Keeping *fake* simple*

Janek Guerrini

December 2022

Abstract

The semantic complexity of 'fake' has pushed some theorists to argue that our traditional understanding of the meaning of NPs is too simplistic. In this paper, I argue that by embedding similarity into the meaning of 'fake', we can avoid this complexification of the meaning of lexical items. Moreover, by basing the meaning of 'fake' on similarity and hard-coding negation into its meaning, we can make fine-grained and correct predictions on iterated application of 'fake', or on the application of 'fake' to expressions containing other intensional modifiers like 'typical'. Extant theories either do not make fine-grained enough predictions, or make wrong predictions for iteration. More precisely, I propose that a fake *P* is (i) intended to resemble a *P* and (ii) is not a *P*. In double application of 'fake', the interaction between this conjunction and the negation hard-coded into clause (ii) yields a complex truth table. I show in detail that all the objects predicted to be fake fake *P*s by this truth table do in fact qualify as such.

*This paper is in part a descendant of a Qualifying Paper that I wrote under the supervision of Salvador Mascarenhas, whom I thank for first encouraging me to work on this topic, as well as for countless discussions that played an essential role in the development of my ideas. I would also like to thank Benjamin Spector for his supervision and his substantial input at a later stage of this project. I am thankful to Gennaro Chierchia for his valuable input. Thank you, finally, to Guillermo Del Pinal, Philippe Schlenker, Emmanuel Chemla, Paul Egré, Jeremy Kuhn, Diego Feinmann, and Amir Anvari for their very helpful comments.

Contents

1 Introduction

What does it take for an object to be called a fake gun? Certainly, our intuition is that it shouldn't be a gun. But not any object that isn't a gun is a fake gun: there is some relation that a fake gun must bear to guns. In this paper, I argue that this relation is best thought of as one of resemblance: a fake *N* is intended to resemble a *N*.

Understanding the meaning and compositional behavior of 'fake' matters in the context of a larger phenomenon known as privative modification. Privative modification occurs whenever a modifier returns a denotation that is not included in the denotation of its input noun.

(1) $[\![Adj\ N]\!] \not\subseteq [\![\ N]\!]$

This class of adjectives constitutes a long-standing problem for semantic theory, for at least two reasons.

First, they do not obey certain *prima facie* intuitive generalizations. In particular, their behavior runs against the appealing idea that $[\![Adj\ N]\!]$ should be a subset of $[\![N]\!]$, which is intuitive if one looks at most adjectives, from so-called interesectives ('French') to so-called subsectives ('good'). Privatives, instead, do not respect this entailment pattern, as they are precisely defined by (1) above. Partee (2010) famously challenged the view that the class of adjectives is not constrained in terms of entailment pattern. She proposed that privatives are actually subsectives in disguise: they contain no semantic negation. Rather, their composition interacts with some pragmatic principles of interpretation, resulting in an emergent privative entailment pattern (cf. section 2.1).

The second insight, owed to privatives and subsectives jointly, is that adjectives cannot be viewed as denoting simple properties: a French lawyer sits at the intersection between the French and the lawyers; this could be seen as the result of Predicate Modification. This is in contrast with 'good lawyer' and 'fake lawyer', where 'good' and 'fake' do not denote independent sets and thus cannot be merely intersected with [lawyer].

(2) a. $[\![$ French lawyer$]\!] = [\![$French$]\!] \cap [\![$lawyer$]\!]$
b. $[\![$ good lawyer$]\!] \neq [\![$good$]\!] \cap [\![$lawyer$]\!]$
c. $[\![$ fake lawyer$]\!] \neq [\![$fake$]\!] \cap [\![$lawyer$]\!]$

Functional application thus seems the most unproblematic way to compose them. Therefore, the argument goes, we cannot but resign to the idea of capturing adjectives as functions from properties to properties (cf. Montague 1970; Clark 1970; Kamp 1975; Parsons 1970), by "generalizing to the worst case", as put by Partee.

Extant theories of 'fake' disagree about whether its semantic complexity forces us to enrich our semantic framework. Some theorists tried to account for 'fake' by keeping "minimal" lexical meanings for nouns, i.e. by taking 'gun' to denote the set of guns. This is the case of Partee's (2010) seminal work: in her contextualist account, where privatives behave like subsectives on broadened lexical meanings: [fake gun] is a subset of a broadened set of guns, which includes both real and fake guns (cf. 2.1).

Other theorists have it that 'fake' forces us to postulate richer lexical meanings, and

specify compositional operations that derive fine-grained predictions of lexical meanings. Del Pinal (2015, 2018) proposes a bi-partition of lexical items into a categorization-relevant component and a diagnosticity-relevant component. In Dual Content Semantics (DCS), the entries for NPs are subdivided into an E-structure and a C-structure. The E-structure contains the extension, and is thus relevant for set membership, while the C-structure contains information *about* the extension, *viz.* the "core facts" about the category, which are not relevant for the determination of category membership, but are *diagnostic* of that category (cf. the arguments from Putnam 1970). The E-structure newly formed by 'fake', then, includes a negation of the E-structure of its input noun as well as certain dimensions of the C-structure of its input noun: a fake gun is not a gun and has certain properties diagnostic of a gun.

In this paper, I propose, like Del Pinal[1] and unlike Partee, that 'fake' contributes a *bona fide* semantic negation, and thus is not a subsective in disguise.

On the other hand, like Partee and unlike Del Pinal, I propose that we can specify a lexical entry for 'fake' with no need, as far as semantic composition is concerned, for rich lexical semantic structure. Put in other terms, I argue that much of the content of the C-structure that Del Pinal posits for 'fake' is actually *compositionally inert*: it is not due to 'fake' operating on an additional layer of meaning. Rather, it emerges because the meaning of 'fake' embeds a form of similarity talk.

Concretely, I propose that a fake gun is an object that (i) is intended to resemble a gun, but (ii) isn't a gun. If we reduce the semantics of 'fake' partly to the semantics of similarity verbs like 'resemble' or 'seem like', we can import the mechanisms by which similarity verbs get access to the diagnostic properties of a category. And if similarity verbs can be accounted for by a one-dimensional semantics, as they recently have been (Guerrini, 2022), we can dispense completely with the C-structure when accounting for 'fake'.

Of course, that Del Pinal's DCS is not needed for 'fake' does not mean it is not needed *tout court*. But once one takes seriously similarity as a tool to capture adjectives that behave like 'fake', there are good reasons to think that a one-dimensional semantics is enough for more than just 'fake'. For instance, Del Pinal (2015) argues that 'counterfeit' and 'artificial' depart in important ways from 'fake' but, crucially, always paraphrases them with expressions involving similarity. Roughly, he proposes that a counterfeit *X* is not an *X*, intended to look *like* an *X*, and intended to function *like* an *X*. An 'artificial *X*', instead, tells us that an articial *X* isn't an *X* but functions *like* an *X*. It is straightforward to see that, provided we have a one-dimensional analysis of 'like', we can capture these paraphrases one-dimensionally.

In this paper, I focus on 'fake', and thus don't assess whether the paraphrases of 'counterfeit' and 'artificial' mentioned above yield the right truth conditions – though I do find Del Pinal's paraphrases convincing. But the larger point made in this paper is that for those privatives that involve an adequate paraphrase in terms of similarity, a one-dimensional semantics should suffice. Once these more recalcitrant cases are treated one-dimensionally, all privatives can plausibly be treated similarly. Other privatives seem indeed to be more well-behaved: among those originally mentioned by Partee there are simple privatives like the prefix *non-*, which does not seem to contribute anything beyond simple negation. Others are modal privatives such as 'imaginary' or

[1] Especially in Del Pinal (2015); Del Pinal (2018) is more non-committal about this.

'would-be'; and temporal privatives such as 'ex-', 'past', 'future', which all seem to be reducible to well-understood modalities, and thus to one-dimensional lexical entries.[2]

Regardless of whether 'fake' forces a more complex notion of meaning upon us, both Partee's and Del Pinal's theory make inadequate predictions concerning the iteration of 'fake'. Iteration is an excellent testbed for theories of privativity in general, as it makes privative entailment patterns unstable (cf. Jespersen *et al.*, 2017). Both 'non-' and 'fake' are privative, but if applied twice, they yield different entailment patterns.

(3) x is a non-non-gun $\models$ x is a gun (4) x is a fake fake gun $\not\models$ x is a gun

If you are unsure about (4), notice that (5a) qualifies as a gun, but not (5b). Yet both are fake fake guns.

(5) a. FAKE FAKE GUN; GUN.
An airsoft gun made to shoot actual bullets and kill in an airsoft game.
b. FAKE FAKE GUN; NOT A GUN.
A cake that was built to resemble a famous model of a toy gun.

This paper is structured as follows: in section 2, I show in detail how both Partee's theory and DCS, as stated, make respectively insufficient and wrong predictions for the iteration of 'fake'. In 3, I state in detail my one-dimensional proposal, show that it makes the right predictions for iteration of 'fake', and illustrate it plugging a semantics for similarity verbs into the meaning of 'fake'. I also show that it makes right predictions for cases in which 'fake' is stacked with another intensional modifier, taking the example of 'typical'. Finally, in 4 I discuss some aspects of my analysis.

[2]One last kind of phenomenon related to relativity is the case of 'emergent' privativity such as the one one observes in noun-noun compounds like 'stone lion'. 'Emergent' (Guerrini & Mascarenhas, 2019) or 'contingent' (Martin, 2018) privativity contrasts with 'grammatical' or 'functional' privative modification (the one that concerns 'fake' *etc.*) as follows. Grammatical privatives privativize whatever input they take: *fake-gun(x)* $\models$ $\neg$*gun(x)*, *fake-statue(x)* $\models$ $\neg$*statue(x)*. On the other hand, whether contingent privatives privativize their input depends on what the input is: *stone-lion(x)* $\models$ $\neg$*lion(x)*, *stone-statue(x)* $\models$ *statue(x)*. Much more flexibility seems to be involved in emergent privativity: a stone lion may be a statue of a lion or, e.g., a lion that lives in a stoney environment (as in 'mountain lion'). Consequently, for this phenomenon a richer representation of meaning may well be needed.

2 Iteration problems in previous accounts

2.1 Partee

Partee (2010) provided an analysis of privative adjectives like 'fake' and 'counterfeit' and privative constructions like 'stone lion' in terms of pragmatic coercion of the NP. In brief, she proposed that privatives are actually subsective adjectives, but that pragmatic principles expand the denotation of the input noun of privative adjectives. Put in intuitive terms, these principles work as follows: a MODIFIER-HEAD should be a HEAD. For instance, a police car is a car. But sometimes there are no HEADs such that they are a MODIFIER. For instance, there are no lions (HEAD) made of stone (MODIFIER): a stone lion is not a lion. In these cases we should extend the HEAD to include individuals that are a MODIFIER, viz. we should extend the denotation of *lion* to include lion-shaped individuals made of stone, such as statues of a lion. It is important to flag that pragmatic relevance has an important place in this account, to determine *what* specifically the lion-related individuals are.

Partee's account of privative constructions proposes that a principle requires speakers to interpret predicates so that neither their positive nor their negative extensions are empty:

> **Non-vacuity principle** (NVP): try to interpret any predicate so that both its positive and negative extension are non-empty.

Another principle requires speakers to interpret a compound expression as relative to the head.

> **The Head primacy principle** (HPP):
> in a modifier-head structure, the modifier is the modifier and the head the modifee.
> *or, equivalently:*
> In a modifier-head structure, the head is interpreted relative to the context of the whole constituent, and the modifier is interpreted relative to the local context created from the former context by the interpretation of the head (Partee, 2010).

However, some expressions prevent speakers from applying both principles. 'Stone lion', for example, violates the Non-Vacuity Principle if speakers hold on to the head primacy principle: if a stone lion is an actual lion, then it denotes an empty set, since there are no actual lions made of stone. Conversely, if speakers hold on to the Non-Vacuity Principle, the head primacy principle gets violated: to keep a non-empty extension of the denotation, 'stone lion' gets interpreted relative to the modifier, and the denotation of 'lion' is coercively expanded. Partee argues for this latter possibility: the Non-Vacuity Principle is in fact higher-ranked. If there is no reasonable way to obey the Non-Vacuity Principle without shifting the noun beyond its normal bounds, then it may be shifted in such a way as to make the compound predicate obey the Non-Vacuity Principle:

(6) 'Stone lion'

Try to interpret the expression as $\{\text{stone lion}\} \subseteq \{\text{lion}\}$, get an empty extension

⇒NVP fails.

Give up HPP: a stone lion is now a stone. The extension of 'lion' is broadened.

Partee submits that a similar reasoning applies to 'fake gun'. However, there are reasons to think that for 'fake' this account is not explicit enough.[3]

To see this, let us go through the reasoning for the case of 'fake'.

(7) There are no fake [fake guns] in the extension of 'fake gun'.

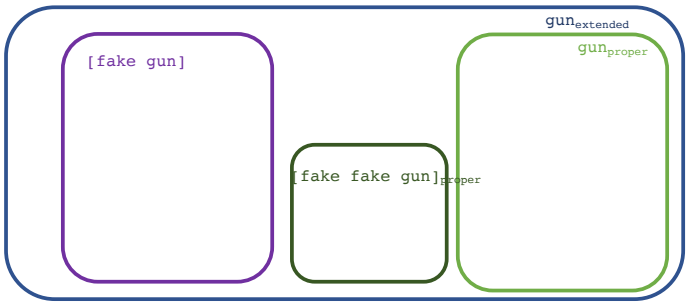

(8) Relax Head Primacy. ⟦ fake gun ⟧ is broadened to include both real and fake [fake guns] (just like in the computation of 'fake gun' the denotation of 'gun'

[3] It is worth mentioning that Del Pinal has already criticised Partee's account of 'fake'. Consider (i). Partee predicts that when the application of a modifier would yield a vacuous extension, NVP would predict (ia) to receive roughly the same reading as (ib). In other words, because there are no fake guns, the set of guns is broadened accordingly: a broadening including toy guns would be especially salient in this context, but isn't observed.

(i) a. I heard some disturbing news. Some terrorists constructed fake guns and planned to use them to attack a halloween party.
b. I heard some disturbing news. Some terrorists constructed fake toy guns and planned to use them to attack a halloween party.

Secondly, consider (iia). Adverbs like 'literally' suspend the application of non-vacuity: a stone lion isn't a statue of a lion made of stone, but a *real* lion made of stone.

(ii) a. Something unbelievable happened at MIT. Scientists discovered a way of making, literally, stone lions and rubber rabbits.
b. Bio-technology is advancing at an astonishing pace. I am convinced that, in the future, we will be able to make, literally, silicon cows.

Because 'fake' and 'stone' are subsectives and both subject to NVP, there should be no difference in their behavior when 'literally' modifies them. In other words, it is predicted that 'literally a fake gun' is interpreted as *fake* $[gun_{literal}]$, not as *fake* $[gun_{extended}]$. This is not what we observe:

(iii) a. Listen to this unbelievable story. Some immoral toy store owner was, literally, selling fake guns at his store.
b. Something amazing happened at MIT. Some engineer managed to make, literally, a fake gun.

was extended to include both real and fake guns).

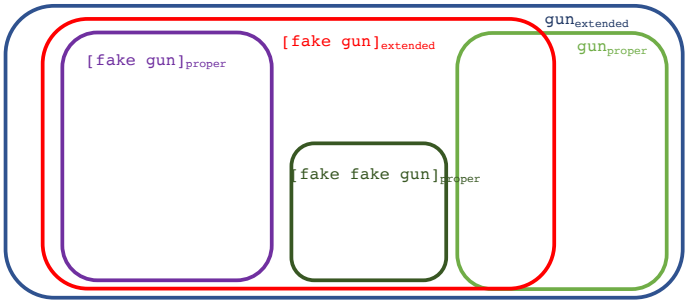

Figure 1: A representation of the step in (8): relax head primacy and broaden the extension of fake-guns-proper to fake-guns-extended, to include both real and fake [fake guns].

But why precisely is it that we have to broaden the set of guns to begin with? Given that we have no specific lexical entry for 'fake', this is unclear. What, in the meaning of 'fake', makes that there are no fake guns? And relatedly, when we broaden the set of guns, what non-guns do we add?

One way to see that Partee's account is at least incomplete is by extending Partee's reasoning for cases of iterated application of 'fake'.

We know from section **??** (but see 3.2 for a detailed discussion of the judgments) that a fake fake gun may or may not be a gun.

This prediction is not made by Partee's account, because we don't know precisely how the denotation of 'fake gun' is broadened to include fake fake guns. That is, we could be in either of the three scenarios in Figures 2, 3, or 4 below.[4]

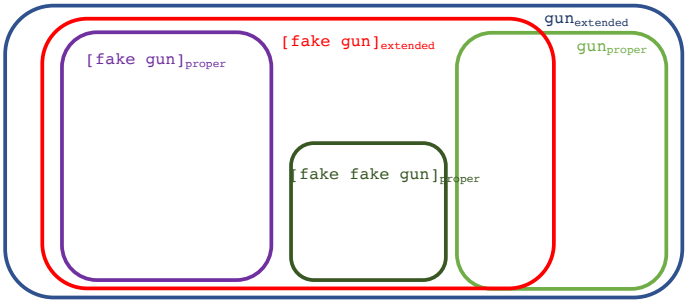

Figure 2: *A scenario in which all fake fake guns are non-guns.*

[4] I leave aside the less important prediction of whether fake fake guns should be in the denotation of [gun]$_{extended}$, which would be Partee theory-internal.

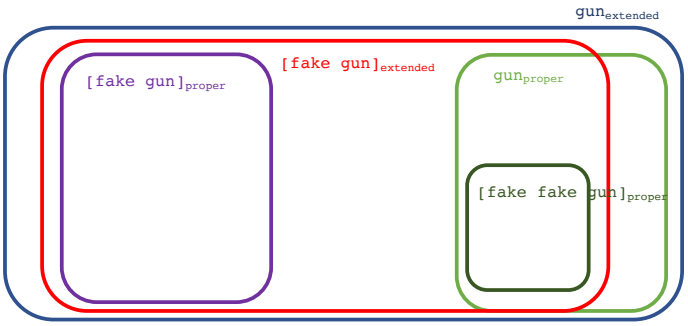

Figure 3: *A scenario in which all fake fake guns are guns-proper.*

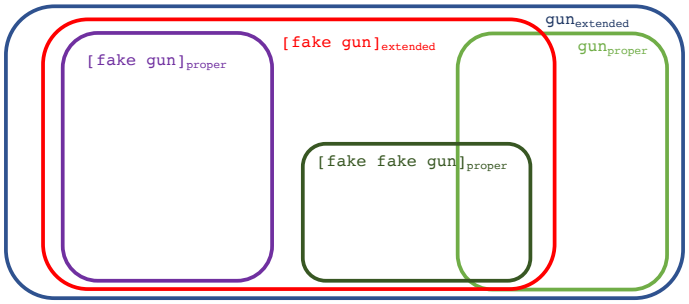

Figure 4: *A scenario in which a fake fake gun may or may not be a gun-proper.*

2.2 Del Pinal

A more recent account that adopted the third solution is Del Pinal's (2015; 2018) Dual Content Semantics (DCS). DCS is an account of privatives and of lexical modulation that posits that lexical meanings are constituted by a fixed set of *qualia* that constrain flexibility. The choice of such a regimented internal structure addresses specific problems faced by contextualist accounts. Contextualists try to account for semantic flexibility by radically liberalizing and loosening the compositional operations of language (Pagin & Pelletier 2007; Recanati 2010; Szabó 2010; Lasersohn 2012). Critics have pointed out that this account overgenerates possible readings that are actually not observed (Asher, 2011; Feinmann, 2020). For instance, we cannot utter sentence (9b) in order to convey the meaning of sentence (9a), although the meaning in (9a) should be the most relevant reading of (9b) (Asher, 2011).

(9) a. Mary stopped eating the apple.
b. Mary stopped the apple.

In DCS, meaning is determined along two lines: the E-structure of nouns determines

their extension, while the C-structure incorporates the related 'core facts', a set of beliefs about the extension, in the form of *qualia* similar to Pustejovsky's. DCS proceeds by addressing two opposing constraints: on the one hand, it seeks a compositional solution for puzzles related to lexical flexibility by recruiting the conceptual structure associated with nouns. On the other hand, it does so while seeking to not free-up the semantics so much as to incur in the overgeneration problems that have affected free modulation in the contextualist framework. The result is a set of non-atomic lexical representations and of combinatorial operations that compositionally derive cases of lexical flexibility without generating unobserved meanings.

In this section, I first present the theory. Then I raise some serious issues with this specific implementation of DCS. DCS makes wrong predictions for multiple applications of modifiers such as 'fake' and 'typical', and there is no straightforward fix for it to make right predictions. I discuss some examples and show that the solutions to this problem proposed in the literature are not fully general.

Consider a lexical entry for *gun* in this framework:

(10) $[\![\text{ gun }]\!]^{c}_{M} =$

E-structure:

$\lambda x.\text{GUN}(x)$

C-structure:

C: $\lambda x.$ PARTS_GUN(x)
P: $\lambda x.$ PERCEPTUAL_GUN(x)
T: $\lambda x.$ GEN $e\big[\text{SHOOTING}(e) \wedge$
$\text{INSTRUMENT}(e,x)\big]$
A: $\lambda x.\exists e_1\big[\text{MAKING } (e_1) \wedge \text{ GOAL}(e_1, \text{ GEN } e(\text{ SHOOTING}(e) \wedge$ INSTRUMENT $(e,x))\big]$

The C-structure encodes information about how entities in the class are typically perceived ('p' for 'perceptual'), what matter they are made of ('c' for 'constitutive'), how they came to being or for what purpose they were created ('a' for 'agentive'), and what their intended and typical function is, if any ('t' for 'telic'). Importantly, whether or not an individual has these features does not determine whether it falls under the concept; that is only determined by the E-structure (cf. Pustejovsky & Bouillon 1995; Moravcsik 1998; Del Pinal 2015).

Other information that is included in the C-structure includes (i) the weight of dimensions as a function of their importance, and (ii) dependency relations between the dimensions and the correspondent relative centrality.[5]

The central point of DCS is that meanings are much richer than traditionally thought. Compositionally, the fundamental insight is that certain expressions carry over contents of the C-structure to the E-structure. To do this, two types of tools are needed: (i) functions that take full meanings and return the value of a particular dimension, and (ii) functions that take full meanings and return combinations of the E-structure with one or more dimensions of the C-structure.

[5] A feature d is central, Del Pinal proposes, in a concept if other dimensions depend on d more than d depends on the other features.

- **Dimensional operators**: partial functions from the full meaning of a term to its fine-grained C-structure denotations, as for instance Q_T, which takes a lexical item as input and returns the value of the TELIC *quale*.
- **Core enrichment operators**: partial functions from full meanings to combinations of their E-structure and C-structure. For instance, a core enrichment operator A may return the conjunction of the E-structure and the value of the AGENTIVE *quale* in the C-structure. The core enrichment operator E takes a full expression and returns only its E-structure.

Some expressions only add descriptive content to the E- and C-structure of their argument. Others upload parts of the C-structure to the E-structure. An example of the former kind is 'steel':

(11) $[\![\text{steel}]\!]^c_M =$

E-structure:

$\lambda P.\lambda x.P(x) \wedge \text{STEEL}(x)$

C-structure:

C: $\lambda P.\lambda x.P(x) \wedge \text{STEEL}(x)$
P: $\lambda P.\lambda x.P(x) \wedge \text{STEEL\_PERCEPTUAL}(x)$
T: $\lambda P.P$
A: $\lambda P.P$

Given a definition of Functional Application that applies "per dimension", *steel* yields the following when applied to *gun*:

(12) $[\![\text{steel gun}]\!]^c_M =$

E-structure:

$\lambda x.\text{GUN}(x) \wedge \text{STEEL}(x)$

C-structure:

C: $\lambda x.\text{PARTS\_GUN}(x) \wedge \text{STEEL}(x)$
P: $\lambda x.\text{PERCEPTUAL\_GUN}(x) \wedge \text{STEEL\_PERCEPTUAL}(x)$
T: $\lambda x.\text{GEN}e\left[\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e,x)\right]$
A: $\lambda x.\exists e_1\left[\text{MAKING}(e_1) \wedge \text{GOAL}(e_1, \text{GEN}e(\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e,x))\right]$

The latter kind of modifiers use the C-structure of their argument. This is the case of 'fake' and 'typical'. 'typical' uploads a variable portion of the content of the C-structure of its argument to the E-structure of the output. Call T a function that takes as input an individual x and a full meaning tuple $\mathbf{G}$ and returns the cardinality of the set of $\mathbf{G}$-attributes whose $\mathbf{G}$-value holds of x. Then in DCS, a typical $\mathbf{G}$ is something that has a big enough number of dimensions of the C-structure of $\mathbf{G}$:

(13) $[\![\text{typical}]\!]^c_M =$

E-structure:

$\lambda \mathbf{G}.\lambda x.E(\mathbf{G})(x) \wedge T(\mathbf{G},x) > s$

C-structure:

C: $\lambda \mathbf{G}.\lambda x.Q_C(\mathbf{G})(x)$

P: $\lambda\mathbf{G}.\lambda x.Q_P(\mathbf{G})(x)$
T: $\lambda\mathbf{G}.\lambda x.Q_T(\mathbf{G})(x)$
A: $\lambda\mathbf{G}.\lambda x.Q_A(\mathbf{G})(x)$

Del Pinal suggests that because DCS includes dimension weights and dependency relations, T can easily be made to count only the most salient or central features. Moreover, the function T accounts for the fact that 'typical' has all the features of a gradable adjective: we can say that an object is more of a typical gun than another object, or that an object is a very typical gun. Del Pinal also observes that while dimensions of the C-structure are not relevant to determine category membership, they are important for typicality. Concretely, if someone says that Alex is a lion, we cannot object to this judgment by saying that it doesn't have a mane. By contrast, if someone says that Alex is a typical lion, it seems befitting to protest that he does not even have a mane.

(14) $[\![\text{typical lion}]\!]^c_M =$

E-structure:
$T([\![\text{ lion}]\!]^c_M, x) > s$
C-structure:
C: $\lambda x.Q_C([\![\text{ lion}]\!]^c_M)(x)$
P: $\lambda x.Q_P([\![\text{ lion}]\!]^c_M)(x)$
T: $\lambda x.Q_T([\![\text{ lion}]\!]^c_M)(x)$
A: $\lambda x.Q_A([\![\text{ lion}]\!]^c_M)(x)$

Besides uploading part of the content of the C-structure of its argument to the E-structure like 'typical' does, 'fake' negates the E-structure of its argument:

(15) $[\![\text{fake}]\!]^c_M =$

E-structure:
$\lambda\mathbf{G}.\lambda x.\neg E(\mathbf{G}(x)\wedge$
$\neg Q_A(\mathbf{G})(x)\wedge$
$\exists e_2\big[\text{MAKING}(e_2)\wedge\text{GOAL}\big(e_2, Q_P(\mathbf{G})(x)\big)\big]$
C-structure:
C: $\lambda\mathbf{G}.\lambda x.Q_C(\mathbf{G})(x)$
P: $\lambda\mathbf{G}.\lambda x.Q_P(\mathbf{G})(x)$
T: $\lambda\mathbf{G}.\lambda x.\neg Q_T(\mathbf{G})(x)\wedge Q_P(\mathbf{G})(x)$
A: $\lambda\mathbf{G}.\lambda x.\exists e_2\big[\text{MAKING}(e_2)\wedge\text{GOAL}\big(e_2, Q_P(\mathbf{G})(x)\big)\big]$

What happens if we apply this entry to the entry for 'gun'? In words, the E-structure of $[\![\text{fake gun}]\!]^c_M$ is satisfied by entities that:

- are not guns
- were not made to be guns
- were made to have the perceptual features of guns.

The C-structure gives more fine-grained information about what fake guns *typically* are: their *telos*, what corresponds to the TELIC dimension, is to not serve as a gun, but as something that looks like a gun:

(16) $[\![\text{fake gun}]\!]^c_M =$

E-structure:

$\lambda x.\neg E\big([\![\text{ gun }]\!]^c_M\big)(x)\wedge$
$\neg Q_A([\![\text{ gun }]\!]^c_M)(x)\wedge$
$\exists e_2\big[\text{MAKING}(e_2)\wedge\text{GOAL}\big(e_2, Q_P([\![\textit{ gun }]\!]^c_M)(x)\big)\big]$

C-structure:

C: $\lambda x.Q_C([\![\textit{ gun }]\!]^c_M)(x)$
P: $\lambda x.Q_P([\![\textit{ gun }]\!]^c_M)(x)$
T: $\lambda x.\lambda x.\neg Q_T([\![\textit{ gun }]\!]^c_M)(x)\wedge Q_P([\![\textit{ gun }]\!]^c_M)(x))$
A: $\lambda x.\exists e_2\big[\text{MAKING}(e_2)\wedge\text{GOAL}\big(e_2, Q_P([\![\textit{ gun }]\!]^c_M)(x)\big)\big]$

2.2.1 *Fake fake*

Martin (2018) pointed out that if we apply 'fake' twice, we get a contradiction in the E-structure. Let us look at the case of 'fake fake gun'.

The E-structure of 'fake gun' reads as follows. A fake gun

- is not a gun
- does not have the origin of a gun
- does not have the purpose of a gun
- is an object x that has a building event that had as a goal that PERCEPTUAL_GUN(x), i.e. that x look like an gun.

Now, applying 'fake' to (16) again, we get a contradiction.

(17) $[\![\text{fake fake gun}]\!]^c_M =$

E-structure:

$\lambda x.\neg E([\![\text{fake gun}]\!]^c_M)(x)\wedge$
$\neg Q_A([\![\text{ fake gun}]\!]^c_M)(x)\wedge$
$\exists e_3\big[\text{MAKING}(e_3)\wedge\text{GOAL}\big(e_3, Q_P([\![\text{ fake gun}]\!]^c_M)(x)\big)\big]$

C-structure:

C: $\lambda x.Q_C([\![\text{ fake gun }]\!]^c_M)(x)$
P: $\lambda x.Q_P([\![\text{ fake gun }]\!]^c_M)(x)$
T: $\lambda x.\neg Q_T([\![\text{ fake gun }]\!]^c_M)(x)$
A: $\lambda x.\exists e_3\big[\text{MAKING}(e_3)\wedge\text{GOAL}\big(e_3, Q_P([\![\text{ fake gun }]\!]^c_M)(x)\big)\big]$

After some calculations, we get something more legible:

(18) $[\![\text{fake fake gun}]\!]^c_M =$

E-structure:

$\lambda x.\neg E([\![\text{ fake gun}]\!]^c_M)(x)\wedge$
$\neg\exists e_2\big[\text{MAKING}(e_2)\wedge\text{GOAL}\big(e_2, Q_P([\![\text{gun}]\!]^c_M)\big)\big]\wedge$
$\exists e_3\big[\text{MAKING}(e_3)\wedge\text{GOAL}\big(e_3, Q_P([\![\text{gun}]\!]^c_M)(x)\big)\big]$

C-structure:

C: $\lambda x.Q_C([\![\text{ gun }]\!]^c_M)(x)$

P: $\lambda x.Q_P([\![\text{ gun }]\!]^c_M)(x)$
T: $\lambda x.\neg Q_T([\![\text{ fake gun }]\!]^c_M)(x)$
A: $\lambda x.\exists e_3\left[\text{MAKING}(e_3) \wedge \text{GOAL}\left(e_3, Q_P([\![\text{ fake gun }]\!]^c_M)(x)\right)\right]$

In other words, there both was and was not an event in which x was made to look like a gun. Why do we get this paradoxical result? Three parts of its meaning interact to yield this undesired contradiction:

1. 'Fake' negates the AGENTIVE of its input and uploads it to the E-structure of the newly formed expression: a fake gun does not have the origin of a gun.
2. It states that there was a making event that had the goal that the denoted object have the formal *quale* of the input. For instance, a fake gun had a making event that had the goal of it having the appearance (formal *quale*) of a gun.
3. 'Fake' passess through the FORMAL and CONSTITUTIVE *qualia* of the C-structure of its input. In other words, a fake gun should have the shape and the material of a gun.

Because of 2. and 3., the second application of 'fake' yields that there was a making event with the goal of the object looking like a gun. But because of 1., the second application of 'fake' negates the AGENTIVE of 'fake gun', which is that there was an event that had the goal of the object looking like a gun.

Martin (2018) proposes a fix for this problem. He proposes that a fake X, instead of having the formal *quale* of an X, has the formal *quale* of *something that has the* TELIC *quale of an* X. For instance, a fake gun has the appearance of something that has the function of a gun. Then, a fake fake gun has the shape of something that has the telos of a fake gun, i.e. it has the shape of something that has the goal of having the shape of a gun.

(19) $[\![\text{fake fake gun}]\!]^c_M =$

E-structure:

$\lambda x.\neg Q_E([\![\text{fake gun}]\!]^c_M)(x) \wedge$
$\neg\exists e_2\left[\text{MAKING}(e_e) \wedge \text{GOAL}\left(e_2, \text{PERCEPTUAL-Q}_T([\![\text{gun}]\!]^c_M)(x)\right)\right] \wedge$
$\exists e_3\left[\text{MAKING}(e_3) \wedge \text{GOAL}\left(e_3, \text{PERCEPTUAL-Q}_T([\![\text{fake gun}]\!]^c_M)(x)\right)\right]$

C-structure:

C: CONSTITUTIVE$([\![\text{gun}]\!])$
P: PERCEPTUAL-$Q_T([\![\text{fake gun}]\!]^c_M)$
T: $\neg Q_T([\![\text{fake gun}]\!]^c_M)$
A: $\lambda x.\exists e_3\left[\text{MAKING}(e_3) \wedge \text{GOAL}\left(e_3, \text{PERCEPTUAL-Q}_T([\![\text{fake gun}]\!]^c_M)(x)\right)\right]$

This fix addresses the problem raised by the contradictions predicted by Del Pinal's entry for 'fake' in recursive application. However, the problem is deeper.

The introduction of PERCEPTUAL-Q_T gives us something that has the appearance of having the telos of Q. Notice that the content of specific *qualia* is a pure extension. Then, although not defined as a function, *de facto* PERCEPTUAL works as a function. It takes as an input a set of individuals (e.g., the entities that have the telos of a gun) and returns all the individuals that have the perceptual appearance of that set of individuals.

Ultimately, it presupposes a notion of similarity. But if we presuppose a notion of similarity, why postulate a second layer of meaning to begin with?

The general lesson of all this is that although we do have to know what a gun is like to understand the meaning of 'fake gun', what a fake gun is is extremely open-ended. For instance, typically fake guns should not be made of the same material as guns, while in Martin's account the CONSTITUTIVE *quale* is the same for guns, fake guns, and fake fake guns. Arguably, guns are typically made of steel, while fake guns are typically made of plastic. This comes out pretty clearly if we look at a similarity statement like (20):

(20) With respect to its material, this object looks like a fake gun.

Certainly a speaker uttering (20) does not mean by this that the object is made of perfect steel, quite the opposite. This explains why (21) does not feel contradictory.

(21) With respect to its material, this thing looks like a fake gun, and not like a gun.

To account for these facts, we would need to postulate a second operator, CONSTITUTIVE-Q_T, which takes as an input a dual concept and outputs the material that makes up things that have the goal of resembling the input concept. For instance, CONSTITUTIVE-Q_T($[\![$ gun $]\!]$) would output the material of which are made things that have the goal of having the shape of a gun.

2.2.2 *Fake typical*

DCS also fails to deliver the right predictions for the stacking of multiple intensional adjectives, both in its original form from Del Pinal (2018) and in its revised form from Martin (2018). Take the example of 'fake' embedding another intensional expression like 'typical'. Recall that, in DCS, 'typical' uploads a variable portion of C-structure to the E-structure.

(22) $[\![\text{typical}]\!]^c_M =$

E-structure:

$\lambda \mathbf{G}.\lambda x.E(\mathbf{G})(x) \wedge T(\mathbf{G},x) > s$

Again, the C-structure is passed along from the input expression to the output expression. In DCS without Martin's fix, we get that the E-structure of a fake typical gun:

(i) does not make true a sufficient number of prototypical gun traits.
(ii) crucially, **there is no making event that had the goal that the denoted object function like a typical gun, thus like a gun.** This is at the root at the problem, as we are about to see.
(iii) there is a making event with the goal that it resemble a gun.

Here's a computation of the E-structure of 'fake typical gun' following Del Pinal's

lexical entries. For brevity, I leave out the C-structure (which is again the same as the C-structure of 'gun').

(23) $[\![\text{ fake typical gun }]\!]^{c}_{M} =$

E-structure:
(i) $\lambda x.\neg|T(\text{GUN}_c, x) > s\wedge$
(ii) $\neg\exists e_1[\text{MAKING}(e_1)\wedge\text{GOAL}(e_1, \text{GEN}_e(\text{SHOOTING}(e)\wedge\text{INSTRUMENT}(e,x))]\wedge$
(iii) $\exists e_2[\text{MAKING}(e_2)\wedge\text{GOAL}(e_2, \text{PERCEPTUAL\_GUN}(x))]$

To restate these three points in simpler words, an object that should qualify as a fake typical gun:

(i) does not have a certain number of traits typical of guns
(ii) is not built to function like a gun;
(iii) is intended to be perceived as a gun.

This makes undesired predictions. Take a gun that was built to function like a gun, i.e. to shoot. It was intentionally built in such a way that from a certain perspective it looks like an ordinary, unremarkable Colt, but from another angle it becomes clear it has all sorts of exotic features, e.g. it has ten barrels. The account incorrectly predicts this object to not qualify as a fake typical gun, because it is built to function like a gun.

This example illustrates how a fake typical gun should not merely be made to resemble a gun. It should be made to resemble a *typical gun*.

The case of 'fake typical' also demonstrates that Martin's (2018) fix isn't merely explanatorily inadequate, but also not fully general. Taking Del Pinal's entry for 'typical', Martin's analysis of 'fake' tells us that a fake typical gun was made to look like it has the telos of a gun, since the telos of a typical gun is the same as that of a gun. This is clearly wrong. Take an object that is intended to look like it shoots, and moreover is intended to look like it has all sorts of exotic features, e.g. it has ten barrels. This object is incorrectly predicted to qualify as a fake typical gun.

What we really want is that a fake typical gun have those traits that are diagnostic of *typical guns*. These traits may not be the same as those that are diagnostic of guns. This will crucially depend on the extension: typical *N*s may all happen to have a feature *F* that is not very typical of *N*s. An individual that simulates precisely *F* is thus a fake typical N, but not a fake N.[6]

One may wonder whether there isn't a revision of the semantics of 'typical' that could make Martin's solution for 'fake' work in cases of concatenation of 'fake' and 'typical'. To do this, we would have to resort to a solution like Martin's (2018) for 'fake': for instance, postulating that the PERCEPTUAL *quale* of 'typical' features a function, say PERCEPTUAL-TYPICAL, which takes as an input a dual concept and returns the typical appearance of *typical* exemplars of that category. But again, we are left with

[6]For instance, suppose that by chance all typical speakers of French are concentrated in a given city, say Rouen. Suppose further that a small linguistic innovation occurs, and that French speakers from Rouen start pronouncing a single phoneme *P* in a slightly different way that isn't found in any other dialect of French. If I am not a typical speaker of French, but pretend like I can naturally produce *P*, I will only qualify as a fake typical French speaker, and not like a fake French speaker. See Guerrini (2021) for an in-depth discussion of such examples.

very little predictions, and with a rapidly increasing number of black-box functions: PERCEPTUAL-Q_T, CONSTITUTIVE-Q_T, PERCEPTUAL-TYPICAL...

This discussion shows, I think, that for DCS to adequately treat iterated intensional modifiers, we need the C-structure to be computed based on the E-structure.[7] But since the C-structure cannot be computed compositionally, why not state a compositional semantics for 'fake' that simply does not refer to the C-structure? In the following, I show that referring to similarity in the meaning of 'fake' gets us the right predictions for iterated application of 'fake'. I do not offer a full account for similarity talk, but illustrate what specific predictions we would get by importing one specific account of similarity talk, Guerrini (2021). I argue that these are reasonable enough predictions to make it a good candidate to integrate my incomplete theory of 'fake'.

3 The proposal

3.1 My proposal: a semantics based on similarity

I propose that 'fake' conjoins two operations:

- It states that an entity intends or is intended to look like an instance of a category and
- it states that that entity is not a part of that category.

This is implemented below:

(24) $[\![\text{fake}]\!] = \lambda P.\lambda x.\text{INTENDED}\big(x, \text{SEEM-LIKE}(x,P)\big) \wedge \neg P(x)$

Here INTENDED captures, informally, the fact that there was an action intentionally performed on x whose goal was that x seem like a P. One way to implement this would be the following:

(25) a. $\text{INTENDED}(x,P) = \exists e.\text{ACTION}(e,x) \wedge \text{GOAL}\big(e, P(x)\big)$
b. $\text{INTENDED}\big(x, \text{SEEM-LIKE}(x,gun)\big) = \exists e.\text{ACTION}(e,x) \wedge \text{GOAL}\big(e, \text{SEEM-LIKE}(x,gun)\big)$

In words, there is an event that is an action performed on x whose goal is that x seem like a P.

This entry is similar to a part of Del Pinal's meaning for 'fake':

(26) $\exists e_2\big[\text{MAKING}(e_2) \wedge \text{GOAL}\big(e_2, Q_P([\![\ gun\ ]\!]^c_M)(x)\big)\big]$

However, Del Pinal's version only applies to artefacts. Take for instance the case of a fake lawyer: a person who pretends to be a lawyer without being one.[8] In this case, there is no making event. Rather, there is an action that the person performs on

[7] Guerrini (2021), for instance, proposes to fix DCS as follows: any time the E-structure is modified, there is some C-structure byproduct, and new diagnostic information is created on the fly.

[8] Of course, there is a reading of the expression under which we can call John a fake lawyer because someone else is causing John to act

themselves to resemble a lawyer. This is adequately captured by the specification of INTENDED in (25a).

Moreover, this specification of INTENDED seems to make sense of speakers' intuition that there must be some intentionality behind fakeness. Take a random configuration of atoms on Mars that looks exactly like a gun, but is not a gun. Suppose we are sure that this is just a random circmustance, i.e. that nobody molded this configuration of atoms. We cannot call this a fake gun. This requirement of intentionality is encoded by GOAL. I leave GOAL underspecified here, but an account of goals in terms of bouletic modality (along the lines of, say, Heim 1992) should be straightforward.

3.2 Testing iteration in the present account

A second application of 'fake' will give us an object that is not a fake gun, and is intended to look like a fake gun.

(27) $[\![\text{fake fake gun}]\!] = \lambda x.$
i. $\neg\text{fake-gun})(x) \wedge$
ii. $\text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{fake-gun}))$

If we develop (28), we get:

(28) $[\![\text{fake fake gun}]\!] = \lambda x.$
i. $\neg\Big(\neg\text{gun}(x) \wedge \text{INTENDED}\big(x, \text{SEEM-LIKE}(x, \text{gun})\big)\Big) \wedge$
ii. $\text{INTENDED}\big(x, \text{SEEM-LIKE}(x, \text{fake-gun})\big)$

which means:

(29) $[\![\text{fake fake gun}]\!] = \lambda x.$
i. $\text{gun}(x) \vee \neg\text{INTENDED}\big(x, \text{SEEM-LIKE}(x, \text{gun})\big) \wedge$
ii. $\text{INTENDED}\big(x, \text{SEEM-LIKE}(x, \text{fake-gun})\big)$

In words, something is a fake fake gun iff

- it is a gun, or wasn't made to look like a gun, or both
- and, moreover, it was made to look like a fake gun

Consequently, the account predicts that three classes of objects should count as fake fake guns:

	$\mathscr{G}_1$	$\mathscr{G}_2$	$\mathscr{G}_3$
$\text{gun}(x)$	1	0	1
$\neg\text{INTENDED}\big(x, \text{SEEM-LIKE}(x, \text{gun})\big)$	0	1	1
$\text{INTENDED}\big(x, \text{SEEM-LIKE}(x, \text{fake gun})\big)$	1	1	1

An object in the $\mathscr{G}_1$ class is:

- a gun
- intended to look like a gun
- intended to look like a fake gun

In other words, the object is intended to look like a gun precisely because *successful* fake guns look like guns. However, it turns out it actually is a gun, thus not a fake gun.

For instance, imagine an airsoft gun that was made to shoot actual bullets and murder someone during a military simulation game played with airsoft weapons.

- This gun is intended to look like a gun, because airsoft guns are almost indistinguishable from real guns.
- Moreover, it was made to look like a fake gun: airsoft guns are fake guns, and this object was made to look like an airsoft gun.
- Finally, it is a gun.

Consequently, the account correctly predicts this to be a fake fake gun.

An object in the $\mathscr{G}_2$ class is:

- not a gun;
- not intended to look like a gun;
- intended to look like a fake gun;

Imagine a visual illusion that appears to have the shape of a colorful toy gun from a certain angle, but when contemplated from a different perspective it becomes clear that it is only a bunch of superposed parts far apart from each other. It is:

- not a gun;
- not intended to look like a gun, as it is intended to look like something that is not a gun, a toy gun;
- intended to look like a fake gun, a toy gun;

A further, more concrete example of this class of individuals is that of a fake fake secret agent.

(30) a. *Context*: The CIA has placed an agent inside the KGB, a precious source of information. The Russians come to know this, and are on the hunt. John, an old CIA sleeper agent decides to sacrifice himself to save the American spy. He decides to start behaving like an American who tries to infiltrate the KGB: he gives excessive detail on his Russian origins, and when manipulating weapons he is very theatrical about having had a Russian weapon training.
b. John is a fake fake KGB agent.

The account correctly predicts that John may qualify as a fake fake KGB agent, as he:

- is not a KGB agent;
- is not intending to look like a KGB agent (otherwise he would be more discreet);
- intends to look like a fake KGB agent.

An object in the $\mathscr{G}_3$ class is:

- a gun;
- not intended to look like a gun;
- intended to look like a fake gun;

Imagine terrorists who want to board a flight with a gun, and to this effect produce a real gun that looks like a toy gun. The account correctly predicts this object to qualify as a fake fake gun, as it is a gun, but not intended to look like one - it is intended to resemble a toy gun. Moreover, it is intended to look like something that qualifies as a fake gun, a toy gun.

3.3 *Fake typical N*

The goal of this subsection is to bring out some desiderata for a semantics of 'typical' and show that, when these desiderata are met, the analysis proposed in this paper for 'fake' makes the right predictions for the concatenation of 'fake' and 'typical', unlike DCS (Del Pinal, 2018) and minimal modifications to it like Martin (2018).

To do this, we will work with an underspecified entry for 'typical' which only refers to two conditions an individual has to meet to be a typical *N*:

- it must be a member of *N*
- it must meet some condition **C** that makes it a typical instance *N*

To see this, I briefly review some data concerning 'typical'. A first, simple insight is that 'typical' is a subsective adjective: simply having many properties of a category would not be enough to be a typical instance of that category. To see this, compare 'typical' to 'look-alike':

(31) - That is a typical lion.
- No, it isn't a lion!

(32) - That is a look-alike lion.
- *No, it isn't a lion!

But what semantic material does 'typical' add? Del Pinal (2018) shows, in his discussion of 'typical', that 'typical' makes at-issue some associated properties that are not usually at issue. To see this, consider (33):

(33) a. - That's a lion.
- No, it's not. It doesn't even have a mane.
b. - That's a typical lion.
OK - No, it's not. It doesn't even have a mane.

While having a mane is not relevant to determine category membership, it is relevant to determine whether something is a typical instance of a category. This justifies the presence of our **C** condition on typicality. We can thus move on to test our entry for 'fake' by applying it to 'typical gun'.

Applying 'fake' on top of 'typical' will negate the conjunction present in the truth conditions of *typical gun*, and add a clause:

(34) $[\![\text{fake typical gun}]\!] =$

(i) $\lambda x.\neg\Big(\text{gun})(x) \wedge C(\text{gun})(x)\Big) \wedge$

(ii) $\textsc{intended}\big(x, \textsc{seem-like}(x, \text{typical gun})\big)$

This means:

(35) $[\![\text{fake typical gun}]\!] =$

(i) $\lambda x.\big(\neg\text{gun}(x) \vee \neg C(\text{gun})(x)\big) \wedge$

(ii) $\textsc{intended}\big(x, \textsc{seem-like}(x, \text{typical gun})\big)$

Then, three classes of objects should qualify as fake typical guns:

	$\mathscr{G}_1$	$\mathscr{G}_2$	$\mathscr{G}_3$
$\neg\text{gun}(x)$	1	0	1
$\neg\mathbf{C}(\text{gun})(x)$	0	1	1
$\textsc{intended}\big(x, \textsc{seem-like}(x, \text{typical gun})\big)$	1	1	1

An object in the $\mathscr{G}_1$ class is:

- not a gun
- has a certain number of traits typical of guns
- is intended to look like a typical gun

For instance, a replica of a Colt (*qua* typical gun) falls within this class, and is correctly predicted to qualify as a fake typical gun.

An object in the $\mathscr{G}_2$ class is:

- a gun
- does not have a certain number of traits typical of guns
- is intended to look like a typical gun

Take a gun that was built in such a way that from a certain perspective it looks like an ordinary, unremarkable Colt, but from another angle it becomes clear it has all sorts of exotic features, e.g. it has ten barrels. This is a gun, it does not have enough typical traits of a gun, and is intended to look like a typical gun, and the account correctly predicts it to qualify as a fake typical gun.

An object in the $\mathscr{G}_3$ class is:

- not a gun
- does not have a certain number of traits typical of guns
- is intended to look like a typical gun

take a toy gun that, just like above, was built in such a way that from a certain perspective it looks like an ordinary, unremarkable Colt, but from another angle it becomes clear it has all sorts of exotic features, e.g. it has ten barrels. The account correctly predicts this object to qualify as a fake typical gun.

4 Discussion

In the remainder of this paper, I address some aspects of the analysis just presented. In sections 4.1 and 4.2, I address two possible objections to the 'SEEM-LIKE' part of my lexical entry for 'fake'. More specifically:

- Nothing rules out that the black box SEEM-LIKE is only accountable within a bi-dimensional semantics. This would make my semantics for 'fake' not obviously simpler than Del Pinal's. In 4.1 I illustrate with Guerrini's (2022) account of similarity constructions, one possible way of spelling out SEEM-LIKE. To the extent that this analysis of similarity is one-dimensional, my analysis of 'fake' is one-dimensional, too.
- In 4.2 I discard a possibly simpler theory of 'fake' which reduces it to well-understood attitudes, namely one that posits that a fake *N* is intended so that some judge believe it is a *N*, but it is not a *N*.

In sections 4.3 and 4.4, I address some possible objections to the '$\neg P(x)$' part of my lexical entry for 'fake'. More specifically:

- In 4.3 I discuss why one can use 'fake watch' to refer to a counterfeit Rolex. Because a counterfeit Rolex is still a watch, at least *prima facie* 'fake' is not negating its argument. I show that this is not a problem for my analyisis once one takes into account some facts about the syntax of adjectival modification.
- In 4.4, I show why the experimental arguments that Martin has recently brought to argue that 'fake' contributes no semantic negation do not actually support this hypothesis. I also show that sentences like 'That gun is fake' do not undermine my claim that 'fake' involves semantic negation.

4.1 One-dimensional 'fake' with one-dimensional similarity

Guerrini (2022) argues that constructions such as the one involved in sentence (36a) are best captured in terms of property sharing, i.e. as in (36b)

(36) a. John is like Mary.
b. John shares relevant properties with Mary.

This is motivated by the fact that similarity talk seems to be sensitive to the type of properties shared by two entities:

(37) With respect to personality, John is like Mary.

We can think of these similarity respects as sets of properties for which the two individuals have the same value. In figure 5, such sets are represented in a simplified way as partitions over the domain of individuals.

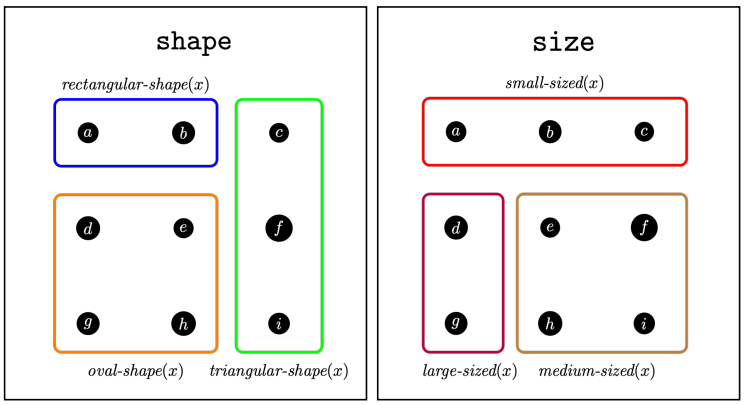

Figure 5: "*With respect to shape and size, object* a *is like object* b."

Guerrini's proposal, roughly, is that x is like y iff for a relevant set of properties, x has the same value as y.

(38) $[\![\text{like}]\!]^{w,\mathbf{D}} = \lambda y.\lambda x.\forall P_{s,\langle e,t\rangle} \in \mathbf{D}_{\langle s,\langle e,t\rangle\rangle,t}.P(x)(w) = P(y)(w)$

Notice again that similarity is contingent on the extension: if lawyers happen to all wear a green suit, then if John wears a green suit, we can say of John that with respect to clothing, he resembles a lawyer.

Notice, too, that expressions such as 'like Mary' denote properties, i.e. the set of individuals that for relevant properties shares the same value with Mary. Then, one can embed this property in other copular constructions, and notably in appearance verbs such as 'seem', 'look', 'sound'.

Rudolph (2019) proposes that sentences like (39a) can be captured as in (39b).

(39) a. 'Bob looks French'
b. $\forall w' \in V(j,w).\ French(bob)$
c. 'at all worlds compatible with the visual experience of the judge j at w, $French(bob)$'

Extending Rudolph's semantics of 'look $P_{e,t}$', then, a sentence like 'Bob looks like Carl' is analyzed as in (40).

(40) $\forall w' \in V(j,w).\forall P \in \mathbf{D}.P(x)(w') = P(y)(w')$

This allows us to capture the subjectivity of 'look': 'To John, Bob looks like Carl' is captured by simply feeding a different judge argument to V to get the worlds w' in $V(john,w)$, i.e. compatible with *John's* visual experience at w. It also illustrates how

'look' restricts the properties that can serve as similarity criteria to purely *visual* sets of properties ('sound' to auditory ones, and so on).

Now consider sentences embedding expressions of the form 'like a *N*'.

(41) John looks like a lawyer.

a. John looks like a specific lawyer I have in mind. SPECIFIC
b. John has the general appearance of a lawyer. GENERAL

The GENERAL reading has a stronger than existential force (looking like one lawyer is not enough) but weaker than universal (looking like a lawyer doesn't involve looking like *every* lawyer). Guerrini proposes that this is because the indefinite goes into the restriction of the generic quantifier, a silent quantificational adverb postulated in Krifka *et al.* (1995):

(42) $[\![(41)_{\text{GENERAL}}]\!]$ = GEN[x is a lawyer][John is like x]

This is because this indefinite patterns with more vanilla generic interpretations in a number of ways. For instance, saying that someone looks like a lawyer is intuitively equivalent to saying that someone looks like a *typical* lawyer - like for vanilla generics.[9]

(43) a. A bird is flying. (∃ **indefinite**)
b. $\not\approx$ A typical bird is flying.

(44) a. A bird flies. (GEN **indefinite**)
b. $\approx$ A typical bird flies.

(45) a. John looks like a lawyer I know. (SPECIFIC)
b. $\not\approx$John looks like a typical lawyer I know.

(46) a. John looks like a lawyer. (GENERAL)
b. $\approx$John looks like a typical lawyer.

Second, constructions involving 'like' PPs are non-monotonic, as illustrated in (47) and (49), just like more vanilla characterizing sentences, cf. (48) and (50).

(47) a. John looks like a British judge. (GENERAL)
b. $\not\models$ John looks like a judge.

(48) a. A British judge wears a wig.
b. $\not\models$ A judge wears a wig.

(49) a. John looks like a bird. (GENERAL)
b. $\not\models$ John looks like a penguin.

(50) a. A bird flies.
b. $\not\models$ A penguin flies.

[9] I illustrate with 'look like a *N*' constructions because they sound more natural, but the same results apply to 'be like a *N*' constructions.

Third, constructions such as 'like a lawyer or a judge' display scope ambiguity. The narrow reading (51c) is almost conjunctive: we get the inference that John has properties that a lawyer and judge *share* – and the sentence becomes almost equivalent to 'John is like a lawyer *and* a judge'.

(51) a. Bob is like a lawyer or a judge.
b. (WIDE) *is-like(B,lawyer)* $\vee$ *is-like(B,judge)*
c. (NARROW) *is-like(B, lawyer-or-judge)*

Fourth, 'like' exhibits the same subtrigging effects that Carlson (1981) discovered for vanilla generic sentences with 'someone'.

(52) Someone should be punctual.
a. Existential
b. #Generic

(53) Someone who respects others should be punctual.
a. Existential
b. Generic

(54) John looks like someone.
a. Existential
b. #Generic

(55) John looks like someone who respects others.
a. Existential
b. Generic

We are equipped to plug in Guerrini's semantics for 'look like a gun' into our lexical entry for 'fake'. Just as a reminder, we want a fake gun to:

(i) be intended to seem like a gun
(ii) not be a gun

Then:

(56) a. $[\![\text{ fake }]\!] = \lambda Q.\lambda x.$
$\text{INTENDED}\big(x,$
$\forall w' \in R_s(j,w).$
$\text{GEN}\big[Q(y)(w')\big]\big[\forall P \in \mathbf{D}.P(x)(w') = P(y)(w')\big]$
$\big)$
$\wedge \neg Q_w(x)$
b. Paraphrase:
"x is intended so that
at all worlds compatible with the experience of j,
x shares relevant properties P with typical instances y of Q
and x is not a Q"

My analysis of 'fake' is, of course, compatible with any appropriated one-dimensional entry for 'like'. This is just an illustration to provide a full specification of the semantics of 'fake', not a committal to Guerrini's proposal. However, it is worth pointing out that the presence of genericity does not seem unjustified: an object is not a fake gun in virtue of resembling a specific, very atypical gun. Instead, it has to be intended to mimic some diagnostic properties that are typical of guns.

4.2 Discarding an obvious alternative: 'fake' grounded in misguided belief

Let us first consider a simple, prima-facie appealing theory that grounds 'fake' in a combination of well-understood propositional attitudes. A fake P, in this theory, is someone/something that intends/is intended to cause an observer to believe it is a P, but isn't one:

(57) a. $[\![\text{fake}]\!] = \lambda x.\lambda P.\text{INTENDED(-TO-CAUSE)}\big[x, \text{BELIEVE}\big(j, P(x)\big)\big] \wedge \neg P(x)$
b. 'x is intended to cause j to believe x is a P, and x is not a P'

This theory doesn't work. To see this, take the following scenario:

(58) Mary wants to pull a prank on John and scare him. She gives him a Toshiba radio of model X which does not actually broadcast true radio programs, but fake program recordings that Mary transmits. In one of the newscasts transmitted by Mary, it is announced that all Toshiba radios of model X are bombs that could explode shortly.

The Toshiba radio is intended to make a judge, John, believe that it is a bomb, and isn't one. However, in this context, the following sentence is infelicitous:

(59) # John's Toshiba radio is a fake bomb.

On the other hand, a sentence about caused belief seems to be utterable:

(60) John's Toshiba radio is intended to cause John to believe that it is a bomb and it is not a bomb.

Notice that judgments about 'fake' completely line up with judgments about appearance and similarity.

(61) a. #John's Toshiba radio seems like a bomb.
b. #John's Toshiba radio sounds like a bomb.

This seems to be a further argument that appearance and similarity (something like

'seem like') are part of the semantics of 'fake'.[10,11]

4.3 Fake Rolexes

One issue that I have not addressed so far is that in some uses of 'fake', 'fake *N*' does not seem to entail 'not *N*'. For instance, a counterfeit Rolex can felicitously be called a fake watch, even though it qualifies as a watch. To accommodate this observation, we need to take into account some facts about the syntax of adjectives.

Cinque (2010, 2014) shows that adjectives have in general two syntactic modification sources: (i) Direct Modification (DM) and (ii) Reduced Relative Clause (RRC). (i) DM is merged closer to the noun and, roughly, associated with interpretive properties compatible with a semantic combination featuring Functional Application. (ii) RRC, instead, is merged farther from the noun, and is associated with a semantic combination featuring Predicate Modification.

These two sources do not come apart in English surface structures, where the word order is the same across the two sources. But they do in Italian, where only the DM can show up pre-nominally:

(62) un falso orologio
a fake watch
'a fake watch'
a. intended to look like a watch and isn't a watch

[10] It is worth mentioning that this is in contrast with sentence-embedding 'seem' and 'seem like', which are completely acceptable in this context.

(i) a. It seems that John's Toshiba radio is a bomb.
b. John's Toshiba radio sounds like it is a bomb.

The contrast between (61) and (i) is parallel to what we find with simple properties:

(ii) **Context:** John has nothing typical of French people: he dresses like an American, speaks French with an American accent, and so on. But the speaker has indirect evidence pertaining to Bob's behavior that John has a French passport.
a. #John seems French.
b. John seems like he's French.

(iii) **Context:** John has nothing typical of French people. But the speaker indirectly gathered, from what John said in a conversation, that he has a French passport.
a. #John sounds French.
b. John sounds like he's French.

This coheres with Rudolph's (2019) finding that when predicates like 'seem like' or 'be like' embed a proposition, they are compatible with more epistemic-like accessibility relations. Their only requirement is that the evidence for the embedded clause be perceptual, i.e. broadly perceptual for 'seem' and acoustic for 'sound like'. The predicate-embedding 'seem like' and 'sound like', instead, strictly require their evidence to be perceptual and be directly diagnostic of the predicate.

[11] Another reason why (59) is degraded may be that misguided belief may be too strong. Toy guns are very often produced with a red cap on top of the barrel, to signal that they are not real guns. Suppose that even children knew that if an object has such a red cap, it cannot be a real gun. Then, a toy gun with a red barrel could not fool anyone into thinking it is a gun. However, it would still qualify as a fake gun. This is surprising under the lexical entry in (57a).

This is predicted under the present account, since the toy gun shares many other perceptually salient properties with guns, such as the shape.

b. * intended to look like a Rolex and isn't a Rolex

(63) un orologio falso
a watch fake
'a fake watch'
a. # intended to look like a watch and isn't a watch
b. intended to look like a Rolex and isn't a Rolex

This pattern can be explained as follows. In DM, 'falso' directly combines with its input noun, 'orologio', thereby giving rise to the truly privative reading in (62a). The 'intersective' reading (62b) is unavailable because, simply, the nominal 'Rolex' is not a surface argument of 'fake'.

In RRC, I assume, following Martin (2022), that the reduced relative clause can only combine with the noun that is the surface argument of the adjective via Predicate Modification. For this to be possible, the adjective must first combine with a covert nominal: this allows the reduced relative clause to be of type $\langle e,t\rangle$ and to compose with the noun via Predicate Modification. Notice that this assumption correctly predicts RRC-resulting readings to be intersective.

(64) $[\![\text{ fake }]\!]([\![N_{covert}]\!]) \cap [\![N_{surface\_argument}]\!]$ (from Martin 2022)

This yields the observed pattern. Reading (63a) is ruled out because if the covert nominal happens to be 'orologio', we get a contradiction:

(65) $[\![(63a)]\!] = \lambda x.\text{INTENDED}\big(x, \text{SEEM-LIKE}(x, watch)\big) \wedge \neg watch(x) \wedge watch(x)$

However, if the covert nominal is 'Rolex', or any other nominal that denotes a relevant property P such that $[\![\text{watch}]\!] \not\subseteq P$, we get a non-contradictory reading.

(66) $[\![(63b)]\!] = \lambda x.\text{INTENDED}\big(x, \text{SEEM-LIKE}(x, Rolex)\big) \wedge \neg Rolex(x) \wedge watch(x)$

The ambiguity observed in English constructions such as (67), then, is a bona-fide *syntactic* ambiguity, since the word order is the same for the two sources (Cinque, 2010, 2014). Reading (67a) is the product of a DM composition, while (67b) is the product of a RRC composition.

(67) fake watch
a. intended to look like a watch and isn't a watch
b. intended to look like a Rolex and isn't a Rolex

4.4 More on semantic negation in 'fake'

4.4.1 Is a fake *N* a *N*?

Is it necessary to include semantic negation in the meaning of 'fake'? Recently, Martin (2022) argued for a DCS-style entry for 'fake' that lacks semantic negation. Instead, the intuition that a fake N is not a N arises as an interaction between considerations

about the C-structure and Partee's pragmatic principles: in a nutshell, he claims that 'fake' uploads a part of the C-structure to the E-structure that, most of the times, is incompatible with the original E-structure. For instance, 'fake' may manipulate and then upload some C-structural features of 'gun' and upload them to the E-structure of 'fake gun', which already contains the E-structure of 'gun', *gun(x)*. Because these manipulated C-structural features are incompatible with *gun(x)*, Partee's Non-Vacuity Principle is triggered. Crucially, however, Martin predicts that when the intersection between the uploaded C-structure features and the E-structure of the input NP is not vacuous, there should be no negative inference.

This move is made on experimental grounds: he presents the results of an experiment in which he asked participants whether a *Adj N* is a *N*, finding great variation both along privative adjectives and along different input nouns. What concerns us here is that, for some *N*s, participants responded "yes" to the question "is a fake *N* a *N*?". However, this experiment presents a significant confound, precisely because of the possibility of covert arguments in the RRC interpretation of 'fake *N*' discussed in the previous section.

Suppose I point at an object that looks like a Rolex but I know isn't one, and say, felicitously:

(68) That is a fake watch.

I am not at all committing to this object not being a watch in this context. But crucially, I am absolutely committing to the object not being a Rolex!

Thus, Martin's experiment does not provide evidence that privatives provide semantic negation; rather, it shows quite clearly that privatives have a certain freedom in the covert nominal they take as an argument when they are combined in a RRC syntactic source. To be informative, an experiment testing these facts should be run in a language which has a word order that unambiguously indicates a Direct Modification source, like Italian. To wit, if Italian speakers were to judge that a 'falso orologio' (*fake watch*) may or may not be a watch, then the theory provided here would be falsified and Martin's claims about the absence of semantic negation in 'fake' corroborated.

In the absence of such evidence, we have some rather strong arguments in favour of semantic negation. First, Italian judgments about privatives in pre-nominal position indicate that a truly privative meaning seems to systematically emerge in such positions (cf. (62a); see Cinque 2010, 2014). Second, as we have seen in section 3.2, the way the conjunction of $\neg P(x)$ and $\text{INTENDED}\big(x, \text{SEEM-LIKE}(x, P)\big)$ enters the compositionality in iterated application of 'fake' delivers fine-grained and right predictions.

4.4.2 *That gun is fake*

Another question concerning semantic negation is raised by sentences such as (69), which are among Partee's arguments for a subsective treatment of 'fake':

(69) That gun is fake.

The covert nominal taken by 'fake' seems to be 'gun', given that (69) is equivalent to (70):

(70) That object is intended to look like a gun but isn't one.

So the original problem of sentences like (69) remains: if 'fake' is truly privative, why is it that we can refer to the object in question as a gun?

In fact, I believe that Partee's original idea is the correct account for these cases. In other words, while it does not account for the compositional *meaning* of 'fake', Partee's account can account for why a Priv-Adj *N* can sometimes be referred to as an *N*. With a slight modification of her account to adapt it to copulas, we can sketch a Partee-style account of sentences like (69) coupled with my semantics for 'fake'.

(71) **Subject Primacy (SP):** in a sentence of the form '*D NP* is *Adj*', *NP* is interpreted relative to the context of the whole sentence, *Adj* is interpreted relative to the local context created from the former context by the interpretation of *NP*.[12]

(72) **Non-vacuity for Copulas (NVC):** interpret a sentence of the form '*D NP* is *Adj*' so that $NP \cap Adj \neq \emptyset$

Keeping in mind that the only syntactic source allowed in the copula is RRC (cf. Cinque 2010), the reasoning can go as follows:

(73) *That gun is fake(gun)*

Try to interpret the sentence so that
$\{x\text{: gun}(x)\} \subseteq \{x\text{: INTENDED}\big(x, \text{SEEM-LIKE}(x, \text{gun})\big) \wedge \neg\text{gun}(x)\}$,
get an empty extension
$\Rightarrow$NVC fails.
Give up SP: extend $\{x\text{: gun}(x)\}$ to include both real and fake guns.

Notice that now that we have a precise lexical entry for 'fake', we know precisely what the broadened denotation of 'gun' is:

(74) $[\![\text{gun}_{broadened}]\!] = [\![\text{gun}_{proper}]\!] \cup [\![\text{fake gun}]\!] =$
$= \{x : gun(x)\} \cup \{x : \text{INTENDED}\big(x, \text{SEEM-LIKE}(x, gun)\big) \wedge \neg gun(x)\}$

Now, of course in a sentence of the form '*D NP* is fake', the noun will only be coerced if NVC fails. This means that if 'fake' combines with a covert nominal N_1 such that $[\![NP]\!] \not\subseteq [\![NP]\!]$, no reasoning like the one above is triggered. To wit, (75) under the reading in (75a) is not contradictory to begin with, so NVC does not fail.

(75) This watch is fake.
a. This watch is fake (as a rolex)

[12] I here take up Partee's terminology; the notion of local context employed here is not necessarily the one we think of today.

This leads us to reassess Italian non-copular modification: why can't a Parteean procedure not rescue the truly privative reading for the RRC modification source? Recall (62) and (63), repeated below in (76) and (77).

(76) **Direct Modification (DM) Source**
un falso orologio
a fake watch
a. intended to look like a watch and isn't a watch
b. # intended to look like a Rolex and isn't a Rolex

(77) **Reduced Relative Clause (RRC) Modification Source**
un orologio falso
a watch fake
a. # intended to look like a watch and isn't a watch
b. intended to look like a Rolex and isn't a Rolex

Of course that (76b) is not available is expected independently: given the syntactic properties of DM, the adjective cannot combine via Predicate Modification, but only via direct Functional Application. And the Parteean mechanism only applies to cases of Predicate Modification.

However, it would surprising that a Parteean mechanism could save copular sentences like (69) but not an RRC reading like (77a). However, there are good reasons to think that (77a) is in fact generally possible, but that it needs a strong context to be accessed. This is because for (77a) we need to go through Partee's pragmatic adjustment to resolve the initial contradiction, while (77b) is good as is.

And indeed, judgments are more unstable than the literature suggests: according to four informants I consulted, (77a) is accessible, although it needs some context.[13]

(78) John thought that he would be able to know what the time was once in the desert, but he realized that...
a. ...era stato equipaggiato con un orologio falso.
...was-3SG equipped with a watch fake.
'...he had been equipped with a fake (as watch) watch.'

This sentence means that the object with which John has been equipped was intended to resemble a watch (successfully, in fact) but isn't one; this is the (77a) reading.

In fact, in some cases there is no salient covert nominal different from the noun available to saturate the argument of 'fake'. In these cases, speakers report that the 'truly privative' reading is the most accessible even in the RRC source. For instance, absent a strong context, (79a) is the only reading of (79)

(79) finestra falsa
window fake
a. an object that is intended to look like a window but isn't a window

[13] As a native speaker of Italian, I myself share these judgments.

All of this suggests that in RRC in Romance, in an NP like ⟦finestra falsa(finestra$_{covert}$)⟧ under a reading like (79a), the overt occurrence of 'finestra' gets broadened via a Parteean procedure, as illustrated below.

(80) *window fake(window$_{covert}$)*
Try to interpret the expression so that
$\{x$: window$(x)\} \subseteq \{x$: INTENDED$\big(x$, SEEM-LIKE$(x$, window$\big) \wedge \neg$ window$(x)\}$,
get an empty extension
⇒Non-Vacuity fails.
Give up Head Primacy: extend $\{x$: window$(x)\}$ to include both real and fake windows.

4.4.3 Taking stock

One may wonder if the conclusions reached in 4.4.1 and those reached in 4.4.2 aren't contradictory: in 4.4.1 I argued that the meaning of 'fake' is truly privative in the sense that it contains a semantic negation of its input. In 4.4.2, I appealed to Parteean mechanisms to explain why a fake *N* may sometimes referred to as an *N*. I think that, in fact, these two conclusions are compatible.

The reason why the meaning of 'fake' should contain the semantic negation of its input is simple: in certain syntactic configurations (Direct Modification in Italian), the inference from "fake N" to "not an N" is systematic and obligatory.

On the other hand, the Reduced Relative Clause syntactic source only allows Predicate Modification. In these cases, we have to allow pragmatic coercion of the noun when the intersection between the noun and the modifier happens to be empty.

The two components are therefore not at odds with each other, quite the opposite. As shown in (74), the truly privative meaning of 'fake' helps us understand how precisely the denotation of 'gun' is broadened in sentences such as (69), reported below in (81):

(81) This gun is fake.

Indeed, in the account argued for here, in such sentences the denotation of 'gun' is broadened to include both guns and objects that are intended to resemble a gun. This explicitness was missing from Partee's original proposal. But crucially, this pragmatic adjustment is not an ad-hoc adjustment to have the theory capture sentences such as (81). The prediction that in Functional Application 'fake' should be truly privative is indeed borne out: when the syntax unambiguously determines Direct Modification, and therefore Functional Application, we get an entailment from 'fake *N*' to 'not an *N*'.[14]

[14]One question that remains is why certain adjectives cannot enter copular constructions. For instance, 'alleged' can appear in NPs, but not in copulas:

(i) an alleged murderer
(ii) *That murderer is alleged

Put in theory-internal terms, why can certain adjectives only combine via DM, and thus determine semantic Functional Application, and not via RRC, and thus determine semantic Predicate Modification? In yet other terms, why can certain adjectives not take a covert N as an argument, but can only fill their ar-

5 Conclusion

I have argued for a one-dimensional, intrinsically privative semantics for 'fake'. Basing 'fake' on similarity allows us to see that the considerable semantic complexity of this adjective is actually mostly compositionally inert. Precisely *what* is fake is very open-ended, because similarity is a contingent notion: a fake *X* may simulate properties that a *X* merely *happens* to have. More precisely, I argued that a fake *X* is (i) intended to look like a *X* and (ii) is not an *X*. I showed that this makes right predictions for iterated application of 'fake', as well as for the stacking of 'fake' on top of 'typical'.

References

Asher, N. (2011) *Lexical meaning in context: A web of words*. Cambridge University Press.

Carlson, G.N. (1981) Distribution of free-choice any. *Papers from the Seventeenth Regional Meeting of the Chicago Linguistics Society*, 17, pp. 8–23.

Cinque, G. (2010) *The syntax of adjectives: A comparative study*, vol. 57

Cinque, G. (2014) The semantic classification of adjectives. a view from syntax. *Studies in Chinese Linguistics, 35.(1) 3-32* .

Clark, R. (1970) Concerning the logic of predicate modifiers. *Nous* pp. 311–335.

Del Pinal, G. (2015) Dual content semantics, privative adjectives, and dynamic compositionality. *Semantics and Pragmatics* **8**, 7–1–53 1937–8912.

Del Pinal, G. (2018) Meaning, modulation, and context: a multidimensional semantics for truth-conditional pragmatics. *Linguistics and Philosophy* **41**.

Feinmann, D. (2020) *Lexical illusions,(non-) maximality, and invisible gaps*. Ph.D. thesis, University of Sheffield Sheffield, UK.

Guerrini, J. (2021) Similarity statements, privative adjectives, and generics: a unified view. *Qualifying paper, Ecole Normale Supérieure* .

Guerrini, J. (2022) *Like a N* constructions and genericity. *Proceedings of the ESSLLI 2022 Student Session*.

gument by explicit composition with the noun? All that I can note, for now, is that the difference does not seem to be determine by meaning. For instance, in French, 'nul', *'bad'*, and 'piètre', *'bad'* are very similar in meaning, but only 'nul' can be used in copular constructions:

(iii) Un piètre avocat
A bad lawyer

(iv) Un avocat nul
A lawyer bad

(v) Cet avocat est {nul/*piètre}
This lawyer is bad

Guerrini, J. & Mascarenhas, S. (2019) Shifting centers: toward a unified view of grammatical and contingent privative modification. *Brain, Language, and Learning Conference, Università di Siena*.

Heim, I. (1992) Presupposition projection and the semantics of attitude verbs. *Journal of semantics* **9**, 183–221.

Kamp, H. (1975) *Two theories about adjectives*, pp. 123–155. Cambridge University Press.

Krifka, M., Pelletier, F.J., Carlson, G., Ter Meulen, A., Chierchia, G. & Link, G. (1995) Genericity: an introduction .

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

Martin, J. (2018) Compositionality in privative adjectives: extending dual content semantics. *European Summer School in Logic, Language and Information*, pp. 93–107, Springer.

Martin, J. (2022) *Compositional routes to (non) intersectivity*. Ph.D. thesis, Harvard University.

Montague, R. (1970) English as a formal language. *Linguaggi nella Società e nella Tecnica*, B. Visentini et al. (eds), Milan, Repr. in Montague 1974.

Moravcsik, J.M. (1998) Meaning, creativity, and the partial inscrutability of the human mind .

Pagin, P. & Pelletier, J. (2007) Content, context, and composition. *Context-sensitivity and semantic minimalism: New essays on semantics and pragmatics* pp. 25–62.

Parsons, T. (1970) Some problems concerning the logic of grammatical modifiers. *Synthese* **21**, 320–334.

Partee, B.H. (2010) Privative adjectives: subsective plus coercion. *Presuppositions and discourse: Essays offered to Hans Kamp* **21**, 273–285.

Pustejovsky, J. & Bouillon, P. (1995) Aspectual coercion and logical polysemy. *Journal of semantics* **12**, 133–162.

Putnam, H. (1970) Is semantics possible? *Metaphilosophy* **1**, 187–201.

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

Rudolph, R.E. (2019) *Talking about appearances: Experience, evaluation, and evidence in discourse*. Dissertation. University of California, Berkeley.

Szabó, Z. (2010) The determination of content. *Philosophical Studies* **148**, 253–272.