

Keeping *fake* simple *

Janek Guerrini[†]

Abstract

In this paper, I argue against two common claims about privatives like ‘fake’: first, I argue against the idea that the semantic complexity of ‘fake’ requires a richer notion of lexical meaning than the standard notion (see, e.g., Del Pinal, 2018); second, I argue against the idea that ‘fake’ does not semantically negate its input and rather behaves like a subsecutive (see, e.g., Partee, 2010). I propose that a fake *P* is (i) intended to resemble a *P* and (ii) is not a *P*. This makes correct predictions for multiple applications of ‘fake’, a task at which other theories fail. In double application of ‘fake’, the interaction between the conjunctive meaning of ‘fake’ and the negation hard-coded into clause (ii) yields a complex meaning, compatible with a variety of objects. While the core meaning of ‘fake’ is quite simple, its mode of composition bears some complexity. I propose that ‘fake’ can alternatively (a) combine directly with the noun via Functional Application or (b) saturate its property argument via an implicit, contextually provided variable via Functional Application and then combine with the noun via Predicate Modification (cf. Martin, 2022 for a similar view). Mode of composition (a) is clearly visible in syntactic parses that only allow for Functional Application: for instance, in Italian, if prenominal, ‘fake’ can only directly take the noun as an input (cf. Cinque, 2010). Positing (b) well-predicts readings where ‘fake’ is not apparently privative: ‘fake watch’ can designate a watch that is made to look like a Rolex but isn’t one, i.e. a watch that is fake-as-a-Rolex. When the intersection between the $\llbracket fake \rrbracket(c)$ complex and the noun is empty, rescuing principles originally proposed by Partee kick in to rescue from vacuous modification: this explains why we can refer to a fake gun as a gun, as in the sentence ‘this gun is fake’. As a result, besides well-predicting iterated ‘fake’, this theory provides clear predictions on when and how Partee’s pragmatic principles of noun modulation apply. I conclude the paper arguing that this view of privatives suggests a novel way of classifying adjectives, in terms of their compositional behavior rather than by their emergent entailment pattern.

*This work 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 Genaro Chierchia for his valuable input. Thank you, finally, to Guillermo Del Pinal, Viola Schmitt, Philippe Schlenker, Emmanuel Chemla, Paul Egré, Jeremy Kuhn, Diego Feinmann, and Amir Anvari for their very helpful comments.

This research was supported by the ANR program PROBASEM (ANR-19-CE28-0004-01, PI: Spector). This research also received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 788077, Orisem, PI: Schlenker). Research was conducted at Institut d’Etudes Cognitives, Ecole Normale Supérieure - PSL Research University. Institut d’Etudes Cognitives is supported by grants ANR-10-IDEX-0001-02 and FrontCog ANR-17-EURE-0017.

[†]Institut Jean Nicod, Department of Cognitive Studies, Ecole Normale Supérieure, EHESS, CNRS, PSL University

Contents

1	Introduction	3
2	Iteration problems in previous accounts	7
2.1	Partee	7
2.2	Del Pinal	11
3	Proposal: similarity and negation	17
3.1	Testing iteration in the present account	18
3.2	<i>Fake typical N</i>	20
4	Negation in ‘fake’	22
4.1	Fake Rolexes	22
4.2	Is a fake <i>N</i> a <i>N</i> ?	24
4.3	‘That gun is fake’	25
5	Similarity in ‘fake’	28
5.1	One-dimensional ‘fake’ with one-dimensional similarity	28
5.2	Discarding an obvious alternative: ‘fake’ grounded in deception	33
6	Conclusion	34

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.

$$(1) \quad \llbracket Adj N \rrbracket \not\subseteq \llbracket N \rrbracket$$

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 $\llbracket Adj N \rrbracket$ should be a subset of $\llbracket N \rrbracket$, which is intuitive if one looks at most adjectives, from so-called intersectives ('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, proposing that privatives are actually subsectives in disguise in that they do not negate their input noun. 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, which is usually 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 $\llbracket lawyer \rrbracket$.

$$(2) \quad \begin{array}{l} \text{a. } \llbracket \text{French lawyer} \rrbracket = \llbracket \text{French} \rrbracket \cap \llbracket \text{lawyer} \rrbracket \\ \text{b. } \llbracket \text{good lawyer} \rrbracket \neq \llbracket \text{good} \rrbracket \cap \llbracket \text{lawyer} \rrbracket \\ \text{c. } \llbracket \text{fake lawyer} \rrbracket \neq \llbracket \text{fake} \rrbracket \cap \llbracket \text{lawyer} \rrbracket \end{array}$$

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 privatives disagree about whether their 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 account: $\llbracket \text{fake gun} \rrbracket$ is a subset of a broadened set of guns (cf. 2.1). More specifically, when trying to restrict 'gun' with 'fake', we systematically end up with a vacuous extension – the intuition being that there are no

bona fide guns that are fake. In order to avoid ending up with an empty extension, we widen the set of guns via specific pragmatic mechanisms so that it includes both guns and gun-like objects.

Other theorists have it that privatives force 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 his 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. 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 Pinal¹ 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, as illustrated informally below.

(3) fake = $\lambda P.\lambda x. x$ is intended to resemble a $P \wedge x$ is not a P

Note that for ease of exposition, I keep this illustration very informal. For instance, ‘resemble a P ’ here is a black box – but this indefinite will end up receiving a generic interpretation. I will write schematized, informal entries like (3) without interpretation brackets, to distinguish them from formal entries. For simplicity, when schematizing entries, I will always work within a purely extensional framework, but ‘fake’ is ultimately of type $\langle\langle s, \langle e, t \rangle \rangle, \langle e, t \rangle \rangle$, as shown in more formal entries later on in the paper (cf. sections 3 and 5.1).

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 semantics² we can dispense with the C-structure entirely 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 the class of privative adjectives as a whole. For instance, Del Pinal (2015)

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

²Cf. for instance Guerrini, 2022a,b.

argues that ‘counterfeit’ and ‘artificial’ depart in important ways from ‘fake’ but, crucially, always paraphrases them with expressions involving similarity. Roughly, Del Pinal proposes that a counterfeit X is not an X , is intended to look *like* an X , and is intended to function *like* an X . An ‘artificial X ’, instead, tells us that an artificial 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 sense, the broader point made in this paper is that for those privatives that involve an adequate paraphrase in terms of similarity, a one-dimensional semantics suffices. Once these more recalcitrant cases are treated one-dimensionally, all privatives can plausibly be brought to heel. Other privatives indeed seem 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.³ Others are modal privatives such as ‘imaginary’ or ‘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.⁴

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’ (cf. section 2). 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.

- (4) x is a non-non- $X \models x$ is an X (5) x is a fake fake $X \not\models x$ is a X

If you are unsure about (5), notice for instance that (6a) qualifies as a gun, but not (6b). Yet both are fake fake guns. Likewise, (7a) qualifies as a KGB agent, but not (7b). Yet both are fake fake KGB agents.

- (6) 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.
- (7) a. FAKE FAKE KGB AGENT; KGB AGENT.
John is a KGB agent who wants the Americans to believe that he’s actually a French agent infiltrated in the KGB (e.g., for them to share information about their own infiltrated agents), so pretends to be a French fake

³Although double application of ‘non-’ may have significant pragmatic effects. On this note, see for instance Tessler & Franke (2019).

⁴One 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, 2019) privativity contrasts with ‘grammatical’ or ‘functional’ privative modification (the one that concerns ‘fake’ *etc.*) as follows. Grammatical privatives privatize 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 privatize 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.

KGB agent.

☞ John is a fake fake KGB agent.

- b. FAKE FAKE KGB AGENT; NOT A KGB AGENT.

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.

☞ John is a fake fake KGB agent.

In my theory, which I present in section 3, this unstable entailment pattern is predicted by the interaction between the negation of the input of ‘fake’ and its conjunctive meaning, as shown informally below:

$$\begin{aligned}
 (8) \quad \text{fake}(\text{fake gun}) &= \\
 &\lambda x. x \text{ is intended to resemble a fake gun} \wedge x \text{ is not a fake gun} = \\
 &\lambda x. x \text{ is intended to resemble a fake gun} \wedge \neg(x \text{ is not a gun} \wedge x \text{ is intended to look like a gun}) = \\
 &\lambda x. x \text{ is intended to resemble a fake gun} \wedge \underbrace{(x \text{ is a gun} \vee x \text{ is not intended to look like a gun})}_{(I)}
 \end{aligned}$$

Note indeed that there are both guns and non-guns that respect (I) above. The entry above, when made explicit, is one where ‘fake’ is an $\langle\langle s, \langle e, t \rangle \rangle, \langle e, t \rangle\rangle$ function – in analogy to Partee, where it is $\langle\langle e, t \rangle, \langle e, t \rangle\rangle -$, which combines with its argument via (Intensional) Functional Application. However, it has a significantly different meaning, in that it semantically negates its input noun. Among other things, this ensures that privatives systematically privatize whatever noun they take as an input. This seems like a desired outcome, in that it distinguishes grammatically privative adjectives from emergent privative modification, e.g. in noun-noun compounds. The former, but not the latter, systematically privatize their input (Franks, 1995; Del Pinal, 2015; Martin, 2018; Guerrini & Mascarenhas, 2019).

- (9) **Grammatical privative modification.**

- a. x is a fake lion $\models x$ is not a lion.
 b. x is a fake statue $\models x$ is not a statue.

- (10) **Emergent privative modification.**

- a. x is a stone lion $\models x$ is not a lion.
 b. x is a stone statue $\not\models x$ is not a statue.

Apparent counterexamples to the systematic privativity of ‘fake’ are cases like (11):

- (11) (*Pointing at a functioning watch that is a counterfeit Rolex*)
 This is a fake watch.

In line with Martin (2022), I show in section 4 that ‘fake’ can either (i) compose

directly via (Intensional) Functional Application with the noun, or (ii) first via (Intensional) Functional Application with a covert argument of type $\langle e, t \rangle$, and then via Predicate Modification with the noun. The latter composition is at work in (11), where ‘fake’ first composes via Functional Application with a covert property argument ‘Rolex’ and then, via Predicate Modification, with the property denoted by the noun ‘watch’.

This view finds corroboration in data from Italian, which has two possible *Adj+Noun* syntactic configurations. One of these is incompatible with Predicate Modification, thus forcing an LF where ‘fake’ directly composes via Functional Application with the noun (cf. Cinque, 2010; Martin, 2022). Indeed, pre-posed adjectives can only have the reading in which ‘fake’ takes as its input the property denoted by the noun, as shown below (the unmarked *Noun-Adj* order is instead ambiguous like in English).

- (12) un falso orologio
 a fake watch
 ‘a fake watch’
- a. an object that is intended to look like a watch but isn’t one
 b. *a watch that is intended to look like a Rolex but isn’t one

We will also see that pragmatic principles like those originally proposed by Partee are still active in composition, but in ways more constrained than was originally proposed. In a nutshell, they are never active in Functional Application. Instead, they are only active in Predicate Modification, and exclusively when the noun is included in the extension of the implicit argument: in those cases the result would be vacuous in case no additional pragmatic weakening were at work.

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 section 3, I state in detail my proposal, showing that it makes the right predictions for iteration 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’. In section 4, I discuss some apparent counterexamples to the presence of semantic negation in the lexical entry for ‘fake’, and show that they fall in line once one fully articulates the ways in which ‘fake’ can compose with its input noun. Finally, in section 5 I discuss some implications resulting from the presence of similarity within ‘fake’: for instance, I illustrate how ‘fake’ can be kept one-dimensional by plugging in a one-dimensional entry for similarity verbs.

Finally, I conclude by arguing that the view defended in this paper suggests a novel way of classifying adjectives, in terms of their compositional behavior rather than in terms of their emergent entailment pattern.

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 this framework, privatives are actually subsecutive 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 goes as follows. The function denoted by ‘fake’ always outputs an extension that is not included in the extension of the input noun. However, a meaning postulate requires the output of any adjective to be included in the extension of the input noun. ‘Fake’ is thus bound to output empty extensions. Partee therefore 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 mod-
iffee.

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, since there are no actual lions made of stone. Conversely, if speakers hold on to the Non-Vacuity Principle, the head primacy principle is 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. Let us start by looking at ‘stone lion’.

- (13) ‘Stone lion’
 Try to interpret the expression as $\{\text{stone lion}\} \subseteq \{\text{lion}\}$, get an empty extension
 \Rightarrow 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.⁵

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

- (14) There are no fake [fake guns] in the extension of ‘fake gun’.



- (15) 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’ was extended to include both real and fake guns).

⁵It is also worth mentioning here Del Pinal’s arguments against this account. 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.

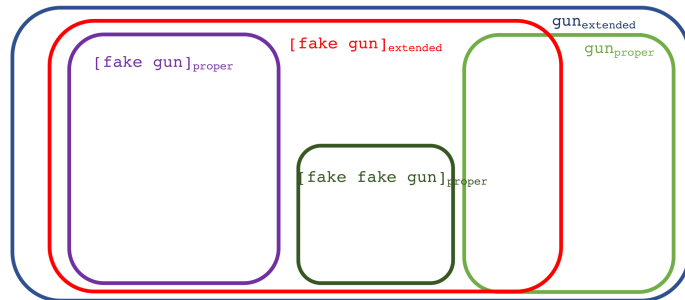


Figure 1: A representation of the step in (15): relax head primacy and broaden the extension of fake-guns-proper to fake-guns-extended, to include both real and fake [fake guns].

But why precisely needs the set of guns to be broadened 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 incomplete is by extending Partee’s reasoning for cases of iterated application of ‘fake’.

We know from section 1 (but see 3.1 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.

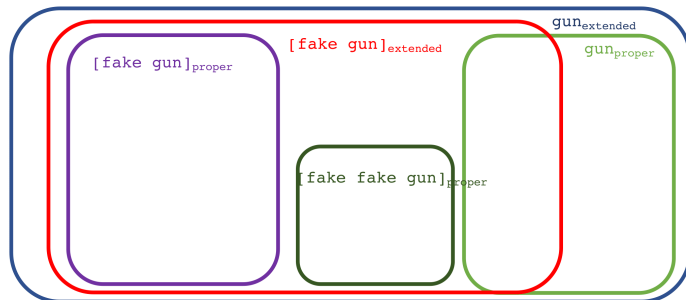


Figure 2: A scenario in which all fake fake guns are non-guns.

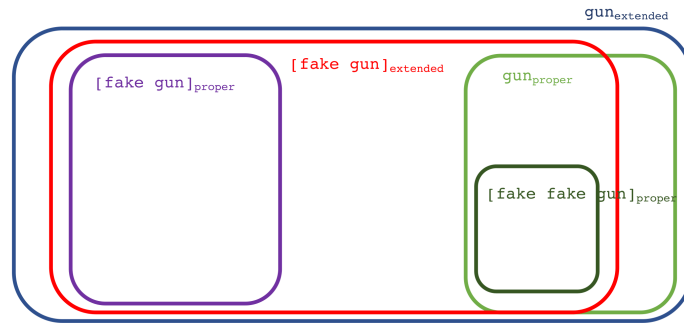


Figure 3: A scenario in which all fake fake guns are guns-proper.

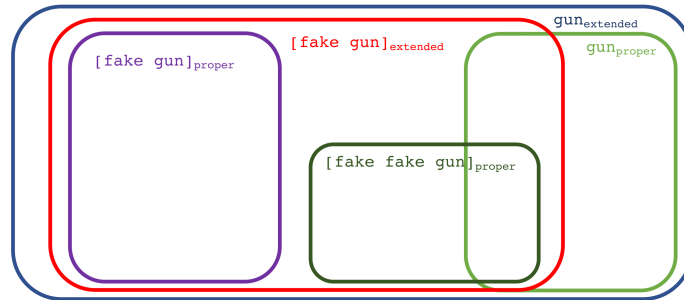


Figure 4: A scenario in which a fake fake gun may or may not be a gun-proper.

This shows that just the idea that ‘fake’ systematically violates Non-Vacuity when applied to a N is not enough to predict its behavior. However, I will ultimately maintain that something like Partee’s principles of interpretation is needed, once we have clear predictions on when Non-Vacuity is violated. I discuss this in depth in section 4.

2.2 Del Pinal

A more recent account 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; Szabo, 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 (16b) in order to convey the meaning of sentence (16a), although the meaning in (16a) should be the most relevant reading of (16b) (Asher, 2011).

- (16) 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 issues with this specific implementation of DCS, which 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 then discuss some possible amendments, and show that in fact they at least cast doubt over the need of a second layer for NP meanings.

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

- (17) $\llbracket \text{gun} \rrbracket_M^c =$
E-structure:
 $\lambda x. \text{GUN}(x)$
C-structure:
 C: $\lambda x. \text{PARTS_GUN}(x)$
 P: $\lambda x. \text{PERCEPTUAL_GUN}(x)$
 T: $\lambda x. \text{GEN } e[\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)]$
 A: $\lambda x. \exists e_1 [\text{MAKING}(e_1) \wedge \text{GOAL}(e_1, \text{GEN } e(\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)))]$

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, 1995; Moravcsik, 1998; Del Pinal, 2015, 2018).⁶

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.

⁶Other information that can be 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.⁷

- **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.

Regular intersective modifiers (e.g. ‘steel’ as in ‘steel gun’) only add descriptive content to the E- and C-structure of their argument. Other, more interesting modifiers like ‘fake’ and ‘typical’ upload parts of the C-structure of their input noun to the E-structure of their output. ‘Typical’, for instance, 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} :

$$(18) \quad \llbracket \text{typical} \rrbracket_M^c =$$

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 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, as captured in (19).

$$(19) \quad \llbracket \text{typical lion} \rrbracket_M^c =$$

E-structure:
 $T(\llbracket \text{lion} \rrbracket_M^c, x) > s$

C-structure:
 C: $\lambda x. Q_C(\llbracket \text{lion} \rrbracket_M^c)(x)$
 P: $\lambda x. Q_P(\llbracket \text{lion} \rrbracket_M^c)(x)$
 T: $\lambda x. Q_T(\llbracket \text{lion} \rrbracket_M^c)(x)$
 A: $\lambda x. Q_A(\llbracket \text{lion} \rrbracket_M^c)(x)$

Besides uploading part of the content of the C-structure of its argument to the E-structure – as ‘typical’ does –, ‘fake’ negates the E-structure of its argument:

$$(20) \quad \llbracket \text{fake} \rrbracket_M^c =$$

E-structure:

$$\begin{aligned} & \lambda \mathbf{G}.\lambda x.\neg E(\mathbf{G}(x)) \wedge \\ & \neg Q_A(\mathbf{G})(x) \wedge \\ & \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_P(\mathbf{G})(x))] \end{aligned}$$

C-structure:

$$\begin{aligned} \text{C: } & \lambda \mathbf{G}.\lambda x.Q_C(\mathbf{G})(x) \\ \text{P: } & \lambda \mathbf{G}.\lambda x.Q_P(\mathbf{G})(x) \\ \text{T: } & \lambda \mathbf{G}.\lambda x.\neg Q_T(\mathbf{G})(x) \wedge Q_P(\mathbf{G})(x) \\ \text{A: } & \lambda \mathbf{G}.\lambda x.\exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_P(\mathbf{G})(x))] \end{aligned}$$

If we apply this to the entry for ‘gun’, we end up with an E-structure that 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:

$$(21) \quad \llbracket \text{fake gun} \rrbracket_M^c =$$

E-structure:

$$\begin{aligned} & \lambda x.\neg E(\llbracket \text{gun} \rrbracket_M^c)(x) \wedge \\ & \neg Q_A(\llbracket \text{gun} \rrbracket_M^c)(x) \wedge \\ & \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_P(\llbracket \text{gun} \rrbracket_M^c)(x))] \end{aligned}$$

C-structure:

$$\begin{aligned} \text{C: } & \lambda x.Q_C(\llbracket \text{gun} \rrbracket_M^c)(x) \\ \text{P: } & \lambda x.Q_P(\llbracket \text{gun} \rrbracket_M^c)(x) \\ \text{T: } & \lambda x.\lambda x.\neg Q_T(\llbracket \text{gun} \rrbracket_M^c)(x) \wedge Q_P(\llbracket \text{gun} \rrbracket_M^c)(x) \\ \text{A: } & \lambda x.\exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_P(\llbracket \text{gun} \rrbracket_M^c)(x))] \end{aligned}$$

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 $\text{PERCEPTUAL_GUN}(x)$, i.e. that x look like an gun.

Now, applying ‘fake’ to (21) again, we get a contradiction.

$$(22) \quad \llbracket \text{fake fake gun} \rrbracket_M^c =$$

E-structure:

$$\lambda x.\neg E(\llbracket \text{fake gun} \rrbracket_M^c)(x) \wedge$$

$$\neg Q_A(\llbracket \text{fake gun} \rrbracket_M^c)(x) \wedge \\ \exists e_3 [\text{MAKING}(e_3) \wedge \text{GOAL}(e_3, Q_P(\llbracket \text{fake gun} \rrbracket_M^c)(x))]$$

C-structure:

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

After some calculations, we get something more legible:

$$(23) \quad \llbracket \text{fake fake gun} \rrbracket_M^c =$$

E-structure:

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

C-structure:

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

In other words, there both was and was not an event in which x was made to look like a gun. Three parts of the meaning of ‘fake’ 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.

It seems, then, 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. This comes out even more clearly when ‘fake’ embeds another intensional modifier like ‘typical’. Recall that, in DCS, ‘typical’ uploads a variable portion of C-structure to the E-structure.

$$(24) \quad \llbracket \text{typical} \rrbracket_M^c =$$

E-structure:

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

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').

$$(25) \quad \llbracket \text{fake typical gun} \rrbracket_M^c =$$

E-structure:

$$\begin{aligned} & \text{(i)} \quad \lambda x. \neg |T(\text{GUN}_c, x)| > s \wedge \\ & \text{(ii)} \quad \neg \exists e_1 [\text{MAKING}(e_1) \wedge \text{GOAL}(e_1, \text{GEN}_e(\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)))] \wedge \\ & \text{(iii)} \quad \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, \text{PERCEPTUAL_GUN}(x))] \end{aligned}$$

To restate these three points in simpler words, an object that is predicted to 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*: we want a fake typical gun to have those traits that are diagnostic of *typical guns*. These traits need not be the same as those that are diagnostic of guns. This will crucially depend on the extension: typical *Ns* may all happen to have a feature *F* that is not very typical of *Ns*. An individual that simulates precisely *F* is thus a fake typical *N*, but not a fake *N*.⁸

There are ways to amend Del Pinal's system. For instance, a reviewer suggests that one could assume, following by Del Pinal (2018), that C-structure dimensions actually encode a measure of similarity to the prototype.⁹

$$(26) \quad \llbracket \text{gun} \rrbracket =$$

E-structure

...

C-structure

...

P: $\lambda x. \text{SIMILARITY}(x, \lambda x'. \text{gun}(x')) > s$

⁸For 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.

⁹The reviewer also suggests that 'fake' may have the effect of *lowering* the standard *s*. This would still leave open the question of how to account for the entailment patterns characteristic of 'fake': one application has a privative effect (x is a fake *N* \models x is not an *N*), but two applications do not take us back to the root (x is a fake fake *N* $\not\models$ x is an *N*).

At this point, multiple options are open. One could imagine that for fake fake guns, this similarity measure is also iterated.

$$(27) \quad \begin{aligned} \llbracket \text{'fakegun'} \rrbracket = & \\ & \mathbf{E\text{-structure}} \\ & \dots \\ & \mathbf{C\text{-structure}} \\ & \dots \\ & P: \lambda x. \text{SIMILARITY}(x, \lambda x''. \text{SIMILARITY}(x'', \lambda x' \text{gun}(x')) > s) > s' \end{aligned}$$

Be this as it may, the main point argued for here is that a multidimensional semantics is not *needed* to account for ‘fake’: since we have to appeal to a notion of similarity, we might as well do this in a simpler framework.

3 Proposal: similarity and negation

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.¹⁰

$$(28) \quad \begin{aligned} \llbracket \text{fake} \rrbracket^w = & \lambda P_{\langle s, \langle e, t \rangle \rangle}. \lambda x. \\ \mathbf{i.} \quad & [\text{in every world } v \text{ compatible with what's intended in } w, x \text{ seems in } v \text{ to} \\ & \text{be like a } P] \wedge \\ \mathbf{ii.} \quad & \neg P(w)(x) \end{aligned}$$

Note that here ‘to be like a P ’ is a black box – in section 5.1, entry (79), we will see one way to make this and other aspects explicit. For now, for simplicity we will stick with (28), and schematize these kinds of entries extensionally, as in (29). As mentioned in the Introduction, for clarity I leave the interpretation brackets out when considering schematized entries.

¹⁰Here 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:

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

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’:

$$(ii) \quad \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_P(\llbracket \text{gun} \rrbracket_M^c)(x))]$$

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.¹¹ In this case, there is no making event. Rather, there is an action that the person performs on themselves to resemble a lawyer. This is adequately captured by the specification of INTENDED in (ia). 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.

$$(29) \quad \text{fake} = \lambda P. \lambda x. \text{INTENDED}(x, \text{SEEM-LIKE}(x, P)) \wedge \neg P(x)$$

Note that the presence of INTENDED makes 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 circumstance, i.e. that nobody molded this configuration of atoms. We cannot call this a fake gun.

3.1 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.

$$(30) \quad \text{fake fake gun} = \lambda x. \begin{array}{l} \text{i. } \neg \text{fake-gun}(x) \wedge \\ \text{ii. } \text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{fake-gun})) \end{array}$$

If we develop (31), we get:

$$(31) \quad \text{fake fake gun} = \lambda x. \begin{array}{l} \text{i. } \neg \left(\neg \text{gun}(x) \wedge \text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{gun})) \right) \wedge \\ \text{ii. } \text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{fake-gun})) \end{array}$$

which means:

$$(32) \quad \text{fake fake gun} = \lambda x. \begin{array}{l} \text{i. } \text{gun}(x) \vee \neg \text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{gun})) \wedge \\ \text{ii. } \text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{fake-gun})) \end{array}$$

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:

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

In what follows, I go through the generated readings, and show that they are in fact attested. An object in the \mathcal{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.¹²

An object in the \mathcal{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.

¹²For this and the following section, I consulted two native speakers of English who are not trained linguists and are not aware of my account or its predictions. I presented both of them with contexts such as the following, asking them if they found the target sentence acceptable:

- (i) **Context:** John wants to murder someone in a military simulation game played with airsoft weapons. To this effect, he produces a gun that can shoot actual bullets that looks like an airsoft gun. You overhear John's friend Bob, who knows about his plans, saying:

- "Yes, he's been busy all week creating this crazy object, a real gun that totally looks like those fake guns you use for airsoft. He wants it to look exactly like one of them. I feel for him..."
- ☞ "... It must be hard to build a fake fake gun."

Besides some remarks about the slight pragmatic oddness of using 'fake' twice in the same expression by one of the two consultants, they both found sentences such as the one above to be acceptable. In general, both reported a clear judgment that this reading is attested, as well as the readings discussed below. I also consulted two trained linguists who were not aware of the present theory, Cécile Crimon and Léo Migotti, whom I thank, about equivalent expressions in French ('faux faux pistolet'). They reported clear judgments and found all readings to be attested.

- (33) 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 \mathcal{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.

3.2 *Fake typical N*

In this subsection, I will show, for the sake of completeness, that the meaning of ‘fake’ I propose is general enough to account for its compositional behavior when stacked with other intensional modifiers, and take Del Pinal’s case study for ‘typical’. For this reason, the reader may skip this section without missing anything essential to the analysis itself.

We want 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 subjective 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’:

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

(35) - 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 (36):

(36) 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:

(37) fake typical gun =
 (i) $\lambda x. \neg(\text{gun}(x) \wedge C(\text{gun})(x)) \wedge$
 (ii) $\text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{typical gun}))$

This means:

(38) fake typical gun =
 (i) $\lambda x. (\neg \text{gun}(x) \vee \neg C(\text{gun})(x)) \wedge$
 (ii) $\text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{typical gun}))$

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

	\mathcal{G}_1	\mathcal{G}_2	\mathcal{G}_3
$\neg \text{gun}(x)$	1	0	1
$\neg C(\text{gun})(x)$	0	1	1
$\text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{typical gun}))$	1	1	1

An object in the \mathcal{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 \mathcal{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 \mathcal{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 Negation in ‘fake’

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, too. This may seem to suggest that ‘fake’ does not semantically negate its input. However, once we take into account some facts about the syntax of adjectives, this will prove to be a surface phenomenon.

4.1 Fake Rolexes

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 the adjective combining directly with the noun via Functional Application. (ii) RRC, instead, is merged further from the noun, and is associated with Predicate Modification between the adjective and the noun.

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:

- (39) un falso orologio
a fake watch
‘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

- (40) un orologio falso
 a watch fake
 'a fake watch'
- a. #¹³ 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 (39a). The 'intersective' reading (39b) is unavailable because, simply, the nominal 'Rolex' is not a surface argument of 'fake'.

- (41) Direct Modification
 $\llbracket \text{fake} \rrbracket_{\langle \langle s, \langle e, t \rangle \rangle, \langle e, t \rangle \rangle}^{w_0} (\lambda w. \llbracket N_{\text{surface_argument}} \rrbracket^w)$

In RRC, I assume, following Martin (2022), that the reduced relative clause can only combine via (Intensional) Functional Application 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.

- (42) Reduced Relative Clause
 $\llbracket \text{fake} \rrbracket^{w_0} (\lambda w. \llbracket N_{\text{covert}} \rrbracket^w) \cap \llbracket N_{\text{surface_argument}} \rrbracket^{w_0}$ (adapted from Martin, 2022)

This yields the observed pattern. If the covert nominal is 'Rolex', or any other nominal that denotes a relevant property P such that $\llbracket \text{watch} \rrbracket \not\subseteq P$, we get a non-vacuous ("non-contradictory") reading.

- (43) a. (40b) = $\llbracket \text{fake} \rrbracket^{w_0} (\lambda w. \llbracket \text{Rolex} \rrbracket^w) \cap \llbracket \text{watch} \rrbracket^{w_0}$
 b. \approx (schematically) **fake(Rolex) \cap watch**
 $= \lambda x. \text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{Rolex})) \wedge \neg \text{Rolex}(x) \wedge \text{watch}(x)$

A common assumption has been that reading (40a) is infelicitous because if the covert nominal happens to be 'orologio', we get a contradiction:

- (44) a. (40a) = $\llbracket \text{fake} \rrbracket^{w_0} (\lambda w. \llbracket \text{watch} \rrbracket^w) \cap \llbracket \text{watch} \rrbracket^{w_0}$
 b. \approx (schematically) **fake(watch) \cap watch**
 $= \lambda x. \text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{watch})) \wedge \neg \text{watch}(x) \wedge \text{watch}(x)$

I will show that readings like (40a) are in fact attested, and that this falls in line when

¹³This is the judgment reported in Cinque (2010) and Martin (2022) – we will see in section 4.3 that things are a bit more complicated. This reading is, although less salient, attested, and we will see that this is entirely expected.

we integrate the lexical entry proposed here with adequately constrained Partee pragmatic mechanisms.

Regardless, the ambiguity observed in English constructions such as ‘fake watch’, as in (45), is then a *bona fide syntactic* ambiguity, since the word order is the same for the two sources in English (Cinque, 2010; 2014). Reading (45a) is (at least) the product of a Direct Modification syntactic parse, while (45b) is the product of a Reduced Relative Clause syntactic parse.

- (45) 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

4.2 Is a fake *N* a *N*?

Recently, (Martin, 2022) argued for a DCS-style entry for ‘fake’ that lacks semantic negation. 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, Martin 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 privative inference.

This move is made on experimental grounds: in an experiment in which participants were asked whether a *Adj N* is a *N*, Martin found great variation both along privative adjectives and along different input nouns. What concerns us here is that, for some *Ns*, 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:

- (46) That is a fake watch.

It is true that I am not committing to this object not being a watch in this context. But crucially, I am absolutely committing to the object not being a Rolex.

Consequently, Martin’s experiment does not provide evidence that privatives do not negate their input noun; 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’, Direct Modification) 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. (39a); see Cinque, 2010; 2014). Second, as we have seen in section 3.1, the way the conjunction of $\neg P(x)$ and $\text{INTENDED}(x, \text{SEEM-LIKE}(x, P))$ enters the compositionality in iterated application of ‘fake’ delivers fine-grained and right predictions.

4.3 ‘That gun is fake’

Another question concerning semantic negation is raised by sentences such as (47), which are among Partee’s arguments for a subjective treatment of ‘fake’:

(47) That gun is fake.

The covert nominal taken by ‘fake’ seems to be ‘gun’, given that (47) is equivalent to (48):

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

The original problem of sentences like (47) 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 pragmatic principles can correctly account for these cases. In other words, while Partee’s theory is not spelled out enough to account for the compositional *meaning* of ‘fake’, it can account for why a Priv-Adj *N* can sometimes be referred to as an *N*. With a slight modification of her principles to adapt them to copulas, we can sketch a Partee-style account of sentences like (47) coupled with the meaning of ‘fake’ proposed in the present theory.

(49) **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*.¹⁴

(50) **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:

(51) *That gun is fake(gun)*
Try to interpret the sentence so that
 $\{x: \text{gun}(x)\} \subseteq \{x: \text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{gun})) \wedge \neg \text{gun}(x)\}$,
get an empty extension
 \Rightarrow NVC fails.

¹⁴I here take up Partee’s terminology; the notion of local context employed here is not necessarily the one we think of today.

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:

$$(52) \quad \{x: \text{gun}_{\text{broadened}}(x)\} = \{x: x \in \text{gun}_{\text{proper}}(x)\} \cup \{x: \text{fake-gun}(x)\} \\ = \{x: \text{gun}(x)\} \cup \{x: \text{INTENDED}(x, \text{SEEM-LIKE}(x, \text{gun})) \wedge \neg \text{gun}(x)\}$$

Now, of course in a sentence of the form ‘*D N* 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 $\llbracket N \rrbracket \not\subseteq \llbracket N_1 \rrbracket$, no reasoning like the one above is triggered. To wit, (53) under the reading in (53a) is not contradictory to begin with, so NVC does not fail.

- (53) This watch is fake.
 a. This watch is fake (as a rolex)

This leads us to reassess Italian non-copular modification. Given this discussion, there is no reason why Partee’s reasoning could not rescue the truly privative reading for the RRC modification source. Recall (39) and (40), repeated below in (54) and (55).

- (54) **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
- (55) **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

It would be surprising if Partee’s mechanism could save copular sentences like (47) but not an RRC reading like (55a). However, as already anticipated at the beginning of this section, there are good reasons to think that (55a) is in fact generally possible, but needs context to be accessed. This is because for (55a) we need to go through Partee’s pragmatic adjustment to resolve the initial contradiction, while (55b) is, so to say, good as is.

And indeed, the relevant judgment is more available than the literature suggests: according to the six informants I consulted, (55a) is completely accessible in the given context.¹⁵

- (56) 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.

¹⁵As a native speaker of Italian, I myself share these judgments.

‘...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 (55a) 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 only accessible even in the RRC source. For instance, absent a strong context, (57a) is the only reading of (57)

- (57) finestra falsa
window fake
a. an object that is intended to look like a window but isn’t a window

All of this suggests that in RRC in Romance, in an NP like $\llbracket \text{finestra falsa}(\text{finestra}_{\text{covert}}) \rrbracket$ under a reading like (57a), the overt occurrence of ‘finestra’ gets broadened as a result of Partee’s reasoning, as illustrated below.

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

Notice that Non-Vacuity is only ever violated when the extension is formally empty: that is, when ‘fake’ first takes as an input a given covert nominal N , and then is intersected with the overt occurrence of N . This yields an empty intersection no matter what N is. This correctly predicts that we can never violate Non-Vacuity (and therefore give up Head Primacy) when ‘fake’ directly combines with the overt noun via Functional Application – regardless how unlikely the resulting meaning is. In other words, Partee’s reasoning never arises when ‘fake’ combines via Direct Modification: this correctly predicts that in Italian Direct Modification, the privative inference is systematic and obligatory.

Summing up, then, the full picture of the ways in which ‘fake’ composes with nouns are in the following table.

Syntactic parse	Semantic composition	Inference “ x is not a watch”?	Pragmatic Partee-broadening of $\llbracket \text{watch} \rrbracket$? (“Give up Head Primacy to save Non-Vacuity”)
Direct Modification	(Intensional) Functional Application: $\llbracket \text{fake} \rrbracket_{\{(s,(e,t)),(e,t)\}}^{w_0} (\lambda w. \llbracket \text{watch} \rrbracket_{(e,t)}^w)$	Yes, systematically	No
Restricted Relative Clause	(Intensional) Functional Application with a covert argument & Predicate Modification with the noun: $\llbracket \text{fake} \rrbracket_{\{(s,(e,t)),(e,t)\}}^{w_0} (\lambda w. \llbracket c \rrbracket_{(e,t)}^w) \cap \llbracket \text{watch} \rrbracket_{(e,t)}^{w_0}$	Only if $\llbracket \text{watch} \rrbracket^{w_0} \subseteq \llbracket c \rrbracket^{w_0}$	Only if $\llbracket \text{watch} \rrbracket^{w_0} \subseteq \llbracket c \rrbracket^{w_0}$, since then $\llbracket \text{fake} \rrbracket^{w_0} (\lambda w. \llbracket c \rrbracket^w) \cap \llbracket \text{watch} \rrbracket^{w_0} = \emptyset$

Table 1: readings of ‘fake watch’.

5 Similarity in ‘fake’

In this section, I discuss two aspects pertaining to the appeal to similarity of the present theory. In 5.1, I illustrate how the semantics of ‘fake’ can be kept one-dimensional, by appealing to Guerrini’s (2021) recent one-dimensional account of ‘seem like’ and ‘look like’ constructions. In section 5.2, I discuss a potentially simpler theory only making use of deceptive intentions, where e.g. a fake lawyer is someone who wants others to believe they are a lawyer, and nothing else. I show that such a theory doesn’t work, and that we indeed have to embed similarity in ‘fake’.

5.1 One-dimensional ‘fake’ with one-dimensional similarity

In Guerrini (2022), it is argued that constructions such as the one involved in sentence (59a) are best captured in terms of property sharing, i.e. as in (59b).

- (59) 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:

- (60) 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.

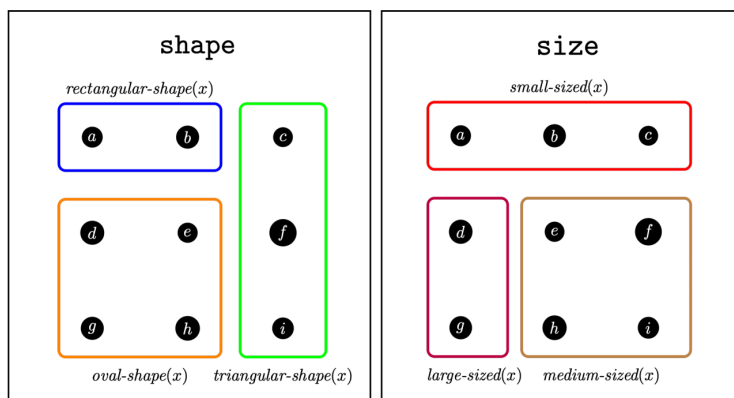


Figure 5: “With respect to shape and size, object a is like object b.”

The proposal, roughly, is that x is like y iff for a relevant set of properties \mathbf{D} , x has the same value as y .

$$(61) \quad \llbracket \text{like} \rrbracket^{w, \mathbf{D}} = \lambda y. \lambda x. \forall P_{\langle s, \langle e, t \rangle \rangle} \in \mathbf{D}_{\langle \langle s, \langle e, t \rangle \rangle, t \rangle}. P(x)(w) = P(y)(w)$$

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 (62a) can be captured as in (62b).

- (62) a. ‘Bob looks French’
 b. $\forall w' \in V(j, w). \text{French}(bob)$
 c. ‘at all worlds compatible with the visual experience of the judge j at w , $\text{French}(bob)$ ’

Extending Rudolph’s semantics of ‘look $P_{\langle s, \langle e, t \rangle \rangle}$ ’, then, a sentence like ‘Bob looks like Carl’ is analyzed as in (63).

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

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 ’.

- (64) 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). This is because the indefinite goes into the restriction of the generic quantifier, a silent quantificational adverb postulated in Krifka (1995):

- (65) $\llbracket (64)_{\text{GENERAL}} \rrbracket = \text{GEN}_x[x \text{ is a lawyer}][\text{John is like } x]$

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.¹⁶

- (66) a. A bird is flying. (\exists indefinite)
 b. \approx A typical bird is flying.

¹⁶I illustrate with ‘look like a N ’ constructions because they sound more natural, but the same results apply to ‘be like a N ’ constructions.

- (67) a. A bird flies. (GEN **indefinite**)
 b. \approx A typical bird flies.
- (68) a. John looks like a lawyer I know. (SPECIFIC)
 b. $\not\approx$ John looks like a typical lawyer I know.
- (69) 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 (70) and (72), just like more vanilla characterizing sentences, cf. (71) and (73).

- (70) a. John looks like a British judge. (GENERAL)
 b. $\not\approx$ John looks like a judge.
- (71) a. A British judge wears a wig.
 b. $\not\approx$ A judge wears a wig.
- (72) a. John looks like a bird. (GENERAL)
 b. $\not\approx$ John looks like a penguin.
- (73) a. A bird flies.
 b. $\not\approx$ A penguin flies.

Third, constructions such as ‘like a lawyer or a judge’ display scope ambiguity. The narrow reading (74c) 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’.

- (74) 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 subtriggering effects that Carlson (1981) discovered for vanilla generic sentences with ‘someone’.

- (75) Someone should be punctual.
 a. Existential
 b. #Generic
- (76) Someone who respects others should be punctual.
 a. Existential
 b. Generic
- (77) John looks like someone.
 a. Existential
 b. #Generic
- (78) John looks like someone who respects others.
 a. Existential

b. Generic

This all argues for the conclusion that ‘like’-constructions feature inherent genericity. At this point, we are equipped to plug in the semantics from Guerrini (2022) for ‘seem like a gun’ into our lexical entry for ‘fake’. As argued above, we want a fake gun to have the following properties:

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

Then we are equipped to write an entry that is explicit with respect to intensionality and to the analysis of similarity embedded by ‘fake’:

- (79) a. $\llbracket \text{fake} \rrbracket^{w,D} = \lambda Q_{\langle s, \langle e, t \rangle \rangle} . \lambda x .$
 $\text{INTENDED}(x,$
 $\quad \forall v \in R_{\text{look}}(j, w).$
 $\quad \text{GEN}_{y,u} [u \in R_{\text{gen}}(v) \wedge Q(y)(u)] [\forall P \in \mathbf{D}_{\langle \langle s, \langle e, t \rangle \rangle, t \rangle} . P(x)(u) = P(y)(u)]$
 $\quad) (w)$
 $\quad \wedge \neg Q(x)(w)$
- b. Paraphrase:
“ x is intended in w so that
at all worlds v compatible with the experience of j in w ,
 x shares relevant properties P with typical instances y of Q
and x is not a Q in w ”

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 this 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.

Before moving on, let me address an issue raised by an anonymous reviewer, who notes that, as things stand, extending Rudolph’s analysis to ‘look like’ makes some inadequate predictions. In a nutshell, taking the semantics in (79) without further assumptions, we would predict that for John to look like Bob it would be enough for the speaker to have *indirect* perceptual evidence that John shares with Bob some properties, which may be non-perceptual. This is a problem for an analysis of similarity: if I say that John looks like Bob, I mean I have *direct* perceptual evidence that they share *visual* properties. Consequently, plugging this analysis into a lexical entry for ‘fake’ would also be problematic. If I say that an object is a fake gun, I mean that the builder intended it to give me *direct* perceptual evidence of a gun via some *perceptual* properties.

I think that this concern can be addressed under this perspective. The issue boils down to the distinction between phenomenal and epistemic uses of appearance verbs like ‘seem like’, ‘look like’, and ‘taste like’. In a nutshell, if an appearance verb is epistemically flavored, it refers to the evidence that is epistemically compatible with the content of perception. This includes evidence from perception that indirectly

points to a proposition being the case. This is the case, for instance, of sentence-embedding ‘taste like’: sentence (80a) is compatible with the context given in (80).

- (80) **Context:** the speaker doesn’t like wine, but is very competent at objectively judging the quality of wine. They don’t directly experience the wine they are drinking as good, but have indirect gustatory evidence that the wine is of a quality that would be taken to be good by wine experts.
- a. This wine tastes like it’s good. *epistemic*

On the other hand, when an appearance verb is phenomenally flavoured, it refers to direct perceptual evidence. For instance, predicate-embedding ‘taste’ is phenomenal: sentence (81a) is not compatible with the context from above.

- (81) **Context:** same as (80).
- a. ?? This wine tastes good. *phenomenal*

On these grounds, Rudolph argues that sentence-embedding ‘taste like’ comes with an epistemically flavored accessibility relation, while predicate-embedding ‘taste’ comes with a phenomenally flavored accessibility relation. In Guerrini (2022a), it is shown that a similar distinction extends to ‘look’ and ‘look like’: sentence-embedding ‘look like’ is epistemically flavored, while predicate-embedding ‘look’ can be phenomenal.

- (82) **Context:** It is very clear that Bob is not French, and much indirect visual evidence points to him not being French: for instance, at some point he dropped his passport and we saw it was a German one. Suppose now that he puts on a Basque beret.
- a. Bob looks French. *phenomenal*
- b. ??Bob looks like he’s French. *epistemic*

This extends to similarity constructions:

- (83) **Context:** same as (82).
- a. Bob looks like a Frenchman.
- b. ??Bob looks like he is a Frenchman

On these grounds, in Guerrini (2022a) it is required that the accessibility relation embedded by expressions such as ‘seem like John’/ *seem like a lawyer* and ‘look like John’/ ‘look like a lawyer’ be phenomenal. This makes sure that the similarity is directly experienced, and in the relevant modality (e.g., the visual modality for ‘look like’). It also makes sure that if I say that Bob looks like John, I mean that they share *visual* properties: I cannot have direct visual evidence of any other property.

We can require the same for the similarity relation in ‘fake’. This makes sure (i) that the properties that a fake gun is intended to share with a gun are perceptual, and (ii) that it is intended that the judge *directly* experience this.

5.2 Discarding an obvious alternative: ‘fake’ grounded in deception

Let us now consider a simple, prima-facie appealing theory that grounds ‘fake’ in a combination of well-understood propositional attitudes without any appeal to similarity. 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:

- (84) a. fake =
 $\lambda x.\lambda P.\text{INTENDED}(-\text{TO-CAUSE})[x, \text{BELIEVE}(j, P(x))] \wedge \neg P(x)$
 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:

- (85) 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:

- (86) # John’s Toshiba radio is a fake bomb.

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

- (87) 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.

- (88) 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’.¹⁷ Notice that this discussion barely

¹⁷ 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 (88) 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.

rules out a theory that is *exclusively* based on deceptive intentions. It may be still necessary to combine deceptive intentions with similarity. In this direction goes a remark made by an anonymous reviewer, who points out that my semantics predicts that the painting described in (89) should qualify as fake, contrary to intuition:

- (89) A copy of a famous painting painted by an art student as a way to learn and better appreciate the artist’s techniques.

As pointed out by the reviewer, this argues for a more complex theory, where an appeal to similarity *coexists* with an appeal to deceptive intentions. “Fake *N*”, then, is true of an object if there is an intention to cause a false belief that the object is an *N* by virtue of the object’s resemblance to an *N*. For simplicity, I will stick to the simpler proposal which only appeals to similarity, but implementing the solution that combines deception with similarity should be straightforward: deception is a well-understood attitude insofar as it boils down to causing a false belief. The *in-virtue-of* relation may perhaps be spelled out as a causal relation between the target appearance of the object – that it look like a *N* – and the goal of the deceptive intention – that someone think that the object is a *N*.

6 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 ‘fake’ is fact 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’.

As summarised in Table 1 (reported again below), ‘fake’ can combine either (i) directly with the noun or (ii) with a covert argument *c*, in both cases via (Intensional) Functional Application.

-
- 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. This seems to constitute further evidence that the ‘seem like’ embedded by ‘fake’ is phenomenally flavoured (cf. section 5.1)

Syntactic parse	Semantic composition	Inference “ <i>x</i> is not a watch”?	Pragmatic Partee-broadening of $\llbracket \text{watch} \rrbracket$? (“Give up Head Primacy to save Non-Vacuity”)
Direct Modification	(Intensional) Functional Application: $\llbracket \text{fake} \rrbracket_{\langle (s, \langle e, t \rangle), \langle e, t \rangle \rangle}^{w_0} (\lambda w. \llbracket \text{watch} \rrbracket_{\langle e, t \rangle}^w)$	Yes, systematically	No
Restricted Relative Clause	(Intensional) Functional Application with a covert argument & Predicate Modification with the noun: $\llbracket \text{fake} \rrbracket_{\langle (s, \langle e, t \rangle), \langle e, t \rangle \rangle}^{w_0} (\lambda w. \llbracket c \rrbracket_{\langle e, t \rangle}^w) \cap \llbracket \text{watch} \rrbracket_{\langle e, t \rangle}^{w_0}$	Only if $\llbracket \text{watch} \rrbracket^{w_0} \subseteq \llbracket c \rrbracket^{w_0}$	Only if $\llbracket \text{watch} \rrbracket^{w_0} \subseteq \llbracket c \rrbracket^{w_0}$, since then $\llbracket \text{fake} \rrbracket^{w_0} (\lambda w. \llbracket c \rrbracket^w) \cap \llbracket \text{watch} \rrbracket^{w_0} = \emptyset$

In the latter case, the $\llbracket \text{fake} \rrbracket^{w_0} (\lambda w. \llbracket c \rrbracket^w)$ complex subsequently combines via Predicate Modification with the noun. This allowed us to capture ambiguities such as the one observed for ‘fake watch’, which could denote a “fake-as-watch object” or a “fake-as-Rolux watch”. We saw that the extension denoted by the noun is sometimes broadened pragmatically, as in Partee’s framework. However, crucially, this only happens when (i) ‘fake’ combines with the noun via the latter type of composition and (ii) *c* is such that the intersection of $\llbracket \text{fake} \rrbracket^{w_0} (\lambda w. \llbracket c \rrbracket^w)$ with the extension denoted by the noun would be vacuous (as, for instance, ‘fake watch’ when $c = \lambda x. \text{watch}(x)$).

From this perspective, the entailment pattern to which an adjective gives rise is a surface phenomenon that results from the interaction between its mode of composition, its core lexical meaning, and Parteean rescuing mechanisms. The most helpful classification of adjectives, then, is likely one that is based on mode of composition, rather than on simple entailment patterns as is standard. To end the paper with an avenue for future research, let me tentatively suggest that such a classification might be based on the empirical tests we used for ‘fake’. These are the following:

- (a) $\langle e, t \rangle$ -test (Copula test)
Can the adjective occur in copular constructions, which only allow for expressions of type $\langle e, t \rangle$?
- (b) $\langle \langle s, \langle e, t \rangle \rangle, \langle e, t \rangle \rangle$ -test (Direct Modification test)
Can the adjective occur in Direct Modification, a syntactic position that only allows for expressions of type $\langle \langle s, \langle e, t \rangle \rangle, \langle e, t \rangle \rangle$?
- (c) ‘Implicit argument’ test (Context-sensitivity test)
If the adjective is of $\langle \langle s, \langle e, t \rangle \rangle, \langle e, t \rangle \rangle$ type (i.e., if the answer to (b) is ‘yes’), can it modify a contextually salient property that is different from the one denoted by the noun?

This classification divides adjectives into three classes. First, ‘fake’ and other adjectives like ‘counterfeit’ and ‘good’, which have the two modes of composition described above and summarised in Table 1, and as a result can occur both in copulas and in Direct Modification.¹⁸ Let us call this the class of “context-sensitive functional adjectives”, and dub it Class 1.

A second class is constituted by adjectives such as ‘former’, ‘alleged’, ‘potential’ which, like Class 1 adjectives, can occur in Direct Modification:

- (90) un potenziale avvocato
a potential lawyer

¹⁸On ‘good’, see von Stechow & Heim, 1999, who clearly show its context-sensitivity. Cf. also Cinque (2010) and Martin (2022).

However, unlike Class 1 adjectives, these adjectives (i) are not context-sensitive (cf. (91)) and (ii) cannot occur in copulas (cf. (92)). They are clearly of type $\langle\langle s, \langle e, t \rangle \rangle, \langle e, t \rangle\rangle$ like ‘fake’, but unlike ‘fake’ cannot saturate their $\langle s, \langle e, t \rangle \rangle$ argument via a covert, contextually provided argument – whence their lack of context-sensitivity. This also means that they cannot come to denote $\langle e, t \rangle$ functions if they stand alone (whence the ungrammaticality in copulas).

- (91) a former lawyer
 a. a former-as-lawyer person (someone who used to be a lawyer)
 b. *a former-as-c lawyer (a lawyer that used to be a c, with c a contextually salient property).

- (92) *That lawyer is former.

Let us call this the class of “non-context-sensitive functional adjectives”, and dub it Class 2.

A third class is constituted by adjectives that can occur in copulas and cannot occur in syntactic positions that specialize for Functional Application. Good examples of this class are nationality adjectives, for instance ‘Greek’, but also adjectives like ‘hexagonal’.

- (93) *Un greco avvocato
 A greek lawyer

- (94) That lawyer is Greek.

Let us call this the class of “pure intersectives”, and dub it Class 3.

This gives rise to the following typology:

Class	Examples	Functional Application with the input noun	Functional Application with a covert argument + Predicate Modification with the input noun	Only Predicate Modification with the input noun
Class 1 Context-sensitive functional adjectives	<i>fake, counterfeit, good</i>	✓	✓	✗
Class 2 Non-context-sensitive functional adjectives	<i>former, potential, alleged</i>	✓	✗	✗
Class 3 Pure intersectives	<i>Chinese, Greek, hexagonal</i>	✗	✗	✓

Table 2: a tentative and possibly non-exhaustive typology of adjectival composition.

This makes it possible to view a number of traditional issues under a new light, such as for instance the difference between ‘fake’ and ‘former’ (see for instance Morzycki, 2016 for an explicit statement of this question): the central challenge becomes understanding why Class 1 adjectives, but not Class 2 adjectives can combine with an implicit argument.

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**. MIT press.
- 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. *Noûs*, pp. 311–335.
- Del Pinal, G. (2015) Dual content semantics, privative adjectives, and dynamic compositionality. *Semantics and Pragmatics* **8**, 7–1–53.
- 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.
- von Stechow, K. & Heim, I. (1999) More on lousy teachers and beautiful dancers. Lecture notes, MIT.
- Franks, B. (1995) Sense generation: A “quasi-classical” approach to concepts and concept combination. *Cognitive science* **19**, 441–505.
- Guerrini, J. (2022a) Genericity in similarity. Manuscript, Ecole Normale Supérieure.
- Guerrini, J. (2022b) Like a N constructions and genericity. *Proceedings of the ESSLLI 2022 Student Session*.
- 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.
- Jespersen, B., Carrara, M. & Duži, M. (2017) Iterated privation and positive predication. *Journal of Applied Logic* **25**, S48–S71.
- 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) *The Generic Book*. The University of Chicago Press.
- Laserson, 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. Stanford: CSLI Publi.
- Morzycki, M. (2016) *Modification*. Cambridge University Press.
- 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.
- Tessler, M.H. & Franke, M. (2019) Not unreasonable: Why two negatives don't make a positive. Manuscript.