Semantic Anomaly, Pragmatic Infelicity and Ungrammaticality

Marta Abrusan

Semantic Anomaly, Pragmatic Infelicity and Ungrammaticality

Márta Abrusán

Institut Jean Nicod, CNRS, ENS, EHESS, PSL, Paris, France, 75005; email: marta.abrusan@gmail.com

Abstract

A major goal of modern syntax has been to find principles that rule out sentences that seem ungrammatical. To achieve this goal, it has been proposed that syntactically odd (or ungrammatical) sentences can be distinguished empirically and theoretically from semantically odd (or semantically anomalous) sentences. However, sometimes it is not clear why a sentence is "weird" and this has repercussions for our syntactic and semantic theories. According to a number of proposals semantic and pragmatic processes can lead to weirdness that empirically feels more like ungrammaticality than semantic oddness. But if this is so then the question arises: what explains the intuitive difference between sentences that feel ungrammatical and those that merely feel semantically (or pragmatically) anomalous? This article addresses this question by describing and comparing various semantic and pragmatic proposals for explaining different types of weirdness: ungrammaticality, semantic anomaly and pragmatic infelicity.

Keywords: semantic anomaly, grammaticality, pragmatic infelicity, natural logic, polarity items, meaning shift

1   Introduction

Linguistics has learned a lot about language by studying when things go wrong, for example when an utterance is "weird". A major goal of modern syntax has

1

been to find principles that rule out sentences that seem ungrammatical. To achieve this goal, it has been proposed that there are different types of "weirdness", in particular that syntactically odd (or ungrammatical) sentences can be distinguished empirically and theoretically from semantically odd (or semantically anomalous) sentences.[1]

(1a)   *Read you a book about modern syntax?

(1b)   #Colorless green ideas sleep furiously.
        (Chomsky 1957)

Initial psycholinguistic evidence seemed to confirm this sharp division. Kutas and Hillyard (1980, 1984) discovered that reading a semantically anomalous word in a sentence (e.g. *#John buttered his bread with socks*) produces what came to be known as the N400 effect: a strong negative ERP 400ms after the presentation of the anomalous word.[2] Simple grammatical errors, for example agreement violations did not produce the N400 effect (cf. Kutas and Hillyard 1983). In contrast, syntactic or morphological anomalies (e.g. **The cat will eating the food*) were found to elicit a large positive wave peaking at 600ms after the presentation of the anomalous word (the P600 effect) as well as a left anterior negativity (LAN) (cf. Osterhout and Holcomb 1992, 1993, Osterhout and Mobley 1995, Osterhout et al. 1996).

However, sometimes it is not clear why a sentence is "weird" and this has repercussions for our syntactic and semantic theories. In particular, it has been proposed that semantic and pragmatic processes can lead to "weirdness" that empirically feels more like ungrammaticality than semantic oddness. For example, the unacceptability of negative polarity items in unsuitable contexts as in example (2) has been argued to follow from semantic or pragmatic processes. But empirically the example feels ungrammatical rather than semantically anomalous.

(2)   *Anyone stole John's sandwich yesterday.

As more and more cases of seemingly ungrammatical sentences started to receive semantic explanations, the dichotomy of syntactic vs. semantic oddness became blurred.

On closer scrutiny it has turned out that processing data do not support a one-to-one mapping between N400 and semantic processing vs. P600 and syntactic

---

[1]I use the terms *ungrammatical* and *semantically anomalous* purely descriptively to capture the intuition of speakers. I use the terms *weird* and *unacceptable* in a neutral way to refer to any of the previous. The symbol '*' is used to mark ungrammaticality, and '#' to mark semantic anomaly (or *category mistakes*, as they are called in the philosophical literature). See Graffi (2002) for the history of the asterisk.

[2]Event related potentials (ERPs) are patterned voltage changes that can be measured on the scalp by electroencephalograms (EEG) in response to specific processing events.

processing either (cf. e.g. Coulson et al. 1998; Osterhout et al. 2012; Kutas and Federmeier 2011). Rather, it is currently thought that N400 is involved in lexical access (lookup) and can be influenced both by semantic and syntactic factors. P600 seems to be present when there is a deviation from some expected pattern or sequence. Neither effects are language specific: stimuli in the visual or musical domain can elicit N400/P600 as well.

If interpretive processes can lead to what feels like ungrammaticality, then the question arises: what explains the intuitive difference between sentences that feel ungrammatical and those that merely feel semantically anomalous or pragmatically infelicitous in a given context? This question is rarely raised in the literature[3] and even more rarely answered.

In this article I describe and compare various semantic and pragmatic proposals that were put forth in order to explain the reasons for the unacceptability of ungrammatical sentences. The first type of proposal assumes that ungrammaticality can follow from a failure to satisfy semantic conditions. The second type of proposal derives ungrammaticality from systematic pragmatic failure. The third type of proposal takes logical triviality to be responsible for various cases of ungrammaticality. I also discuss, where relevant, how these explanations differ (or not) from explanations for semantic anomaly and different types of pragmatic infelicity. I touch upon a number of different linguistic examples, but the running theme of this article will be polarity items: a number of competing semantic and pragmatic analyses were given to explain their distribution and so they can serve as test case for comparing different types of analyses. In the last section of this paper I outline a recent proposal that aims to explain what leads to the difference of intuitions between cases of semantic anomaly and ungrammaticality.

NB: In this paper I do not discuss processing studies of semantic or syntactic oddness any further, nor proposals that explain ungrammaticality via processing factors (e.g. Kluender and Kutas 1993 and much subsequent work) because the topic is so rich that it would deserve a paper on its own.

2   Semantic conditions for grammaticality

2.1   Lexical semantic conditions

Perhaps the best-known examples that argue that semantic conditions can underlie ungrammaticality judgments are polarity items. Let's take the example of the so-

---

[3]See some discussion in Giannakidou (2011) in the context of analyses of polarity items. The comparison between syntactic violations and semantic anomaly played an important role in the early days of generative syntax (Chomsky 1957, 1965) and of generative semantics, see Harris (1995) for an overview.

called negative polarity items (NPI's), shown in italics in the sentences below:

(3a)   John didn't see *anything*.

(3b)   *John saw *anything*.

(4a)   Few people *lifted a finger* to help.

(4b)   *Someone *lifted a finger* to help.

(5a)   Everyone who *budged an inch* was rewarded.

(5b)   *Three students who *budged an inch* were rewarded.

NPI's such as *anything*, *lift a finger* or *budge an inch* have a restricted distribution: they are acceptable in certain types of sentences but not in others. In example (3a), the element that licenses the NPI is negation, in example (4a) it is the quantifier *few* and in example (5a) it is the restriction of the quantifier *every*. The most influential analysis of NPIs was given by Ladusaw (1979), who (building on Fauconnier 1975) argued that the environments that allow NPI's to occur should be characterised in semantic terms: He proposed that NPIs are licensed in downward entailing (DE) contexts. This type of explanation was truly novel at the time as the first approaches to the problem were syntactic (cf. Klima 1964, Baker 1970). But now a purely semantic property was proposed to be the cause for an apparent syntactic infelicity.

Definition A context C[ ] is downward entailing iff for all α and β such that $\alpha \subseteq \beta$, C[β]$\rightarrow$ C[α]

In other words, downward entailing contexts license inferences from supersets to subsets. As the examples below illustrate, the scope of negation, the scope of the quantifier *few* and the restrictor of the universal quantifier are downward entailing, but the scope of the quantifier *some* or the restriction of the numeral *three people* are not:

(6a)   zebra$\subseteq$ animal

(6b)   John did not see an animal$\rightarrow$ John did not see a zebra.

(6c)   John saw an animal$\nrightarrow$ John saw a zebra.

(6d)   Few people saw an animal $\rightarrow$ Few people saw a zebra.

(6e)   Some people saw an animal $\nrightarrow$ Some people saw a zebra.

(6f)   Everyone who saw an animal was happy $\rightarrow$ Everyone who saw a zebra was happy.

(6g)   Three people who saw an animal were happy $\nrightarrow$ Three people who saw a zebra were happy.

The property of downward entailingness does not delineate a particular syntactic environment (*pace* Klima 1964, Linebarger 1980, Progovac 2005); it is a potential semantic characteristic of the sentences in which NPIs occur. This is also shown by the fact that NPIs are acceptable in the complement of negative-implying verbs, e.g. *lack*, as in *He lacks any sense of humour*. On Ladusaw's view, the requirement on NPIs that they should occur in DE environments is formulated as a special semantic licensing condition that must be met in order to use these expressions in a grammatical way. This means that syntax and semantics need to conspire in order to produce a grammatical output, grammaticality results from syntactic *and* semantic well-formedness (cf. also Ladusaw 1983). Thus, if a correct syntactic structure is not semantically well-formed, then such a structure can be filtered out by the semantics. The result of this semantic filtering, in the case of NPIs, is a sentence that feels ungrammatical, and not only semantically anomalous in the sense of example (1b).

A lot has been written on polarity items since the seminal work of Ladusaw, both the empirical generalisations and the accuracy of downward entailingness as a licensing condition have been challenged (See Linebarger 1980, Progovac 2005, Zwarts 1998, Kadmon and Landman 1993, Krifka 1995, von Fintel 1999, Giannakidou 1998, Chierchia 2013 for some influential proposals). But much work on polarity items has preserved the idea of explaining their distribution in terms of a semantic licensing condition. Some examples beside downward entailingness include Strawson DE (von Fintel 1999), intensionality, non-veridicality (Zwarts 1998, Giannakidou 1998), and most recently scope-licensing (Barker 2018). The licensing condition might be expressed as a filtering condition on the whole structure in which PI's appear (e.g. Ladusaw 1979, Barker 2018) or as a rule of lexical semantics (e.g. Zwarts 1998, von Fintel 1999, Giannakidou 1998). Either way of thinking entails that the unacceptability of examples (3b), (4b) and (5b) is the result of a lexical semantic condition gone wrong.

2.2 Failures of semantic composition

In the case of NPI licensing, one can identify a single lexical item (e.g., *anyone*) that requires its environment to meet special semantic conditions. When these conditions are not met, composition fails. In other cases there is no single culprit; a general failure of semantic composition can lead to ungrammaticality. An example is the analysis of so-called weak islands by Szabolcsi and Zwarts (1993). Wh-movement was considered a classical example of a syntactic process, but it shows certain restrictions that appear to be semantic in nature. For example, negation and some other negative elements block the formation of degree-questions:

(7)   *How fast didn't John swim?

But negation is not an absolute block for *wh*-movement: for example, questions about individuals are not sensitive to negation:

(8)   Which painting did John not see?

According to Szabolcsi and Zwarts's (1993) algebraic semantic analysis, example (7) is not acceptable because its logical form, represented schematically in (9), is not meaningful.

(9) $WH_{degreeX}[...NEG(...x...)]$

The reason is that negation requires a domain for which complements are defined, but such a domain is not provided in the case of questions formed over degrees. More generally, weak islands arise whenever the operation associated with a scope-taking element needs to be performed on a domain for which it is not defined.

An analysis in a similar spirit for island effects in the framework of dynamic semantics was given by Honcoop (1998). Beck (2006) explains so-called intervention effects in Korean by appealing to uninterpretability: in this case composition cannot be successful because of the interplay of the interpretation of questions and focus leads to uninterpretable structures.

In these examples there is no lexical element that is responsible for the observed ungrammaticality. The closest we can get to finding a designated culprit is to say that ultimately uninterpretability in these cases follows from the interpretation of questions.

2.3   Comparison with semantic anomaly

Szabolcsi and Zwarts (1993) explicitly liken the deviance of weak-islands to cases such as the example below:

(10)   #six mists

They argue that in this example, as in the case of (7), we observe failure of semantic composition. In example (10) we cannot apply a numeral to a mass noun: counting is not defined in the mass domain. However, there is something to explain about the difference of intuitions with respect the two examples.[4]

More generally, there is an open question why semantic failure such as that observed with polarity items or intervention gives rise to a different intuition than semantic failure that results from composing a predicate with an unsuitable argument.

---

[4]Szabolcsi and Zwarts (1993) mark both examples of weak islands and (10) with '*', but this does not quite render justice to the intuitions of my informants.

(11)   #My toothbrush is pregnant.

Intuitively, the reason why example (11) is unacceptable is because my toothbrush is not the sort of object to which the predicate *pregnant* can apply. In the linguistic and philosophical literature many competing ways to capture this intuition formally were proposed (see Magidor 2013 for a recent overview). The different types of proposals make different predictions about the unacceptability of ungrammatical vs. semantically anomalous sentences. Let's briefly review here three major types of proposals for semantic anomaly.[5]

According to one popular view, in examples such as (11) semantic composition fails because the argument is not in the domain of the function denoted by the predicate, and thus the result is uninterpretable (or 'meaningless') (cf. Russell 1908, Strawson 1952, Chomsky 1957, Drange 1966, Beall and Fraassen 2003, among others). This view does not predict a difference between cases of semantic anomaly and the ungrammatical examples discussed above.

According to another popular approach to semantic anomaly, such examples have meanings that speakers can intuitively grasp (hence are not 'meaningless'), but they cannot be used in a truthful way in any situation. Semantic composition can succeed, but the selectional restrictions of predicates (modelled as preconditions for having a truth value) are not met (cf. e.g. Halldén 1949, Goddard 1966, Thomason 1972, Lappin 1981, Asher 2011, Shaw 2015). Such selectional restrictions might be expressed as type-presuppositions that guide semantic composition cf. Asher (2011). Technically, this approach does predict a difference between the ungrammatical sentences discussed above and semantic anomaly because in the latter case the semantic failure is at the level of presuppositions. However, precisely because the difference between semantic conditions and semantic presuppositions is merely technical (and semantic conditions are often formulated technically as semantic presuppositions), we are still in need of an explanation of the difference.

On a variant of the presuppositional view, selectional restrictions are modelled as purely pragmatic presuppositions; failure to meet them results neither in meaninglessness nor in a lack of truth value. Instead, the interpretive failure of sentences such as (11) only produces pragmatic infelicity (Magidor 2013). This view predicts a clear difference of intuitions between examples of semantic anomaly and the ungrammatical examples discussed above, at least when the latter are derived via semantic conditions. We are still left with the following question: What makes sortal restrictions of lexical predicates so fundamentally different from other types of semantic conditions?

The last section of this paper outlines a proposal that aims at explaining the difference between semantic anomaly on the one hand and ungrammaticality due

---

[5]This summary approximately follows Magidor's (2013) classification.

to unmet semantic conditions/presuppositions on the other.

3 Systematic pragmatic infelicity

The pragmatic context, the norms of assertion, the assumptions of the conversational partners all influence and modulate what sentences are taken to mean in context. Traditionally, pragmatics was assigned the role of a conversational "post-processor"; but increasingly, pragmatic processes are thought to interact with other components of the language system (syntax, semantics, prosody and phonology) in a fundamental way. Most interesting for the purposes of this review is that *systematic* failure to observe general conversational principles can also lead to ungrammaticality.

3.1 Krifka's analysis of polarity items

An influential example is Krifka's (1990, 1995) analysis of polarity sensitive items. Krifka's goal was to derive the distribution of polarity items from their semantic structure and independently motivated pragmatic principles. Here I illustrate the general idea with his analysis of the NPI *any*. The analysis has a semantic and a pragmatic component. As for the semantics, NPI's are assumed to introduce a property and a set of alternative properties. In order to capture this dependency on alternatives, the semantics of NPI's makes use of structured meanings that were originally developed for the semantic analysis of focus. According to this, the denotation of a sentence containing an NPI is a triple $\langle B,F,A \rangle$, where B stands for background context, F for focus value, i.e. the relevant property introduced and A a set of alternative properties that are more restricted:

(12) $\langle \underbrace{\text{a function from properties to propositions}}_{B}, \underbrace{\text{a property P}}_{F}, \underbrace{\text{the set of subsets of P}}_{A} \rangle$

Thus the BFA-structure of the (unacceptable) sentence **Mary saw anything* is the following:

(13) **B:** a function from properties Q to the proposition that Mary saw Q

     F: the property of being a thing

     A: a set of properties that are stronger than the most general property of being a thing

If we apply the background B to the focus F, we get B(F), the proposition that Mary saw a thing. If we apply the background B to some alternative F' in A, we get B(F'), for example, the proposition that Mary saw a spoon. Note that there

8

is an ordering wrt. informativeness between B(F) and B(F'): in this case B(F') is more informative (stronger) than B(F).

Now, to the pragmatic component: The kind of BFA-structure exemplified above triggers a special pragmatic rule called SCALAR.ASSERT. This rule takes a BFA-triple and adds a strengthened proposition to the common ground. This strengthened proposition consists of (i) the proposition expressed by B(F) (ii) the negation of the alternative propositions that are informationally at least as strong as B(F).

Definition SCALAR.ASSERT (BFA)(c)=
      (i) add B(F) to the common ground c
      (ii) add the negation of every alternative B(F') that is informationally at least
      as strong as B(F), to the common ground c.

Thus SCALAR.ASSERT hardwires Gricean scalar implicatures into a special assertion operator that applies whenever a relevant scalar structure is present. Applying Scalar.Assert in the case of (13) gives us, informally,

(14)   SCALAR.ASSERT (BFA-structure for *Mary saw anything*)(c)=
      (i) add to the common ground c that Mary saw a thing
      (ii) add to the common ground c that Mary did not see any specific thing
      (e.g. that Mary did not see a spade, that Mary did not see a spoon, and so
      on for any F' in A.)

But the two conditions are contradictory because whenever there is some thing that Mary saw, that thing must also have some more specific property P. Since the two conditions are contradictory, adding both of them to the common ground c reduces c to the empty set and communication fails. In contrast, in the case of a negative sentence *Mary did not see anything* the second condition of SCALAR.ASSERT is vacuous because for every P, such that P is a subset of things, the proposition that Mary did not see a P is weaker than the proposition that Mary did not see a thing.

A crucial aspect of the approach of Krifka (1995) is that there is no licensing condition involved, the behaviour of NPIs follows from general conversational principles. Thus this approach goes one step further than approaches based on licensing conditions: it seeks to answer the question of *why* these contexts license NPIs.

The idea that uninformativity due to a systematic clash of implicatures can lead to ungrammaticality has been taken up by a number of researchers. In the domain of polarity sensitive items, Lee and Horn (1994), Lahiri (1998) and Van Rooy (2003) have proposed analyses in the spirit of Krifka (1995). The ungrammaticality of certain so-called syntactic islands has also been analysed as a result of

implicature clash, see for example Kuno and Takami's (1997) analysis of negative islands, Oshima's (2006) analysis of factive islands or Simonenko's (2015) proposal for explaining DP islands.

Another line of pragmatic approaches attributes the unacceptability of certain constructions to a clash of presuppositions. For example, Dayal's (1998) analysis of free choice *any* assumes that the reason why *any* cannot be used in episodic sentences is due to conflicting presuppositions. The peculiar distribution of certain focus particles (e.g. *auch nur, even*) has also been proposed to follow from incompatible presuppositions, which can be resolved in some semantnic environments (cf. Guerzoni 2003, Abrusán 2007b, Crnic 2011). Likewise, Abrusán (2007a, 2014) proposed that certain island constraints can be derived from systematically unmet or outright contradictory presuppositions. In another domain, Portner (2003) derives the incompatibility of the present perfect with temporal adverbials such as *yesterday* (e.g. **Mary has read Middlemarch yesterday.*) by assuming that the presupposition of the latter contradicts the presupposition of the former. Another example is Zucchi's (1995) analysis of the definiteness effect observed in *there*-existential sentences.[6]

3.2 Systematic vs. occasional pragmatic infelicity

One question that is often raised about the idea that systematic pragmatic infelicity can lead to ungrammaticality is the following: What distinguishes the ungrammatical cases from better known and more innocent pragmatic violations such as implicature cancellation or presupposition failure? For example, uttering the sentence *John stopped paying taxes* in a context in which it is known that John has never paid taxes forces the hearer to accommodate a presupposition that clashes with what is known. Similarly, the ignorance implicature of the sentence *John is at home or in the bar* will clash with the context if it is known that John is in the bar. The result in such cases might be a communicative breakdown, or a cancellation of the relevant presupposition/implicature, or an adjustment of the context, but it does not seem to involve ungrammaticality.[7]

It is for this reason that I have emphasised the word *systematic* above. In simple cases of presupposition failure or violations of Gricean maxims it is not possible to accomplish a successful speech-act because what is presupposed/implicated

---

[6]One might ask whether the presuppositions assumed in these accounts are necessarily pragmatic or whether they should or could be thought of as semantic presuppositions. The answer, it seems to me, has more to with one's general world-view and favourite theory of presuppositions than with the empirical facts.

[7]A theory for various cases of oddness (e.g. *A father of the victim arrived*) resulting from computing obligatory implicatures that conflict with common knowledge was offered in Magri (2009).

clashes with the particular context in which they are uttered. Under different circumstances, the relevant sentences could have been felicitously uttered. These cases are similar to Grice's *particularised conversational implicatures*. In contrast, the pragmatic infelicity of the type proposed by Krifka and others does not depend on the particular context chosen, any context will lead to an infelicitous assertion. Thus these cases represent a special type of pragmatic infelicity, similar to Grice's *generalised conversational implicatures*. (Oshima 2006 calls such cases 'pragmatic anomaly'.)[8] The same point holds for the analyses cited above where unacceptability was based on systematic presupposition failure that did not depend on the discourse context.

One might wonder what explains the difference of intuitions between the ungrammatical examples referred to in this section and semantic anomaly understood as a failure of pragmatic presuppositions. A possible answer is that examples of semantic anomaly, as it will be discussed in the last section, show a high degree of context sensitivity, unlike the cases above.

3.3 Recursive pragmatics

A thorny issue for the approaches based on pragmatic infelicity (or, 'pragmatic anomaly') is that (at least some of them) need to assume that the relevant pragmatic process can apply in embedded contexts, i.e., at the sub-sentential (sub-utterance) level. For example, making Krifka's (1995) approach work, in general, requires (as he points out) that the SCALAR.ASSERT operator apply to embedded constituents. One case is double licensing exemplified in (15b); here SCALAR.ASSERT needs to apply to the restrictor of the conditional:

(15a) If he knows anything about logic, he will know Modus Ponens.

(15b) If he doesn't know anything about logic, he will (still) know Modus Ponens.
(Hoeksema 1986)

Krifka himself proposes to embrace this radical assumption. As he puts it (page 245), "we must develop a framework in which illocutionary operators are part of the semantic recursion", in other words, that embedded constituents can count as illocutionary acts. This might be done, for example, if we assume

---

[8]One problem that is sometimes raised against Krifka's analysis is that trivial sentences, e.g. tautologies are also uninformative but are acceptable. However, trivial sentences can be used informatively in a given context, e.g. mathematical proofs; further we have a natural drive to shift the interpretation of the predicates in trivial sentences to make them more informative, cf. Kamp and Partee (1995). This argument has more force when raised against the approach discussed in Section 4, see discussion there.

some form of dynamic semantics. Others took such facts to show that certain implicature-like processes should be encoded at the level of truth-conditional semantic interpretation. We turn to some such approaches in the next section.

4 Ungrammaticality from triviality: a case for natural logic?

Another line of research has proposed that triviality of meaning can also lead to ungrammaticality. A well-known case concerns the acceptability of quantifiers in *there*-existential sentences. Milsark (1977) observed that certain quantifiers, that he called *strong* quantifiers, such as *every, both, neither, etc.* are unacceptable in this construction, as shown in example (16a). In contrast, weak quantifiers such as *some, few, three, etc.* are acceptable.[9]

(16a) *There is every fly in my soup.

(16b) There are some flies in my soup.

There is no obvious syntactic difference between the two examples above, and indeed most explanations of this contrast reduce it to some difference in the semantic and pragmatic meaning of the quantifiers involved (cf. Milsark 1977, Barwise and Cooper 1981; Zucchi 1995, Comorovski 1995, Keenan 2003, McNally 1998, etc.)

A prominent early explanation of the contrast by Barwise and Cooper (1981) proposed that the reason for the unacceptability of strong quantifiers in the *there*-existential construction is because the meaning of such sentences is trivial: they are true (or false) independently of the properties of the context. In contrast, a *there*-existential sentence with a weak quantifier has a meaning that is contingent on the properties of the context. This conclusion follows from the theory of generalised quantifiers and the crucial assumption that the expletive *there* denotes the domain of individuals D_e. If we grant this assumption, then example (16a) can be paraphrased as 'The set of flies in my soup is a subset of D_e' which cannot be false since the empty set is also a subset of D_e. In contrast, the meaning of (16b) can be paraphrased as 'the intersection of the set of flies in my soup and the domain D_e is not empty', which can be false if there are no flies in my soup.

This proposal, although controversial, has also been quite influential in the literature. Similar explanations have been given to a number of phenomena, for example the partitive constraint (Barwise and Cooper 1981; Ladusaw 1982, the *but*-exceptive construction (von Fintel 1993), restrictions on comparatives (Gajewski 2008), polarity items (Chierchia 2013), free-choice items (Menéndez-Benito

---

[9]A similar contrast holds for definites vs. indefinites, hence the puzzle is also sometimes called the *definiteness effect*.

2005), various island constraints and intervention effects (Fox and Hackl 2007, Abrusán 2007a, 2014, Mayr 2013). Most recently the idea has been used to explain restrictions on complement selection (Uegaki and Sudo 2017, Theiler et al. 2017).

In the case of many of these proposals, there is a natural connection to pragmatic approaches. If one assumes, following Chierchia et al. (2011), that Gricean conversational implicatures can (or should) be computed at the level of semantics via a silent operator that 'exhaustifies' alternatives (i.e. negates all stronger alternatives), then a Krifka-style pragmatic clash involving contradictory implicatures will appear at the level of truth-conditions. This is what is proposed by Chierchia (2013). Similarly, analyses for syntactic island constraints that rely on the necessary exhaustification of answers could be, in principle, cast in pragmatic or semantic terms, depending on whether exhaustification of alternatives is assumed to happen in the pragmatic or in the semantic module. This is the case in (Fox and Hackl 2007, Abrusán 2007a, 2014).

Typical semantic anomalies of the type *My toothbrush is pregnant* are not predicted to be ungrammatical by this proposal since they are not logically trivial. But a question arises in connection with sentences that express meanings that are logically trivial, for example (17a) and (17b) below: If logical triviality leads to ungrammaticality, then what distinguishes the ungrammatical examples from the examples in (17a) and (17b), which are syntactically acceptable or at worst semantically anomalous?[10]

(17a)   Every woman is a woman.

(17b)   This table is red and not red.

Intuitively, these sentences need not be tautological if the predicates *woman* and *red* are understood differently in their occurrences. For example, *woman* may be understood in a biological sense in its first occurrence, and in a psychological sense in its second occurrence. Such apparent tautologies even are a common rhetorical figure, cf. *War is war* or *Boys will be boys*. It seems that in order for such sentences to be interpreted, the occurrences of the the expressions have to be interpreted differently. A version of this idea was developed in Gajewski (2002), reviewed in the next subsection.

4.1   L-triviality

In an unpublished but widely cited paper, Gajewski proposed that there is a formally specifiable subset of trivial sentences, called 'L-trivial', whose members

---

[10]See for example Kamp and Partee (1995), Alxatib et al. (2013).

are systematically unacceptable (cf. Gajewski 2002).[11] On the standard, Tarskian conception, logical truth and consequence are defined in terms of variation of truth across all interpretations of the logical form. The logical form of sentences, as usually understood, is a level of semantic representation that preserves all semantically significant features of the sentence (e.g. its lexical content) but is free of imperfections such as ambiguity. L-triviality, in contrast, is calculated on modified logical forms in which all lexical material has been 'bleached', i.e. replaced non-uniformly by a fresh variable of the appropriate type. Gajewski (2002) calls these impoverished logical forms 'logical skeletons' (LS):[12]

Definition Logical skeletons (LS) are obtained from the logical form (LF) α as follows:
(i) Identify the maximal constituents of α containing no logical items.
(ii) Replace each such constituent with a distinct variable of the same type.

Thus the logical form (LF) of the sentence *It is raining or it is not raining* is something akin to (18a) (where *raining* stands for the proposition that it is raining), but its logical skeleton (LS) is (18b), where p and q stand for propositional variables:

(18a) LF: *raining* $\vee \neg$ *raining*

(18b) LS: $p \vee \neg q$

Given logical skeletons, L-triviality can be defined as follows:

Definition An LF constituent a of type t is L-trivial iff a's logical skeleton receives the denotation 1 (or 0) under all interpretations.

Notice that according to this definition, the sentence *It is raining or it is not raining*, even though it is a tautology, is not L-trivial. L-triviality and unacceptability are linked by the following hypothesis:

Hypothesis A sentence is ungrammatical if its LF contains an L-trivial constituent.

According to Gajewski (2002), the triviality of (16a) can be derived even after destroying the identity of non-logical predicates, simply by looking at the logical skeleton. The determiner *every* as well as the particle *there* are elements of the logical vocabulary and are preserved in the logical skeleton. Lexical predicates such as *every fly in my soup* are replaced by a fresh variable of the appropriate type. Therefore the logical skeleton of (16a) looks as follows:

---

[11]Gajewski's (2002) original condition is called 'L-analyticity', but this might be slightly misleading, hence we use L-triviality, similarly to some other authors, e.g. Del Pinal (2017).

[12]Gajewski (2002) works with the notion of logical form that arises from generative grammar, cf. Heim and Kratzer (1998).

(19a)   LS of (16a): there is every P.

(19b)   Interpretation: $[\![\text{every}]\!]$ (I(P)) (D$_e$)

Whatever the interpretation of P is in the model, the meaning assigned to the logical skeleton in (19a) is tautologous.

In contrast, a sentence such as (17a) is not L-analytical: once we remove the identity of the non-logical expressions (in this case, *woman*) when constructing the logical skeleton, we cannot deduce the triviality any more.

(20a)   LS of (17a): every P is Q

(20b)   Interpretation: $[\![\text{every}]\!]$ (I(P)) (Q)

As a consequence, such sentences are not predicted to be ungrammatical by Gajewski's (2002) system.

Under the resulting picture, grammar itself is endowed with the capacity to calculate L-triviality. This means that the grammar of natural languages has to include (or at least interact with) a system of 'natural logic', or a 'natural deductive system' (see Fox and Hackl 2007, Chierchia 2013, Del Pinal 2017). By assumption, this deductive system is blind to conceptual information and cannot 'see' non-logical terms, it operates only on the basis of functional terms. This in turn presupposes that terms can be sorted into two non-overlapping classes, lexical terms and functional (or logical) terms.

In the remainder of this section I first discuss the application of L-triviality to Chierchia's analysis of polarity items. Second, I go on observing some problems that motivate us to look for an alternative explanation.

4.2   Application to Chierchia's analysis of polarity items

Chierchia's (2005, 2013) analysis of polarity items is based on Krifka's (1990, 1995) approach, but with some important differences. Chierchia adopts the idea that NPIs contribute a property whose alternatives are subproperties. He also adopts the idea that there is a communicative pressure to make the strongest claim possible with scalar items. However, instead of the SCALAR.ASSERT operator, this is achieved by a silent exhaustivity operator (O) whose meaning is roughly equivalent to *only*. Similarly to SCALAR.ASSERT, the presence of an exhaustivity operator is triggered by the presence of grammatically-encoded alternatives. In contrast to Krifka's approach, exhaustification happens at the level of semantics, and is not tied to making a speech act.

Thus Chierchia assumes that *any* associates with focus, causing it to trigger alternatives. *Any* also has a lexical requirement that its active alternatives must range over subdomains D' of the domain D that *any* associates with. *Any* furthermore requires that these alternatives be exhaustified by an exhaustive operator O:

the operator O negates every alternative in the domain that is not entailed by the original sentence. In the case of a sentence such as *Anything is blue* its logical form, truth conditional import and focus alternatives are as follows:

(21a)   LF of *Anything is blue*: O[any$_{F,D}$ thing is blue]

(21b)   the truth conditional import of (21a): $\exists x \in D[thing(x) \wedge blue(x)]$

(21c)   the alternatives of (21a): { $\exists x \in D'$ [thing(x)$\wedge$blue(x)] | D'$\subseteq$ D }

Just as we saw in the analysis of Krifka, each alternative in (21c) entails the assertion in (21b). As the alternatives in (21c) are not entailed by (21b), exhaustifying the assertion via the O operator will negate them. This leads to a contradiction because the resulting statement implies that there there is some thing that is blue (= the meaning of (21b)) but there is no particular thing such that is blue (=the negation of each alternative in (21c)).

The interpretation of the LF in (21a) is thus a contradiction, but is it also L-trivial in Gajewski's sense? Only with a number of auxiliary assumptions. Firstly, *any* and the silent operator O need to be part of the logical vocabulary. This is a stipulation, but not problematic in itself. However, Chierchia's derivation is encoded by a set of lexical restrictions introduced by *any*. These restrictions, namely that *any* associates with focus, the requirement that its active alternatives must range over subdomains of the domain D that *any* associates with, the requirement that the alternatives have to be exhaustified by the operator O are presumably encoded as presuppositions of *any*. Accordingly, Gajewski's proposal needs to be modified in order to take presuppositions into account. Such a modification was offered in Gajewski (2008) (for slightly different reasons):

Definition A sentence S is L-trivial iff S's logical skeleton receives the truth value 1 (or 0) on all interpretations <u>in which it is defined</u>.

Another necessary auxiliary condition is that when constructing the logical skeleton, the variables that replace lexical predicates have to be given the same assignment in all the alternatives as in the LS of the original sentence. For example, in the case of (22c), the variables P and Q that replace the non-logical vocabulary have to be given the same assignment in each of the alternatives that they had in the LS of the original sentence; otherwise triviality does not follow.

(22a)   LS of *Anything is blue*: O[any$_{F,D}$ P is Q]

(22b)   the truth-conditional import of (22a): $\exists x \in D [P(x) \wedge Q(x)]$

(22c)   the alternatives of (22a): { $\exists x \in D'$ [P(x)$\wedge$Q(x)] | D' $\subseteq$ D }

Similar auxiliary assumptions need to be used for many accounts that rely on L-triviality, e.g. Gajewski (2008), Fox and Hackl (2007) Abrusán (2014).

4.3 Some problems

Defining logical words Gajewski's (2002) proposal hinges on distinguishing two types of vocabulary, logical and non-logical vocabulary. Despite intuitions about the existence of the two classes of words, finding a precise semantic difference has proven difficult. The most frequently cited idea, borrowed from a tradition in logic aimed at defining logical constants, is that function words have meanings that are invariant across certain algebraic transformations of their domains.[13] Examples of transformations that have been proposed to diagnose logical constants include invariance under permutations (Tarski and Givant 1987, van Benthem 1989, Sher 1991), invariance under surjective functions (Feferman 1999), invariance under potential isomorphisms (Bonnay 2008), etc. The underlying idea is that logical meanings are topic-independent: the validity of a logical inference should not depend on the particular properties of what we are talking about. The appropriateness of the above ideas for diagnosing logical constants is a subject of lively debate; but they are clearly unsuitable for diagnosing function (logical) words of natural language (see Gajewski 2002, van Benthem 2002). This is because they predict certain lexical items to be logical (e.g. the predicates *self-identical, exist*), and they also predict that certain intuitively logical elements of natural language, e.g. the quantifier *every* or *each* are not logical since they have a lexical restriction that they need to quantify over countable objects, hence **Every/*Each milk is in the fridge*.

It seems that there is no foolproof method that can distinguish logical words from non-logical ones that also makes the cut in a linguistically intuitive way. The difficulties mentioned above suggest that the logical/grammatical aspects and the conceptual aspects of meaning do not map neatly onto two different classes of words. Instead, both functional and lexical words have logical and conceptual aspects of meaning, packaged together (see also Abrusán et al. 2018)

An alternative possibility, one that Gajewski also considers, is to replace the logical/non-logical distinction with the functional-lexical distinction familiar from the linguistic literature (see for example Abney 1987, von Fintel 1995). This too suffers from difficulties, as some words, for example prepositions or the word *there* are not clear cases of either category.

An exotic deductive system Gajewski's (2002) proposal has profound implications for how we should think about the language system and its interaction with other cognitive systems in general. If L-triviality can have implications for grammaticality, grammar needs to interact with a natural deductive system (cf. also Fox

---

[13]Another idea that was advanced in the linguistic literature is that function words involve higher types than lexical items (cf. Partee 1992). See also MacFarlane (2017) for a review of the philosophical literature on logical constants.

and Hackl 2007, Chierchia to appear). The idea that grammar might be related to some system of natural logic is not in itself new; it has been much discussed in the context of the semantics of quantifiers, monotonicity reasoning, polarity items etc. (see e.g. van Benthem 1986, 2008, Dowty 1994, Moss 2015). What is new is that L-triviality suggests a radical form of modularity of language: grammar is insulated not only from conceptual systems, general world knowledge, but also from most of the information encoded in lexical items. At the same time, grammar needs to have a privileged access to some aspects of the meaning of lexical items, namely their types and (other) presuppositions that they induce.

As it was emphasised in Del Pinal (2017), a deductive system that operates on logical skeletons is a rather exotic system for which most classical formulas and rules of inference are invalid. It is conceivable that the properties of the natural deductive system, as used by grammar, could indeed be radically different from classical systems. However, Del Pinal (2017) argues that certain key accounts depending on logical skeletons, e.g. Chierchia's (2013) account of polarity-sensitive items, cannot be maintained if the Law of Non-Contradiction is invalid at the level of representation where the deductive system determines grammaticality.

Del Pinal (2017) therefore proposes a modification of Gajewski's (2002) idea which allows triviality to be calculated on standard logical forms instead of skeletons. He defends a semantic version of contextualism according to which the interpretation of lexical predicates (or 'open class items') can be modulated by a polymorphic operator RESCALE. This covert operator allows for a more specific interpretation of lexical items and can be freely inserted at LF.

Definition RESCALE
    For any open class term P, argument of suitable type x and context c,
    $\{x : \text{RESCALE}_c(P)(x)\} \subseteq \{x : P(x)\}$

Crucially, the context sensitive parameter of RESCALE can be fixed differently at each position in a sentence and RESCALE cannot modulate the meaning of closed-class items. Given this, we can see that a sentence such as (17a) can have the LF in (23a), which might get a non-tautologous reading. Notice that the two RESCALE operators are independent of each other and we have a non-tautologous meaning if they select two different specific meanings of the predicate. On the other hand, RESCALE does not help with (16a) since it can only modulate lexical predicates, whose meaning is irrelevant for deriving the triviality of the sentence:

(23a)   Every RESCALE_c(woman) is a $\text{RESCALE}_{c'}$(woman).

(23b)   There is every RESCALE_c(fly in my soup).

Del Pinal's (2017) solution has a number of advantages: It resolves the problems that arise with exotic deductive systems and allows us to stick to better known

systems. Il also resolves certain empirical problems with L-triviality concerning bound variables and reflexives (see Gajewski 2009, Sauerland 2017, Chierchia to appear). However, it still crucially depends on a sharp division of lexical vs. logical vocabulary. Moreover, the proposal is slightly unintuitive in the case of examples such as (24a), in which semantic modulation of the meaning of the predicates *raining* in the antecedent would clash with the consequent:

(24a) John believes that it is raining and it is not raining, therefore he has inconsistent beliefs.

(24b) *possible LF of (24a)*: John believes that it is RESCALE$_c$(raining) and it is not raining, therefore he has inconsistent beliefs.

According to Del Pinal (2017), the first clause of the above sentence is saved from being ungrammatical because its meaning can be made non-contradictory, in theory, by inserting a RESCALE operator that modulates at least one occurrence of the predicate *raining*. Grammar is sensitive to the presence of this operator, which allows for the possibility of interpreting the the two occurrences of the predicate *raining* differently and thus resolving the trivial meaning. However, the actual fixing of the context parameter of the RESCALE operator happens at the pragmatic level and it might, in theory, pick out a non-specific reading of the predicate *raining*. This is what explains why the second clause in (24a) does not seem to contradict the first clause. The insertion of the RESCALE operator is vacuous from the point of view of the actual interpretation of the sentence, but its presence is enough to make the antecedent potentially-non trivial and hence grammatical.

5   Future directions

Where do we stand now? We have seen a number of different types of semantic and pragmatic analyses of ungrammatical sentences. We have seen how systematic pragmatic infelicity ('pragmatic anomaly') differs from occasional pragmatic infelicity. However, no clear answer emerged as to what distinguishes sentences that are felt merely as semantically anomalous from other cases of semantically or pragmatically deviant sentences that are felt ungrammatical.

In what follows I outline a recent proposal by Abrusán et al. (2019) that argues that the underlying cause for both semantic anomaly (e.g. *#My toothbrush is pregnant*) and ungrammaticality is failure to observe semantic compositional rules (including semantic-presuppositions).[14] According to this proposal the difference

---

[14]This proposal uses Asher's (2011) formal framework in which presuppositions of lexical items are thought of as semantic (type)-presuppositions. However, the gist of the idea could also be

comes about because cases of semantic anomaly, in contrast to ungrammatical examples, could be repaired by meaning shift, however far-fetched a scenario that would require.[15]

This approach, in a sense, asks the reverse of the question we have been asking so far: Why don't the examples of semantic anomaly feel ungrammatical? In the case of (1b), semantic composition fails because the meaning and the selectional restrictions of the predicates give us a meaning that is nonsensical. But in contrast to many ungrammatical examples, we know exactly where the locus of inconsistency lies and we also have the ability to shift the meaning of the predicates in an interpretive context in order to try to resolve this inconsistency. (See the results of a Stanford competition in 1985 to find contexts in which (1b) could be meaningfully uttered.[16]) Contextual shifting of the meaning of words has been a major topic in the philosophy of language.[17] Some of the most extreme cases of meaning shift are demonstrated by metaphors (25a) and metonymy (25b) (e.g. coercion, i.e. a reinterpretation process triggered by a mismatch between the selectional restrictions of a function and its argument):

(25a)   Language is the breath of God, the dew on a fresh apple, it's the soft rain of dust that falls into a shaft of morning sun when you pull from an old bookshelf a half-forgotten volume of erotic memoirs. (Stephen Fry[18])

(25b)   Hugh enjoyed the monologue. (e.g. 'listening to the monologue')

More mundane meaning shifts happen with most examples of predication, e.g. in adjective-noun composition. As shown by the examples below, the meaning of the adjective *heavy* is modulated depending on the noun that it modifies:

(26a)   heavy box

(26b)   heavy traffic

---

expressed in a framework in which presuppositions are thought of as pragmatic. The usefulness of a formal system using rich types comes, in part, from the fact that it is easier to integrate it with a computational semantic framework.

[15]Logically trivial sentences such as *It is raining an not raining* do not show a compositionality problem and are not semantically anomalous. I return to this below.

[16]An example entry from the competition, from C.M. Street, is: "It can only be the thought of verdure to come, which prompts us in the autumn to buy these dormant white lumps of vegetable matter covered by a brown papery skin, and lovingly to plant them and care for them. It is a marvel to me that under this cover they are labouring unseen at such a rate within to give us the sudden awesome beauty of spring flowering bulbs. While winter reigns the earth reposes but these colourless green ideas sleep furiously".

[17]A growing body of work assumes that the meaning of lexical words can be shifted or modulated in one way or another: either within the semantics (cf. e.g. Martí 2006, Stanley 2007, Alxatib et al. 2013) or within the pragmatics (e.g. Kamp and Partee 1995, Recanati 2010, Lasersohn 2012)

[18]see https://www.youtube.com/watch?v=0gjWUiVtldk

(26c)   heavy bleeding

(26d)   heavy drinker

Meaning shifts can happen with more or less of an ease. It has been observed that semantic anomaly is a graded phenomenon (cf. Drange 1966, Magidor 2013, Lappin 1981). For example, it is increasingly more difficult to interpret the examples below:

(27a)   #Tigers are human.

(27b)   #Tigers are buildings.

(27c)   #Tigers are Zermelo-Fraenkel sets.

One way to think about this gradience is in terms of distance within a semantic space. Intuitively, tigers and humans are fairly close in meaning in that they share a number of features, e.g. they are both animate. In contrast, tigers and Zermelo-Fraenkel sets do not have much in common. Let's assume that selectional restrictions of predicates and other lexical conditions are modelled formally as type-presuppositions in a richly-typed system as in Type-Composition Logic, TCL (cf. Asher 2011). Types are organised into a type hierarchy that represents the conceptual organisation of semantic categories and allows for a calculation of semantic distances. Applying meaning shift to the examples in (27a-c) means that we can try to relax the type presuppositions in question. Depending on the semantic distance between the actual and the target types, type shift allows us to make sense of examples of semantic anomaly to a greater or smaller degree. The higher the type that is involved in the type clash, the more difficult is it to shift the meaning of predicates to resolve type conflict. Nevertheless, even with examples such as (27c), we can glean what could have been said. The contradiction between type presuppositions in these examples is somehow localized. And by relaxing those type presuppositions, for instance by supposing that the predicate in example (27c) simply is seeking an entity of general type E, the composition could actually have succeeded.

The problems of composition in ungrammatical sentences are worse in that no shifting of the type presuppositions seems possible. These types are somehow constitutive of the construction in a way that simpler type presuppositions of open class expressions, nouns, verbs and adjectives, are not. Abrusán et al.'s (2019) hypothesis is that the type presuppositions that are violated in ungrammatical sentences are constitutive of the type of denotation at the highest level of the type hierarchy. For instance, it is something about the nature of existential predication that leads to the uninterpretability of *there is every girl*; it is something about the semantic type of questions that leads to the uninterpretability of weak island sentences, and so on. (I came back to these ideas below.)

Another property that might underlie the difference between semantic anomaly and certain ungrammatical examples is the locality of the type-conflict. In the case of classic examples of semantic anomaly, type conflict (presupposition failure) arises at the level of predicate-argument composition. If shifting were to occur, it would also happen at this level, where the nature of the type-conflict is clear and lowest common types are easy to calculate. In the case of many ungrammatical examples the type conflict arises at a more global level with more linguistic elements (and types, as a consequence) that are involved. Calculating lowest common types is harder, and it is more likely there are simply none.

Meaning shifts in Distributional Semantics   Abrusán et al. (2019) assume that meaning shift can be diagnosed and calculated precisely with the help of a computational method called distributional semantics (DS). A simple-minded way to think about word meaning in this framework is to assume that it is a vector in space V whose dimensions are contextual features.[19] This computational method can find low-level types and calculate distances in semantic space as well as meaning shift involved in semantic composition (cf. Asher et al. 2016). It predicts, e.g. shifts that occur with the various occurrences of *heavy* in (26a-d) or meaning shifts enforced by the context, e.g. the shift from *Julie enjoyed the monologue* to *Julie enjoyed listening to the monologue*. It does not, however, predict shifts with high-level types (e.g. the types for questions or quantifiers) for at least two reasons. Firstly, these types denote context-invariant logical meaning that is simply invisible for distributional methods. Second, the conflicting type-clash is attached to a type so high in the type hierarchy that the type has no neighbours that share the same syntactic/semantic dependencies and so there is nowhere in the space for its meaning to shift. Thus types that are at a maximally general level, for example questions, whose type is a family or set of propositions, propositions themselves, the general type E of entities, the general type of determiner phrases DP, quantifiers or second order properties etc. don't have any neighbours in a vector space, in the way that say a common noun like *tiger* or *idea* does. This is because other expressions that are not of say DP type will belong to a different syntactic category with different syntactic/semantic dependencies; they will be in a different space (though they may share the same latent dimensions as particular DPs). As a result, there is *nowhere in the space* for such a type to shift and preserve its corresponding syntactic category. DS methods *can't shift* those types, in the way that they can shift the meaning of lexical words in context. Furthermore, because these types don't have neighbours in the vector space if indeed they inhabit a vector

---

[19]More complex methods attempt to build a distributional representation in step with the syntactic derivation (Coecke et al. 2010, Baroni et al. 2014), or use neural networks (e.g. Socher et al. 2012).

space at all, these semantic principles will not show up in particular latent dimensions of the vector space. DS methods, at least of the sort Asher et al. (2016) and Abrusán et al. (2019) have employed, will not be able to see these principles. For expressions that have type presuppositions *at this level*, it follows that they cannot be shiftable.[20]

On the other hand, one can imagine another kind of shift — one that rescues some sort of content from a predication in which there is an irresolvable type clash between a predicate and its argument. (1b) or (27a-c) are examples of this. If we simply move to the supertype of the type presuppositions, we can see that the author of (1b) or (27a-c) was predicating a property of some object; she was just confused about or wilfully misusing the meaning of the property or the object expression. The distance between the type presuppositions of the predicate and arguments may be great, but it is still defined, as *tiger* and say *ordinal* are both in the same vector space DS defines for common nouns. And thus DS and TCL together provide us with a means to distinguish (1b) and (27a-c) from ungrammatical examples. [21]

Examples   This approach allows for an integration of various perspectives mentioned above, in particular, approaches based on a violation of a semantic condition and presupposition clash. For example, in the case of polarity items such as *any* we might assume, in accordance with various authors cited in Section 2, that its distribution is governed by a lexical semantic condition. (Possibly, Krifka's (1995) account might be viewed as providing a functional motivation for this semantic condition.) As for existential sentences, we might adopt McNally's (1998) proposal according to which the existential predicate has a sortal restriction on its postverbal argument, requiring it to be a nonparticular (formalised as properties). An argument in favour of this view comes from examples such as (28b), which shows that the restriction is not simply on the type of quantification:

(28a)   *There was every doctor at the convention.

---

[20]An interesting question raised by an anonymous reviewer is what happens in the case of number words. These have logical meanings, but they also have neighbours in vector space. Possibly we might observe shifting in the granularity standard associated with numerical expressions, e.g. as in *5.3 grams of radioactive material* vs. *50 thousand protesters on the street*. But one would need to perform experiments to be able to tell.

[21]Abrusán et al.'s (2019) proposal is explicitly about semantic anomaly and ungrammatical sentences for which semantic/pragmatic analyses have been proposed. It remains to be seen whether this approach can be generalised to all cases of ungrammaticality. The idea that ungrammaticality is a compositionality problem at the syntactic level, is a rather classical one. If one assumes a categorial syntax accompanied by type-driven composition, e.g. Klein and Sag (1985), than syntactic compositionality problems ultimately boil down to a type-clash of some very basic types. Such compositionality problems are very similar to the type clash of unshiftable high-types observed in TCL and thus a unified analysis seems promising.

(28b)   There was every kind of doctor at the convention.

This view also allows for integrating pragmatic accounts, e.g. those based on systematic presupposition failure such as Abrusán's (2014) analysis of weak islands, as these can be expressed as type conflicts. See Abrusán et al. (2019) for more details.

Meaning shift vs. triviality   Triviality of meaning, on the other hand, does not play any role in this explanation: trivial sentences do not show a compositionality problem and are thus predicted to be acceptable. (Meaning shift might still apply to them in order to increase their informativity, as in the case of (17a).) Nor is there any reason to assume a sharp conceptual division between lexical and logical vocabulary. The conceptual content of lexical as well as logical words might have both logical aspects (their model theoretic meaning) and lexical ('worldly') aspects. For example, the conceptual content of a universal quantifier such as *every* might be composed from its logical content (i.e. universal quantification) and a more lexical content, e.g. it's lexical restriction to countable predicates. We cannot neatly separate grammar and conceptual knowledge because they are packaged together within lexical entries. However, we can distinguish conceptual content that is contextually invariant from shiftable conceptual content (cf. Abrusán et al. 2018). Conceptual content that supports logically valid inferences for first order definable quantifiers (whose conceptual content in the form of proof rules can determine all logical consequences of such quantifiers) should always be contextually invariant, since a particular context should not render logically valid inferences incorrect. We would never expect meaning components that lead to logically valid inferences to shift in context. Other types of conceptual content of logical words, however, can be shiftable; that is, context might affect the conceptual content and the extension of a predicate or its argument as well.

6   Conclusion

We find out how things work when they are broken. In linguistics, it has often been assumed to be intuitively clear to which domain the source of the problem of an unacceptable sentence can be attributed: phonology, syntax, semantics, processing, etc. However, the history of the field shows that problems that appear at first sight to be belong to one domain (e.g. syntax) are sometimes later shown to have a semantic, pragmatic (or processing) motivation. One example, discussed in this paper, is the case of NPIs.

Once we accept that different types of "weirdness" can have similar interpretive sources, a new question arises: what leads to a difference in intuitions if the

source of the problem is the same? In particular, what distinguishes ungrammatical examples from cases of semantic anomaly or pragmatic infelicity? This paper has reviewed various proposals that address this question and also outlined a possible new direction.

Acknowledgments

References

Abney, S. (1987). *The English Noun Phrase in Its Sentential Aspect.* Ph. D. thesis, MIT.

Abrusán, M. (2007a). *Contradiction and Grammar: the Case of Weak Islands.* Ph. D. thesis, M.I.T.

Abrusán, M. (2007b). Even and free choice any in Hungarian. In *Proceedings of Sinn und Bedeutung*, Volume 11, pp. 1–15.

Abrusán, M. (2014). *Weak Island Semantics.* OUP Oxford.

Abrusán, M., N. Asher, and T. V. de Cruys (2018). Content vs. function words: The view from distributional semantics. In *Proceedings of Sinn und Bedeutung 22*.

Abrusán, M., N. Asher, and T. V. de Cruys (2019). Grammaticality and meaning shift. In G. Sagi and J. Woods (Eds.), *The Semantic Conception of Logic*. Cambridge University Press, forthcoming.

Alxatib, S., P. Pagin, and U. Sauerland (2013). Acceptable contradictions: Pragmatics or semantics? *Journal of Philosophical Logic 42*, 619–634.

Asher, N. (2011). *Lexical Meaning in Context.* Cambridge University Press.

Asher, N., T. Van de Cruys, A. Bride, and M. Abrusán (2016). Integrating type theory and distributional semantics: A case study on adjective–noun compositions. *Computational Linguistics 42*(4), 703–725.

Baker, C. L. (1970). Double negatives. *Linguistic inquiry 1*(2), 169–186.

Barker, C. (2018). Negative polarity as scope marking. *Linguistics and Philosophy to appear*.

Baroni, M., R. Bernardi, and R. Zamparelli (2014). Frege in space: A program for compositional distributional semantics. *Linguistic Issues in Language Technology 9*, 241–346.

Barwise, J. and R. Cooper (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy 4*(2), 159–219.

Beall, J. C. and B. Fraassen (2003). *Possibilities and Paradox.* Oxford.

Beck, S. (2006). Intervention effects follow from focus interpretation. *Natural Language Semantics 14*(1), 1–56.

van Benthem, J. (1989). Logical constants across varying types. *Notre Dame Journal of Formal Logic 30*(3), 315–342.

van Benthem, J. (2002). Invariance and definability: two faces of logical constants. essays in honor of sol feferman. In W. Sieg, R. Sommer, and C. Talcott (Eds.), *Reflections on the Foundations of Mathematics.*, ASL Lecture Notes in Logic 15., pp. 426–446.

Bonnay, D. . (2008). Logicality and invariance. *Bulletin of Symbolic Logic 14*(1), 29–68.

Chierchia, G. (2004). Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In A. Belletti (Ed.), *Structures and Beyond*, Volume 3. Oxford University Press.

Chierchia, G. (2013). *Logic in grammar: Polarity, free choice, and intervention.* OUP Oxford.

Chierchia, G. (2019). On being trivial: Grammar vs. logic. In G. Sagi and J. Woods (Eds.), *The Semantic Conception of Logic*. Cambridge University Press.

Chierchia, G., D. Fox, and B. Spector (2011). The Grammatical View of Scalar Implicatures and the Relationship between Semantics and Pragmatics. In K. von Heusinger, C. Maienborn, and P. Portner (Eds.), *Handbook of Semantics*. Mouton de Gruyter. Draft version available at `http://semanticsarchive.net/Archive/WMzY2ZmY/CFS_EmbeddedSIs.pdf`.

Chomsky, N. (1957). *Syntactic Structures*. Mouton.

Chomsky, N. (1965). Aspects of the theory of syntax. *Cambridge: MIT Press*.

Coecke, B., M. Sadrzadeh, and S. Clark (2010). Mathematical foundations for a compositional distributed model of meaning. *Lambek Festschrift, Linguistic Analysis, vol. 36 36*.

Comorovski, I. (1995). On quantifier strength and partitive noun phrases. In *Quantification in natural languages*, pp. 145–177. Springer.

Coulson, S., J. W. King, and M. Kutas (1998). Expect the unexpected: Event-related brain response to morphosyntactic violations. *Language and cognitive processes 13*(1), 21–58.

Crnic, L. (2011). *Getting even*. Ph. D. thesis, MIT.

Dayal, V. (1998). Any as inherently modal. *Linguistics and philosophy 21*(5), 433–476.

Del Pinal, G. (2017). The logicality of language: Triviality and logical form. *Noûs (forthcoming)*.

Dowty, D. (1994). The role of negative polarity and concord marking in natural language reasoning. In *Semantics and Linguistic Theory*, Volume 4, pp. 114–144.

Drange, T. (1966). *Type crossings*. The Hague: Mouton.

Fauconnier, G. (1975). Pragmatic scales and logical structure. *Linguistic Inquiry*, 353–375.

Feferman, S. (1999). Logic, logics, and logicism. *Notre Dame Journal of Formal Logic 40*(1), 31–54.

von Fintel, K. (1993). Exceptive constructions. *Natural Language Semantics 1*(2), 123–148.

von Fintel, K. (1995). The formal semantics of grammaticalization. In *Proceedings of NELS 25*, pp. 175–189.

von Fintel, K. (1999). Npi licensing, strawson entailment, and context dependency. *Journal of Semantics 16*(2), 97–148.

Fox, D. and M. Hackl (2007). The Universal Density of Measurement. *Linguistics and Philosophy 29*, 537–586.

Gajewski, J. (2002). On Analyticity in Natural Language. Ms., MIT.

Gajewski, J. (2008). More on quantifiers in comparative clauses. In *Semantics and Linguistic Theory*, Volume 18, pp. 340–357.

Gajewski, J. (2009). L-triviality and grammar. Handout, UConn Logic Group, University of Connecticut.

Giannakidou, A. (1998). *Polarity sensitivity as (non) veridical dependency*. Amsterdam: John Benjamins.

Giannakidou, A. (2011). Negative polarity and positive polarity: Licensing, variation, and compositionality. In K. von Heusinger, C. Maienborn, and P. Portner (Eds.), *The handbook of natural language meaning*, pp. 1660–1712. Mouton de Gruyter Berlin.

Goddard, L. (1966). Predicates, relations and categories. *Australasian Journal of Philosophy 44*(2), 139–171.

Graffi, G. (2002). The asterisk from historical to descriptive and theoretical linguistics: An historical note. *Historiographia linguistica 29*(3), 329–338.

Guerzoni, E. (2003). *Why even ask?: on the pragmatics of questions and the semantics of answers*. Ph. D. thesis, Massachusetts Institute of Technology.

Halldén, S. (1949). The logic of nonsense.

Harris, R. A. (1995). *The linguistics wars*. Oxford University Press.

Heim, I. and A. Kratzer (1998). *Semantics in Generative Grammar*. Oxford: Blackwell.

Hoeksema, J. (1986). *Monotonicity phenomena in natural language*. na.

Honcoop, M. (1998). *Dynamic Excursions on Weak Islands*. Ph. D. thesis, Leiden.

Kadmon, N. and F. Landman (1993). Any. *Linguistics and philosophy 16*(4), 353–422.

Kamp, H. and B. Partee (1995). Prototype theory and compositionality. *Cognition 57*(2), 129–191.

Keenan, E. L. (2003). The definiteness effect: Semantics or pragmatics? *Natural Language Semantics 11*(2), 187–216.

Klein, E. and I. A. Sag (1985). Type-driven translation. *Linguistics and Philosophy 8*(2), 163–201.

Klima, E. S. (1964). Negation in English. In *The Structure of Language*. Englewood Cliffs: Prentice-Hall.

Kluender, R. and M. Kutas (1993). Subjacency as a processing phenomenon. *Language and cognitive processes 8*(4), 573–633.

Krifka, M. (1990). Polarity phenomena and alternative semantics. In *Proceedings of the seventh Amsterdam Colloquium*, pp. 277–301. ITLI Amsterdam.

Krifka, M. (1995). The semantics and pragmatics of polarity items. *Linguistic analysis 25*(3-4), 209–257.

Kuno, S. and K. Takami (1997). Remarks on negative islands. *Linguistic inquiry 28*(4),

553–576.

Kutas, M. and K. D. Federmeier (2011). Thirty years and counting: finding meaning in the n400 component of the event-related brain potential (erp). *Annual review of psychology 62*, 621–647.

Kutas, M. and S. A. Hillyard (1980). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science 207*(4427), 203–205.

Kutas, M. and S. A. Hillyard (1983). Event-related brain potentials to grammatical errors and semantic anomalies. *Memory & cognition 11*(5), 539–550.

Kutas, M. and S. A. Hillyard (1984). Brain potentials during reading reflect word expectancy and semantic association. *Nature 307*(5947), 161.

Ladusaw, W. (1982). Semantic constraints on the English partitive construction. In *Proceedings of the first West Coast Conference on Formal Linguistics*.

Ladusaw, W. A. (1979). *Polarity Sensitivity as Inherent Scope Relations*. Ph. D. thesis, University of Texas at Austin.

Ladusaw, W. A. (1983). Logical form and conditions on grammaticality. *Linguistics and Philosophy 6*(3), 373–392.

Lahiri, U. (1998). Focus and negative polarity in hindi. *Natural language semantics 6*(1), 57–123.

Lappin, S. (1981). *Sorts, ontology, and metaphor: the semantics of sortal structure*. Walter de Gruyter.

Lasersohn, P. (2012). Contextualism and compositionality. *Linguistics and Philosophy 35*, 171–189.

Lee, Y.-S. and L. Horn (1994). Any as indefinite plus even. *Manuscript, Yale University*.

Linebarger, M. (1980). *The grammar of negative polarity*. Ph. D. thesis, Massachusetts Institute of Technology.

MacFarlane, J. (2017). Logical constants. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017 ed.). Metaphysics Research Lab, Stanford University.

Magidor, O. (2013). *Category Mistakes*. Oxford University Press.

Magri, G. (2009). A theory of individual-level predicates based on blind mandatory scalar implicatures. *Natural language semantics 17*(3), 245–297.

Martí, L. (2006). Unarticulated constituents revisited. *Linguistics and Philosophy 29*.

Mayr, C. (2013). Intervention effects and additivity. *Journal of Semantics 31*(4), 513–554.

McNally, L. (1998). Existential sentences without existential quantification. *Linguistics and Philosophy 21*(4), 353–392.

Menéndez-Benito, P. (2005). *The Grammar of Choice*. Ph. D. thesis, University of Massachusetts Amherst.

Milsark, G. L. (1977). Toward an explanation of certain peculiarities of the existential construction in english. *Linguistic analysisnalysis 3*, 1–31.

Moss, L. (2015). Natural logic. In S. Lappin and C. Fox (Eds.), *The handbook of contemporary semantic theory*. John Wiley & Sons.

Oshima, D. Y. (2006). On factive islands: pragmatic anomaly vs. pragmatic infelicity. In *Annual Conference of the Japanese Society for Artificial Intelligence*, pp. 147–161. Springer.

Osterhout, L. and P. J. Holcomb (1992). Event-related brain potentials elicited by syntactic anomaly. *Journal of memory and language 31*(6), 785–806.

Osterhout, L. and P. J. Holcomb (1993). Event-related potentials and syntactic anomaly: Evidence of anomaly detection during the perception of continuous speech. *Language and Cognitive Processes 8*(4), 413–437.

Osterhout, L., A. Kim, and G. R. Kuperberg (2012). The neurobiology of sentence comprehension. *The Cambridge handbook of psycholinguistics*, 365–389.

Osterhout, L., R. McKinnon, M. Bersick, and V. Corey (1996). On the language specificity of the brain response to syntactic anomalies: Is the syntactic positive shift a member of the p300 family? *Journal of Cognitive Neuroscience 8*(6), 507–526.

Osterhout, L. and L. A. Mobley (1995). Event-related brain potentials elicited by failure to agree. *Journal of Memory and language 34*(6), 739–773.

Partee, B. (1992). Syntactic categories and semantic type. *Computational linguistics and formal semantics*, 97–126.

Portner, P. (2003). The (temporal) semantics and (modal) pragmatics of the perfect. *Linguistics and Philosophy 26*(4), 459–510.

Progovac, L. (2005). *Negative and positive polarity: A binding approach*, Volume 68. Cambridge University Press.

Recanati, F. (2010). *Truth-conditional pragmatics*. Clarendon Press Oxford.

Russell, B. (1908). Mathematical logic as based on the theory of types. *American journal of mathematics 30*(3), 222–262.

Sauerland, U. (2017). A note on grammaticality and analyticity. *Snippets 31*.

Shaw, J. R. (2015). Anomaly and quantification. *Noûs 49*(1), 147–176.

Sher, G. (1991). *The Bounds of Logic*. Cambridge, Mass: MIT Press.

Simonenko, A. (2015). Semantics of dp islands: the case of questions. *Journal of Semantics 33*(4), 661–702.

Socher, R., B. Huval, C. D. Manning, and A. Y. Ng (2012, July). Semantic compositionality through recursive matrix-vector spaces. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, Jeju Island, Korea, pp. 1201–1211. Association for Computational Linguistics.

Stanley, J. (2007). *Language in context: Selected essays.* Oxford: Clarendon Press.

Strawson, P. F. (1952). *Introduction to Logical Theory*. London: Methuen.

Szabolcsi, A. and F. Zwarts (1993). Weak islands and an algebraic semantics for scope taking. *Natural Language Semantics 1*, 235–284.

Tarski, A. and S. R. Givant (1987). *A formalization of set theory without variables*, Volume 41. American Mathematical Soc.

Theiler, N., F. Roelofsen, and M. Aloni (2017). What's wrong with believing whether. In *Semantics and Linguistic Theory*, Volume 27, pp. 248–265.

Thomason, R. (1972). A semantic theory of sortal incorrectness. *Journal of Philosophical Logic 1*, 209–258.

Uegaki, W. and Y. Sudo (2017). The anti-rogativity of non-veridical preferential predicates. In *Proceedings of the 21st Amsterdam Colloquium*, pp. 492–501.

van Benthem, J. (1986). *Essays in logical semantics*. Springer.

van Benthem, J. (2008). A brief history of natural logic. In *Logic, Navya-Nyaya & applications: Homage to Bimal Krishna Matilal & applications: Homage to Bimal Krishna Matilal*, pp. 21–42. London: College Publications.

Van Rooy, R. (2003). Negative polarity items in questions: Strength as relevance. *Journal of semantics 20*(3), 239–273.

Zucchi, A. (1995). The ingredients of definiteness and the definiteness effect. *Natural Language Semantics 3*(1), 33–78.

Zwarts, F. (1998). Three types of polarity. In *Plurality and quantification*, pp. 177–238. Springer.