259–276.

Sauerland, Uli. 2014. Making fuzzy logic work for language. In *Semantics and Philosophy in Europe 7*, .

Sauerland, Uli. 2017. A note on ungrammaticality and analyticity. *Snippets* 22–23. http://dx.doi.org/10.7358/snip-2017-031-saue.

Sher, Gila. 2003. A characterization of logical constants is possible. *Theoria* 18. 189–197.

Stanley, Jason. 2007. *Language in context: Selected essays*. Oxford: Clarendon Press.

Szabó, Zoltan & Jason Stanley. 2000. On quantifier domain restriction. *Mind & Language* 15(2-3). 219–61. http://dx.doi.org/10.1111/1468-0017.00130.

van Benthem, Johan. 1989. Logical constants accross varying types. *Notre dame journal of formal logic* 30. 315–42.

von Fintel, Kai. 1993. Exceptive constructions. *Natural language semantics* 1. 123–148.

von Fintel, Kai. 1995. The formal semantics of grammaticalization. In *Proceedings of NELS 25*, vol. 2, 175–189.

Williamson, Timothy. 1994. *Vagueness*. New York, NY: Routledge.