

Dual Content Semantics, privative adjectives, and dynamic compositionality *

Guillermo Del Pinal
Columbia University &
Ruhr-Universität Bochum

Submitted 2014-04-01 / 2014-06-02 / Revision received 2014-07-16 / Accepted 2014-09-01 / Final version received 2014-11-14 / Published 2015-03-27

Abstract This paper defends the view that common nouns have a dual semantic structure that includes extension-determining and non-extension-determining components. I argue that the non-extension-determining components are part of linguistic meaning because they play a key compositional role in certain constructions, especially in privative noun phrases such as *fake gun* and *counterfeit document*. Furthermore, I show that if we modify the compositional interpretation rules in certain simple ways, this dual content account of noun phrase modification can be implemented in a type-driven formal semantic framework. In addition, I also argue against traditional accounts of privative noun phrases which can be paired with the assumption that nouns do not have a dual semantic structure. At the most general level, this paper presents a proposal for how we can begin to integrate a psychologically realistic account of lexical semantics with a linguistically plausible compositional semantic framework.

Keywords: compositionality, conceptual content, lexical semantics, lexical decomposition, atomism, privative adjectives, modification

* I am extremely grateful to Daniel Rothschild for extensive comments and discussions of various drafts of this paper. The paper also benefited enormously from the sharp and constructive comments of Louise McNally and three anonymous referees. I would also like to thank Akeel Bilgrami, Barbara Partee, Brian H. Kim, Marco J. Nathan, Achille Varzi, Anubav Vasudevan and members of the UCL Linguistics Department, where an earlier version of this paper was presented.

1 Introduction

Are lexical items the atoms of meaning or do they have internal semantic structure? If the latter, what kind of structure — definitions, prototypes, some sort of holistic structure? These questions have received a lot of attention in philosophy, linguistics and psychology, in part because of their intrinsic interest, but mostly because they are intimately connected to questions about nativism, the structure of lexical concepts, the nature of language acquisition, and the relation between language and thought.

We will explore questions about lexical decomposition by focusing on noun phrases (NPs).¹ We will try to uncover the semantic structure of nouns (Ns) by examining their compositional behavior in *privative NPs* such as *fake gun* and *counterfeit document*. This is a systematic and productive class of NPs which have privative adjectives (Adjs) as main modifiers and artifact terms as head Ns. Privatives have puzzled theorists due to their unique compositional behavior. For example, privatives behave differently from the intersectives and the subsectives. Example (1) displays the characteristic inference pattern of intersectives:

- (1) a. $\frac{x \text{ is a red gun}}{x \text{ is a gun}}$ b. $\frac{x \text{ is a red gun}}{x \text{ is red}}$

Example (2) displays the characteristic inference pattern of subsectives, where # marks incorrect inferences:

- (2) a. $\frac{x \text{ is a good gun}}{x \text{ is a gun}}$ b. $\frac{x \text{ is a good gun}}{\# x \text{ is good}}$

Privatives have inference patterns like (1b) but not like (1a) or (2a):

- (3) a. $\frac{x \text{ is a fake gun}}{\# x \text{ is a gun}}$ b. $\frac{x \text{ is a fake gun}}{x \text{ is a fake}}$

Privatives also differ from two important classes of non-subsective Adjs, the plain non-subsectives (mainly constituted by certain epistemic modals) and the temporal and alethic modals. Example (4) displays the characteristic inference pattern of plain non-subsectives:

¹ Recent discussions of lexical decomposition tend to focus on verbs. Although some theorists defend atomism (Fodor 1998, Fodor & Lepore 2002), the emerging consensus is that most verbs have internal semantic structure, with much disagreement about the specifics of this structure (Wunderlich 2012, Harley 2012, Hinzen 2012, Pustejovsky 2012).

- (4) a. $\frac{x \text{ is an alleged criminal}}{\# x \text{ is not a criminal}}$ b. $\frac{x \text{ is an alleged criminal}}{\# x \text{ is a criminal}}$

Example (5) displays the characteristic inference pattern of temporal modals:

- (5) a. $\frac{x \text{ is a former/future criminal}}{\# x \text{ is not a criminal}}$ b. $\frac{x \text{ is a former/future criminal}}{x \text{ was/will be a criminal}}$

Unlike plain non-subjectives and temporal modals, privatives allow inferences like (4a) and (5a).² In addition, they do not allow inferences like (5b):

- (6) a. $\frac{x \text{ is a fake gun}}{x \text{ is not a gun}}$ b. $\frac{x \text{ is a fake gun}}{\# x \text{ was/will be a gun}}$

Note, finally, that from *x is a former/future criminal* you cannot infer anything about what *x* does or seems like; but from *x is a fake gun* you can infer something about what *x* seems like or about the intention for which *x* was created: a fake gun either seems like a gun or was created to seem like a gun.³

The inferential behavior of privative NPs presents substantial challenges for compositional semantic theories (Franks 1995, Partee 2007, 2010, Recanati 2010); some theorists even argue that they force us to abandon that framework (Lakoff & Johnson 1980, Coulson & Fauconnier 1999). In contrast, I will argue that privative NPs are not a threat to compositional formal semantics. What they do suggest, however, is that we abandon the assumption that most common Ns are semantically atomic.

In what follows, I defend a compositional account of privative NPs which can be implemented in a type-driven formal semantic framework, once this framework is modified in ways that I will present in detail in Section 3. This account is based on two ideas about the semantic structure of Ns. Assume *e* stands for an arbitrary common N (e.g., an artifact or natural kind term).

² One might question (5a): it might seem that from *x is a former/future criminal* you can infer that, at the moment the assertion is made, *x is not a criminal*. However, this is not the case. It is easy to think of examples in which we assert that *x* is a former/future criminal and know that *x* is also currently a criminal.

³ There have been various attempts to reduce the subclasses of Adjs, and especially to try to lump as many Adjs as possible into the intersectives (Heim & Kratzer 1998, Partee & Kamp 1995, Partee 2007, 2010, Recanati 2010). Here we are mainly concerned with privatives, so we will discuss attempts to reduce these classes only when they bear on the privatives, especially when they attempt to reduce them to the subjectives or intersectives.

- i. e has a complex semantic structure. Part of this structure determines e 's extension (call this its *E-structure*), and part of it does not determine e 's extension (call this its *C-structure*).
- ii. e 's C-structure includes, at least, representations of perceptual features, functional features and genealogical features related to e , and it determines e 's compositional contribution to certain complex constructions.

According to this “DUAL CONTENT” account, to fully grasp the literal meaning of artifact and other types of lexical terms such as natural kinds, speakers need to grasp both their E-structure and their C-structure. The E-structure plays a role, for instance, in predicative uses of Ns. The role of C-structure is less straightforward, but no less important. I will argue that to explain why speakers have a productive and systematic competence with privative NPs we need to assume that Adjs such as *fake* operate on the C-structure of the Ns they modify. I will also point out other types of complex expressions in which C-structure seems to play a crucial compositional role.

DUAL CONTENT is based on a non-atomistic theory of lexical meaning, but it should be distinguished from the two most famous non-atomic theories: prototype and definitional theories. Unlike pure prototype theories, in DUAL CONTENT the meaning of lexical terms includes E-structure, an extension-determining atomic component. In addition, the C-structure of terms is, as we will see, richer than the structures assumed by most prototype theories. Unlike definitional theories, in DUAL CONTENT C-structure often determines neither the E-structure nor the extension of terms. C-structure is part of the meaning of terms because it is a necessary part of what language-users need to grasp to have full (literal) linguistic competence.

It might seem that I'm using the notion of linguistic meaning in an idiosyncratic way; but my usage is appropriate in the context of debates about lexical decomposition. As I use the term, the “meaning” of a lexical item e includes that which e contributes to the determination of the meaning of complex expressions of which e is a constituent. More generally, meaning is what speakers grasp, or part of what they grasp, when they learn how to “use” e , which includes the systematic and productive use of e in complex expressions such as privative NPs. Some prominent theorists also assume that the meaning of e completely coincides with the extension-determining semantic features of e . I take this claim as a hypothesis about linguistic

meaning and not as part of the meaning of “linguistic meaning”, and part of the task of this paper is to undermine it.

2 DUAL CONTENT and privative NPs: Philosophical motivations

According to DUAL CONTENT, we can account for the inference pattern of privative NPs if we assume

- i. that the head Ns have a complex semantic structure which includes E-structure and non-extension-determining C-structure, and
- ii. that privative Adjs operate in certain ways on the C-structure of the modified Ns.

Before developing this account in Section 3, I want to discuss its two basic ideas, in part because the details of the implementation can be intricate and only if we are clear about the sort of general account we are trying to develop will the technical decisions make sense.

The idea that some types of Ns have a complex semantic structure which includes E-structure (which determines their range of correct application) and C-structure (which includes associated beliefs), and that the reason why we posit C-structure as part of the meaning of Ns is to explain certain aspects of our linguistic competence, was originally defended by [Putnam \(1970\)](#).

When Putnam first defended this idea, versions of the definitional theory of lexical terms and concepts were quite popular. According to definitional theories, it usually holds that if e is a N the meaning of e can be given by specifying a conjunction of properties $P_1 \dots P_n$ that (mature speakers know) define or determine e 's extension. For each P_i , the statement e *has the property* P_i is an analytic truth. Putnam argued that these theories are clearly incorrect for a rather simple reason: for any property that supposedly “defines” an artifact or natural kind for a mature speaker, one can (as a representative mature speaker) imagine a counterfactual situation in which something does not have that property and yet still correctly falls under the kind. At around the same time, Kripke was developing a similar objection to definitional theories of natural kinds ([Kripke 1980](#)).

Putnam’s criticism of definitional theories is widely accepted by philosophers and linguists, but it is sometimes ignored or too lightly discussed by psychologists. Among theorists who take Putnam’s view seriously, most have abandoned the idea that the meaning of terms can be represented as non-

atomic structures with semantically significant parts. One famous example is Fodor (1998), who argues that the failure of definitional theories and their descendants — essentially for the reasons Putnam originally presented — forces us to accept atomistic theories of lexical terms and concepts. Theorists like Fodor, who on the strength of Putnam’s criticisms of definitional theories see atomism as the only route, remain attached to the idea that there can be nothing more to the meaning of an expression than its E-structure.

However, there is another way to proceed from Putnam’s criticisms of definitional theories, somewhat lost in subsequent discussions, but actually suggested by Putnam himself: hold that most Ns have a complex semantic structure, but deny that the role of that structure is only to determine their extension. Putnam suggests that the meaning of an arbitrary natural-kind term *e* is a complex involving two representational dimensions: a representation that *e* stands for a natural kind, and a representation of the *stereotype* associated with *e* — that is, a structured set of properties which speakers associate with normal members of the kind *e*. According to Putnam, *e*’s stereotype does not determine *e*’s extension, but it is part of its core meaning and not just a bit of associated encyclopedic knowledge because it is part of what speakers need to know in order to have full linguistic competence with *e*.

Putnam thinks that the view that *e*’s stereotype is part of its meaning becomes clear when we consider questions about language acquisition — about what we learn when we learn the meaning of *e* such that knowing that enables us to use *e*. According to Putnam, most language theorists fail to be impressed by the fact that, in many cases, the use of Ns can be taught by presenting speakers with the word’s stereotype. More precisely:

The hypothesis is that there are, in connection with almost any word (not just “natural kind” words), certain core facts such that (1) one cannot convey the normal use of the word (to the satisfaction of native speakers) without conveying those core facts, and (2) in the case of many words and many speakers, conveying those core facts is sufficient to convey at least an approximation to the normal use. In the case of a natural kind word, the core facts are that a normal member of the kind has certain characteristics, or that this idea is at least the stereotype associated with the word. (Putnam 1970)

Putnam's basic idea is that when you show a speaker the stereotype — the set of core facts — associated with a term, you enable the speaker to competently use the term, and you cannot achieve this in any other way. In other words, Putnam's claim is that part of what we learn when we learn the meaning of a kind term is the associated stereotype, and this knowledge is what enables us to competently use the term. Still, as revealed by their response to various counterfactual scenarios, speakers know that the associated stereotypes and core facts do not, in general, determine the range of correct application of kind terms. This is why we must also assume that lexical terms include an extension-determining E-structure, which C-structure in some ways traces, but does not fully determine.⁴

Putnam's idea that the semantic structure of lexical terms can be divided into E-structure and C-structure, and that C-structure is part of the literal meaning of terms because it is an essential part of our linguistic competence, is a move in the right direction, fully embraced by DUAL CONTENT. However, Putnam's account is incomplete in two important respects, which DUAL CONTENT tries to rectify. First, we need a detailed example of a *type* of linguistic competence that essentially involves C-structure. Otherwise one can agree with Putnam that we learn the meaning of terms *via* their associated stereotype (Putnam's C-structure), but deny that these stereotypes are part of their meaning.⁵ Second, we need a more detailed account of the C-structure of terms, since it is unclear what information Putnam would include under his "core facts".

This is where privative NPs come into the picture. They provide us with a concrete example of a kind of (literal, non-pragmatic) general linguistic competence — the capacity to systematically and productively use privative NPs — that seems to essentially involve the C-structure of terms. In addition, privative NPs yield important insights about the details of C-structure, including that it seems to contain more information than just a representation of the associated stereotype.

⁴ For a detailed discussion of E-structure, C-structure, and their relation, see Section 7 below. For now, the basic idea is just that E-structure is a semantically atomic component of the meaning of lexical terms which, via some primitive relation, determines their extension or range of correct application.

⁵ For example, Margolis (1998) and Fodor (2008) acknowledge that stereotypes might be "sustaining mechanisms" involved in the acquisition and even application of concepts, but still deny that stereotypes are constituents of concepts. This same view was also defended in a very influential paper by Rey (1983).

The idea that we can uncover the semantic structure of common Ns by investigating compositions involving privative Adjs is not new (Franks 1995, Coulson & Fauconnier 1999). One of the first and most insightful presentations of this idea is found in Lakoff & Johnson 1980. They agree that the compositional behavior of privative NPs such as *fake gun* is adequately captured by the inference patterns presented in Section 1. So any adequate treatment of privative Adjs such as *fake* has to explain why, for instance, from knowing that x is a fake gun we can infer

- i. that x is not a gun, and
- ii. that x was created with the intention that it only seem like a gun,

but not that it can be used to shoot. Lakoff and Johnson argue that to explain (i) and (ii) — generalized to include other artifact kind terms — we need a non-atomistic account of the lexical semantics of the modified Ns.

Lakoff and Johnson's account of privative NPs is based on the idea that privative Adjs operate over the complex semantic structure of artifact terms. In the case of *fake gun*, they propose that *gun* has a complex semantic structure which includes at least three dimensions: perceptual, functional, and genealogical representations. They then take the semantic potential of *fake* to be that of an operation over this complex semantic structure: the operation negates both the function and the genealogical representations of *gun* (since fake guns are not supposed to and were not made to function like real guns, i.e., to shoot), keeps the perceptual representation unmodified (since fake guns are supposed to look like guns), and takes this perceptual representation as the value of the genealogical dimension (since fake guns are made to look like guns). This last point is important, since something made to be a gun which doesn't function like a gun (because it is old or doesn't work) is a malfunctioning gun but not a fake gun. Fake guns are artifacts made with the intention that they be fake guns.

There is an important difference between DUAL CONTENT and the account of lexical decomposition defended by Lakoff & Johnson (1980). Both accounts agree that, in general, lexical terms do *not* decompose into sets of necessary and sufficient conditions. For example, we may categorize something as a gun even if it doesn't function like a gun or doesn't look like a typical gun. In other words, we accept that there can be guns that don't shoot, or guns that don't look like typical guns. Indeed, we can imagine cases in which something that was not made to be a gun nevertheless is used and functions

so well for shooting that we would categorize it as a gun. So the perceptual, functional, and genealogical features associated with *gun* are independent of its E-structure.⁶ At the same time — as I will argue in detail below — there are constructions in which functional or genealogical information is introduced into the E-structure for complex expressions. Although Lakoff and Johnson do not explicitly argue against this possibility, it cannot be captured in a framework, such as theirs, which does not preserve the distinction between E-structure and C-structure.

DUAL CONTENT, which we will now present in detail, takes Putnam's insights about the basic division of internal lexical semantic structure, and Lakoff and Johnson's insights about the operations performed by privative Adjs, and develops and implements them in a way that should be attractive to formal semantic theorists.

3 DUAL CONTENT and privative NPs: Implementing the account

The proposal that artifact terms have a semantic structure which includes E-structure and C-structure, and that privative Adjs are operators that modify that complex structure, can be implemented in a number of ways without affecting the basic idea behind DUAL CONTENT. The implementation I will present is attractive because it is an extension of the type-driven formal semantic framework of Heim & Kratzer 1998, widely used by linguists and philosophers. I suspect that the main reason why many theorists don't accept lexical decomposition, especially into E-structure and C-structure, is simply that our traditional formal semantic theories are not sensitive to the C-structure of lexical terms. However, only after working out particular accounts can we really determine what complications would result from the required changes. The following account illustrates how a type-driven compositional theory can access the C-structure of terms, and shows that the required changes do not substantially complicate the general framework.

To begin, let us introduce the format we will use to represent lexical entries for Ns, which according to DUAL CONTENT include E-structure and C-structure. The C-structure of a lexical item *e* represents some common properties of *e*'s extension which, as we argued above following Putnam, Lakoff and Johnson, go beyond what is captured by the stereotype associated with *e*. We will assume that this rich C-structure is organized around four semantic dimensions or *qualia roles* (Moravcsik 1998, Pustejovsky 1995):

⁶ For further discussion, see Sections 7.2 and 7.3 below.

- CONSTITUTIVE: the relation between the entity and its parts, what it is part of, or what it is made of.
- FORMAL: that which we use to distinguish or pick out the entity (e.g., shape, orientation, taste, magnitude, and color).
- TELIC: the purpose or function of the entity.
- AGENTIVE: factors connected with the origin or way in which the entity came into being.

Qualia roles provide us with a powerful and flexible template for representing the meaning of Ns.⁷ Not all lexical terms need to have a value for each quale, and distinct types of terms (e.g., artifact vs. natural kind vs. perceptual terms) give different weight or relative importance to each quale. For example, perceptual terms have a value for the FORMAL and, in some cases, for the CONSTITUTIVE quale, but do not have a value for the TELIC or AGENTIVE quale. Artifact terms usually have a value for the FORMAL, and assign high importance to the AGENTIVE and TELIC quale, but in many cases do not have a value for the CONSTITUTIVE quale. Natural kind terms have a value for the FORMAL, and assign high importance to the CONSTITUTIVE and AGENTIVE quale, but in many cases do not need to have a value for the TELIC quale.⁸

⁷ I should note at outset that it is not essential to DUAL CONTENT to commit to only (or precisely these) dimensions of C-structure. For the types of expressions we will consider here, these dimensions seem particularly important. However, some persuasive recent accounts on the family of DUAL CONTENT propose additional dimensions. For example, McNally (2006) argues that gradable adjectives might have a scale structure dimension, and Knobe, Prasada & Newman (2013) argue that certain concepts seem to have a normative dimension which is distinct from the TELIC dimension. I take these interesting proposals to be compatible with DUAL CONTENT.

⁸ Two clarifications are in order. First, these suggestions are intended to reflect recent research on categorization and language acquisition, which suggest that what is given most weight when categorizing an artifact is the purpose for which it was created and how this relates to its function, whereas what is given most weight when categorizing a natural kind is what it is made of and/or what its natural origin is. For overviews of the data see Bloom 2002, Keleman & Carey 2007, Keil, Greif & Kerner 2007, and Margolis & Laurence 2007. To properly account for categorization patterns, it is likely that we have to add weights to the dimensions or qualia roles of C-structure. For further discussion of these issues, see Section 7. Secondly, the way in which qualia roles are generalized to represent the C-structure of lexical terms that do not denote kinds of objects, such as adjectives and mass nouns, will become clear below once we deal with expressions such as *counterfeit money*, *fake lawyer*, and *fake red gun*.

The idea that the meaning of lexical terms is organized around qualia roles is extensively defended by Pustejovsky (1995). However, Pustejovsky sometimes represents qualia roles as if they were “essential” or “definitional” properties of the extension of terms. This aspect of Pustejovsky’s theory is at odds with Putnam’s lessons against definitional theories summarized above, and has been for similar reasons criticized by Asher (2011), following Fodor & Lepore (2002). For these reasons, it should be clear from the outset that, by incorporating them into C-structure, DUAL CONTENT entails that qualia roles are *not*, in general, part of the extension-determining aspect (E-structure) of lexical terms. In addition, as will become clear below, the specific compositional operations which, according to DUAL CONTENT, have access to qualia roles are also entirely different from the combinatorial operations used in Pustejovsky’s Generative Lexicon theory. The latter operations have been persuasively criticized by Asher (2011).

To illustrate DUAL CONTENT, consider the following simplified lexical entries to which we will often refer:

(7) *lion*E-structure: $\lambda x.LION(x)$

C-structure:

CONSTITUTIVE: $\lambda x.SUBSTANCE-LION(x)$ FORMAL: $\lambda x.PERCEPTUAL-LION(x)$

TELIC:

AGENTIVE: $\lambda x.\exists e_1[BIOLOGICAL-BIRTH-LION(e_1, x)]$ (8) *gun*E-structure: $\lambda x.GUN(x)$

C-structure:

CONSTITUTIVE: $\lambda x.PARTS-GUN(x)$ FORMAL: $\lambda x.PERCEPTUAL-GUN(x)$ TELIC: $\lambda x.GEN e[SHOOTING(e) \wedge INSTRUMENT(e, x)]^9$ AGENTIVE: $\lambda x.\exists e_1[MAKING(e_1) \wedge GOAL(e_1, GEN e[SHOOTING(e) \wedge INSTRUMENT(e, x)])]$

⁹ Note that we formulate the TELIC as involving generic or typical events in which the associated objects can be used. This is what is captured by the generic quantifier over events GEN.

The E-structure of a lexical term e has two main roles. First, it specifies the number and types of arguments which e takes. Secondly, it stands for the (atomic) part of the meaning that determines e 's extension. For now, we take the relation between the atomic E-structure and the extension of lexical terms as primitive (for further discussion, see Section 7 below). The C-structure of e , as illustrated by (7) and (8), is richer than e 's associated prototype or stereotype, at least as the latter are traditionally understood. The C-structure of e can be thought of as a restricted set of general beliefs associated with e 's extension, which includes, roughly, information about how entities that fall under e tend to look, feel or taste, what materials or parts they are made of, how or for what purpose they came into being or were created, and what their intended and typical function is.

The key question that any formal implementation of DUAL CONTENT needs to answer is this: how do the compositional operations interact with semantically complex lexical entries such as (7) and (8)? Type-driven theories use general compositional rules — that is, rules that are not restricted to particular types of syntactic constructions. For our purposes, the most important are Functional Application (*FA*) and Predicate Modification (*PM*), here presented in their simplest form:

(*FA*) If α is a branching node, $\{\beta, \gamma\}$ is the set of α 's daughters, and $\llbracket \beta \rrbracket$ is a function whose domain contains $\llbracket \gamma \rrbracket$, then $\llbracket \alpha \rrbracket = \llbracket \beta \rrbracket(\llbracket \gamma \rrbracket)$.

(*PM*) If α is a branching node, $\{\beta, \gamma\}$ is the set of α 's daughters, and $\llbracket \beta \rrbracket$ and $\llbracket \gamma \rrbracket$ are both in $D_{(e,t)}$, then $\llbracket \alpha \rrbracket = \lambda x \in D_e. \llbracket \beta \rrbracket(x) \wedge \llbracket \gamma \rrbracket(x)$.

To extend the type-driven framework so that it can properly interact with lexical entries which have a dual semantic structure such as (7) and (8), we have to modify *FA* and *PM*. The first obvious modification is that we need two interpretation functions, one which determines the E-structure of complex terms, which we will designate $\llbracket \cdot \rrbracket_E$, and one which determines their C-structure, which we will designate $\llbracket \cdot \rrbracket_C$. In addition, we will assume that the original $\llbracket \cdot \rrbracket$ retrieves the entire meaning of expressions as an ordered set of the E-structure and C-structure. For example, $\llbracket e \rrbracket_E$ is just the E-structure of e , $\llbracket e \rrbracket_C$ is an ordered set of the C-structure of e , and $\llbracket e \rrbracket$ is an ordered set of the E-structure and the CONSTITUTIVE, FORMAL, TELIC and AGENTIVE of e .¹⁰

Following Krifka et al. (1995), we assume that the generic quantifier involves a contextually supplied restriction of the set of relevant events.

¹⁰ For now, we will assume that terms that do not have a value for one or more of these dimensions — that are, in this sense, unconstrained — represent this via the identity function.

For clarity, we first discuss how the E-structure of complex NPs is determined by $\llbracket \cdot \rrbracket_E$. Here is the first version of the new rules, specified only for E-structure, and focusing only on *FA*:

(FA^{DC}) If α is a branching node, $\{\beta, \gamma\}$ is the set of α 's daughters, and $\llbracket \beta \rrbracket_E$ is a function whose domain contains $\llbracket \gamma \rrbracket$, then $\llbracket \alpha \rrbracket_E = \llbracket \beta \rrbracket_E(\llbracket \gamma \rrbracket)$ and $\llbracket \alpha \rrbracket_C = \dots$

This minor modification of *FA* is required to specify what aspect of the semantic structure of lexical terms the interpretation function is retrieving and calculating. As we will see, for some types of complex NPs the E-compositional operations work just like the traditional compositional operations. In other cases, the E-compositional operations are more complex, which is why they need to take the entire meaning of expressions as arguments.

How do we bring the C-structure of lexical terms into the E-compositional operations in cases when we need to do so, such as in certain types of adjectival modifications? To do this, we introduce two tools:¹¹

- i. *Qualia functions*: partial functions from the meaning of terms into their respective C-structure denotations, namely, CONSTITUTIVE, FORMAL, TELIC, and AGENTIVE. The qualia functions are Q_C, Q_F, Q_T, Q_A . Consider the following examples, based on lexical entries (7) and (8):

$$Q_C(\llbracket lion \rrbracket) = \lambda x. \text{SUBSTANCE-LION}(x)$$

$$Q_T(\llbracket gun \rrbracket) = \lambda x. \text{GEN } e[\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)]$$

In addition, we will include a function Q_E which takes the meaning of terms and returns their E-structure denotations.

- ii. *Semantic re-structuring operators*: partial functions from the meaning of terms into certain combinations of their E-structure and C-structure.

If the terms are predicate modifiers instead of predicates, we can represent the empty value as a function maps properties unto themselves.

¹¹ Vikner & Jensen (2002) also use *similar* tools in their account of genitives. Although their account and DUAL CONTENT are roughly in the same family, there are some important difference between their version of qualia functions and semantic restructuring operators and the versions presented below.

The operators are C , T , A , and F . Consider the following examples:

$$\begin{aligned} A(\llbracket lion \rrbracket) &= \lambda x. Q_E(\llbracket lion \rrbracket)(x) \wedge Q_A(\llbracket lion \rrbracket)(x) \\ &= \lambda x. LION(x) \wedge \exists e_1 [\text{BIOLOGICAL-BIRTH-LION}(e_1, x)] \\ T(\llbracket gun \rrbracket) &= \lambda x. Q_E(\llbracket gun \rrbracket)(x) \wedge Q_T(\llbracket gun \rrbracket)(x) \\ &= \lambda x. GUN(x) \wedge \text{GEN } e [\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)] \end{aligned}$$

Using these tools, let us consider the case of privative NPs such as *fake gun*. As we saw above, the E-compositional effect of *fake* is more complex than the E-compositional effect of intersective and subsective Adjs. The main difference is that, according to DUAL CONTENT, *fake* interacts with various aspects of the E and C-structure of the term it modifies. For this reason, I propose that we treat *fake* as a semantic re-structuring operator along the lines of C , T , A , and F . But unlike those operators, *fake* and other privatives are *lexicalized* semantic restructuring operators. To allow lexicalized semantic restructuring operators into the composition process, we need to incorporate Qualia functions into the specification of their E-structure.

Following this idea, we can take $\llbracket fake \rrbracket_E$ as an operator which takes the meaning of the noun it modifies and restructures it by using — via Qualia functions — the noun’s C-structure. Here’s a first try:

$$(9) \quad \llbracket fake \rrbracket_E = \lambda D_C. [\lambda x. \neg Q_E(D_C)(x) \wedge \neg Q_A(D_C)(x) \wedge \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(D_C)(x))]]]$$

Note that D_C is a variable that ranges over ordered sets of the E-structure and C-structure of common Ns. So according to (9), *fake* takes the entire meaning, as an ordered set of the E-structure and all the elements of the C-structure, and outputs a new predicate of type $\langle e, t \rangle$ that is not a $Q_E(D_C)$, was not created with the same goal specified in $Q_T(D_C)$, but was created to have the perceptual features specified in $Q_F(D_C)$. Consider then $\llbracket fake \text{ gun} \rrbracket_E$, which by FA^{DC} , is equal to $\llbracket fake \rrbracket_E(\llbracket gun \rrbracket)$:

$$\begin{aligned} (10) \quad \llbracket fake \rrbracket_E(\llbracket gun \rrbracket) &= \lambda x. \neg Q_E(\llbracket gun \rrbracket)(x) \wedge \neg Q_A(\llbracket gun \rrbracket)(x) \wedge \\ &\quad \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(\llbracket gun \rrbracket)(x))] \\ &= \lambda x. \neg GUN(x) \wedge \\ &\quad \neg \exists e_1 [\text{MAKING}(e_1) \wedge \text{GOAL}(e_1, Q_T(\llbracket gun \rrbracket)(x))] \wedge \\ &\quad \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, \text{PERCEPTUAL-GUN}(x))] \end{aligned}$$

According to (10), $\llbracket \textit{fake gun} \rrbracket_E$ is a predicate of type $\langle e, t \rangle$, which is satisfied by objects that are not guns, were not made with the goal of being shooting instruments, and were made to have the perceptual features specified in $Q_F(\llbracket \textit{gun} \rrbracket)$ (i.e., to look like guns). This is exactly the result we wanted to get. Given appropriate lexical entries, this will get the intuitively correct results for *fake knife*, *fake headphones*, *fake pen*, and so on.

The entry presented in (9) captures several uses of *fake* that might initially seem quite different. For example, there is a sense of *fake* that is close to that of *counterfeit*. This use is illustrated in expressions such as [*fake [Chanel handbag]*] and [*fake [Polo shirt]*]. To account for these cases, the key thing to note is that (9) does *not* entail that a fake ϕ cannot have the function of a ϕ : in (9) the TELIC of the modified N is not negated. What (9) entails, however, is that a fake ϕ cannot have the same source as ϕ (i.e., cannot have the same AGENTIVE as ϕ). In the case of brand name artifacts, certain aspects of their institutional source are crucial, hence are likely included in their AGENTIVE. So (9) gives the intuitively correct results: fake Polo shirts and fake Chanel bags do not come from the required authorized sources, although they can be used as if they did. Now, phrases headed by artifact Ns modified by brand name terms are productive: we can understand *Chanel motorcycle* and *Polo cell-phone* even if the relevant products do not yet exist. This suggests that, in general, brand-name modifiers add a condition to the AGENTIVE of the modified N. Since *fake* negates that AGENTIVE, this entails that [*fake [Chanel handbag]*] can also mean a fake handbag made by Chanel. That this reading should be available is perhaps more obvious in the case of novel products. We can easily imagine that Chanel made a fake motorcycle for an exhibition, or that Polo made a fake cell-phone for some runway show.¹²

In addition, (9) also works well for expressions that denote fake gestures and fake social roles, such as *fake smile* and *fake lawyer*. A fake smile is intentionally produced to look like a smile, just like the role of a fake lawyer is usually to behave or look like a lawyer. Both cases, however, clearly have a different source or AGENTIVE than their real counterparts. The source of a real smile involves a certain sensation that is lacking in the fake version, although the intention is to exhibit the same outward features. The source of a real lawyer — how one becomes a lawyer — includes certain socially prescribed milestones (e.g., finishing law school, passing the bar, etc.). A fake lawyer has

¹² To be clear, this second reading is not possible under the structure [*fake Chanel*] *handbag*], which is also available. We discuss recursive NPs in more detail below — including modifications of complex privative NPs and different syntactic structures.

the outward appearance and behavior of a lawyer (e.g., “looks like a lawyer”), and can even carry out some of the same or similar functions of a lawyer. However, a fake lawyer lacks the AGENTIVE of a real lawyer — and since that is what gives lawyers their authority, fake lawyers do not have the relevant legitimacy.

These examples deserve more discussion, but for now what I want to highlight is that DUAL CONTENT, combined with an entry like (9), can account for the sense that *fake* is doing something at least slightly different in examples such as *fake gun*, *fake smile*, and *fake Chanel handbag*. This is explained on the basis of the important differences in the AGENTIVE of *gun*, *smile*, and *Chanel handbag*. In each case, *fake* itself is carrying out essentially the same operation, which includes negating the AGENTIVE of the modified N. However, different types of terms, including subtypes of artifacts, have different sorts of information in their AGENTIVE, which accounts for the intuitive differences.¹³

Privative Adjs in general can be treated as lexicalized semantic restructuring operators. Indeed, this approach has special advantages when their differences are subtle. Take the cases of *counterfeit* and *artificial*, which contrast in interesting ways with *fake*. An important difference between many paradigmatic uses of *counterfeit* and *fake* is that, unlike a fake, a counterfeit ϕ is usually made to look and function like a ϕ . For example, a counterfeit Rolex is made both to look like and function like a Rolex. *Counterfeit*, in its most usual sense, can be applied to terms for artifacts whose origin gives them an added value, with the result that this appropriate origin is lacking in the counterfeit versions. Consider the following entry:

$$(11) \quad \llbracket \textit{counterfeit} \rrbracket_E = \lambda D_C. [\lambda x. \neg Q_E(D_C)(x) \wedge \neg Q_A(D_C)(x) \wedge \\ \exists e_2 [\text{MAKING}(e_2) \wedge \\ \text{GOAL}(e_2, Q_F(D_C)(x) \wedge Q_T(D_C)(x))]]]$$

If we apply (11) to a plausible lexical entry for *Rolex*, we get that a *counterfeit Rolex* is not a Rolex, does not come into being in the manner in which a Rolex comes into being (as specified in $Q_A(\llbracket \textit{Rolex} \rrbracket)$), and is an artifact that was made to look *and* function like a Rolex (as specified in $Q_F(\llbracket \textit{Rolex} \rrbracket)$) and

¹³ Even if we accept DUAL CONTENT, there are cases that crucially involve some degree of context-induced meaning modulation. For discussion of cases when the meaning of *fake* itself is modulated, and more generally on the relation between DUAL CONTENT and contextualism, see Section 7.1 below.

$Q_T(\llbracket \text{Rolex} \rrbracket)$ respectively).¹⁴ Consider, on the other hand, the case of *artificial*, as used in expressions such as *artificial leg* and *artificial heart*. Unlike a fake heart, an artificial heart is made with the intention that it function like a heart. In addition, unlike fakes and counterfeits, an artificial heart need not be made to look like a heart, as long as it was made to function like one. The following entry captures this important distinction between *artificial* on the one hand and *fake/counterfeit* on the other:

$$(12) \quad \llbracket \text{artificial} \rrbracket_E = \lambda D_C. [\lambda x. \neg Q_E(D_C)(x) \wedge \neg Q_A(D_C)(x) \wedge \\ \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_T(D_C)(x))]]]$$

Unlike the case of *fake* and *counterfeit*, in this case we might want to eliminate $\neg Q_E(D_C)(x)$ from (12). Indeed, some people I have informally queried about whether, say, artificial hands and legs are really arms and legs report mixed and unstable intuitions. This is of course an empirical question, and both options are compatible with DUAL CONTENT. We reserve further discussion of the descriptive accuracy of these proposals, including a discussion of the ways in which DUAL CONTENT and contextualism are compatible, for Section 7.1 below.

In addition to privatives, other types of Adjs can also be treated as lexicalized semantic restructuring operators. Consider subsecutive Adjs such as *typical*, *representative*, or *normal*, as used in expressions such as *typical lion/gun*, *representative French/Harvard Alumni*, *normal lion/gun*. We can treat these Adjs as operators which take the full meaning of Ns and return a predicate of type $\langle e, t \rangle$ which includes in its E-structure a part of the original N's C-structure. Consider the case of *typical*:

$$(13) \quad \llbracket \text{typical} \rrbracket_E = \lambda D_C. [\lambda x. Q_E(D_C)(x) \wedge Q_C(D_C)(x) \wedge \\ Q_F(D_C)(x) \wedge Q_T(D_C)(x) \wedge Q_A(D_C)(x)]$$

According to (13), *typical* takes D_C and outputs a new predicate that is satisfied by objects that fall under $Q_E(D_C)$, $Q_F(D_C)$, $Q_T(D_C)$, and so on. So $\llbracket \text{typical gun} \rrbracket_E$, which by FA is equal to $\llbracket \text{typical} \rrbracket_E(\llbracket \text{gun} \rrbracket)$, would be:

¹⁴ The difference between *fake* and *counterfeit*, on this account, is quite subtle: counterfeit artifacts are explicitly made to function like the artifacts they are counterfeits of. This is not required in the case of fakes, although in some cases it is possible to so use them.

$$\begin{aligned}
(14) \quad & \llbracket \text{typical} \rrbracket_E(\llbracket \text{gun} \rrbracket) \\
& = \lambda x. Q_E(D_C)(x) \wedge Q_C(D_C)(x) \wedge \\
& \quad Q_F(D_C)(x) \wedge Q_T(D_C)(x) \wedge Q_A(D_C)(x) \\
& = \lambda x. \text{GUN}(x) \wedge \text{PARTS-GUN}(x) \wedge \\
& \quad \text{PERCEPTUAL-GUN}(x) \wedge \\
& \quad \text{GEN } e[\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)] \wedge \\
& \quad \exists e_1[\text{MAKING}(e_1) \wedge \\
& \quad \quad \text{GOAL}(e_1, \text{GEN } e[\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)])]
\end{aligned}$$

According to (14), a *typical gun* is a gun that has all the common features of a gun, as represented in the C-structure of *gun*: it looks like a gun, can be used for shooting, was made to be used for shooting, and is made out of whatever materials and parts guns are typically made out of (which, for simplicity, I have represented as PARTS-GUN). Now, we said before that some default lexical entries might lack a quale value in their C-structure. In addition, *typical* does not seem to have any restrictions with the type of common Ns it can combine with. However, in (13) *typical* is defined over all the values of the qualia. There are two simple ways to deal with this. One is to assume that *typical* is ambiguous: one of its senses operates over the qualia roles defined for artifacts, another sense over the qualia roles defined for natural kinds, and so on. Another way to deal with it is to assume that *typical* has the meaning proposed in (13) but that it coerces the meaning of the lexical terms it modifies to specify a value for each of the qualia roles, and that we use our background (encyclopedic) knowledge to do this online.

The discussion of *typical* raises two further questions about *fake*. The first is how to deal with cases when *fake* modifies a N that is not an artifact term, or more precisely, one that is undefined with respect to its TELIC or AGENTIVE. Just as in the case of *typical*, there are at least two ways to deal with this, which are not mutually exclusive and could sometimes work together. The first is to assume that the meaning of the modified head N is modulated online: using general encyclopedic information, it is enriched to include a value for the TELIC and AGENTIVE. The second is to assume that the meaning of *fake* is modulated so that it can apply to the head N. This could be proposed to deal with modifications of mass terms such as *gold* and *titanium*, which, some would argue, don't seem to have a value for the AGENTIVE. The modulated meaning of *fake* in these cases could be something like:

$$(15) \quad \llbracket fake_m \rrbracket_E = \lambda D_C. [\lambda x. \neg Q_E(D_C)(x) \wedge \neg Q_C(D_C)(x) \wedge \exists e_2 [\text{MAKING}(e_2, x) \wedge \text{GOAL}(e_2, Q_F(D_C)(x))]]$$

One difference between *fake* and *fake_m* is that the former negates the AGENTIVE of the term it modifies and the latter negates the CONSTITUTIVE. This would explain why a fake gun can be made of steel but could not have been made to shoot perfectly, whereas fake gold cannot be made of gold but could be used, in some cases, to pass for gold. However, we need not assume this kind of meaning modulation of *fake*. We can still explain cases such as *fake gold* and *fake titanium* by insisting that the meaning of *fake* is as in (9). To do so we need to assume that, either in the context of privative NPs or in general, substance terms such as *gold* and *titanium* have an AGENTIVE. This is a reasonable move; for we could represent in their AGENTIVE the information that gold and titanium “come from” particular molecular compositions. In this case, (9) would give the intuitively correct results for *fake gold* and *fake titanium*, namely, that they have some of the perceptual features of gold and titanium, but are not made up of the same stuff as gold and titanium.¹⁵

The other question raised by the discussion of *typical* is how to deal with modifications of complex privative NPs. We must account for the meaning of expressions such as *typical fake gun*. We said that *typical* takes the full meaning of a lexical term, which involves operations on its C-structure. If this is correct, it follows that *fake gun* must itself have a C-structure. To account for this, we first have to fully define FA^{DC} , since so far we have only defined it for the E-structure of terms.

¹⁵ I would argue that this second route is the better bet. If we follow this route, what would be stored in the CONSTITUTIVE is information about the identifiable part-whole structure; in this case, that mass terms do not have identifiable parts. Treating the CONSTITUTIVE in this way would also allow us to incorporate into DUAL CONTENT some of the suggestions for color and degree modifications presented by McNally (2006). Then the information that we intuitively regard as “what substances are made of” (e.g., their molecular composition) would, when available, go in the AGENTIVE. Furthermore, it is worth noting that this option interacts nicely with some recent views on artifactual substances. For example, Bloom (2007) argues that some substance terms can be used in an artifactual sense. One way to represent this in DUAL CONTENT is to hold that their AGENTIVE has two conditions: the molecular composition of the substance and the production source. In this case, DUAL CONTENT would predict that some substance could be *fake Coke* or *fake Sprite* either because of differences in composition with the originals, or because it does not come from, say, the companies that are authorized to make Coke or Sprite. If we follow Bloom, this result seems along the right lines.

(FA^{DC}) If α is a branching node, $\{\beta, \gamma\}$ is the set of α 's daughters, and $\llbracket \beta \rrbracket_E$ is a function whose domain contains $\llbracket \gamma \rrbracket$, then $\llbracket \alpha \rrbracket_E = \llbracket \beta \rrbracket_E(\llbracket \gamma \rrbracket)$ and $\llbracket \alpha \rrbracket_C = \langle Q_C(\llbracket \beta \rrbracket)(\llbracket \gamma \rrbracket), Q_F(\llbracket \beta \rrbracket)(\llbracket \gamma \rrbracket), Q_T(\llbracket \beta \rrbracket)(\llbracket \gamma \rrbracket), Q_A(\llbracket \beta \rrbracket)(\llbracket \gamma \rrbracket) \rangle$

FA^{DC} captures a rather simple idea about C-structure combinatorics, namely, that the relevant quale of the modifier applies via functional application to the corresponding quale of the modified expression.¹⁶ To do this, we assume that each quale in the C-structure of semantic-restructuring operators can take as argument the full meaning of the N they modify. Consider then the full lexical entry for *fake*:

(16) *fake*

E-structure: $\lambda D_C. [\lambda x. \neg Q_E(D_C)(x) \wedge \neg Q_A(D_C)(x) \wedge \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(D_C)(x))]]$

C-structure:

CONSTITUTIVE: $\lambda D_C. Q_C(D_C)$ ¹⁷

FORMAL: $\lambda D_C. Q_F(D_C)$

TELIC: $\lambda D_C. \neg Q_T(D_C)$

AGENTIVE: $\lambda D_C. [\lambda x. \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(D_C)(x))]]$

Given (16) and (8), by FA^{DC} the full entry for *fake gun* is:

¹⁶ A reviewer suggests that this pointwise formulation of FA^{DC} might overcommit DUAL CONTENT to a particularly strong view of how the C-structures of complex NPs in general are determined. For the cases we are considering, this version of FA^{DC} works well. In addition, in DUAL CONTENT modifiers take the full meanings of heads and can use qualia functions to shuffle the internal semantics, so this gives some extra flexibility to this account even if we stick to the pointwise combinatorics. However, it is undeniable that other types of NPs might involve different combinatorial dynamics. If this turns out to be the case, we can replace the current pointwise formulation of FA^{DC} with a more general formulation, which just says that $\llbracket \alpha \rrbracket_C$ is a function f of $\llbracket \beta \rrbracket_C$ and $\llbracket \gamma \rrbracket_C$. We might then further specify that for the types of expressions we are considering (i.e., subsectives and privatives), what f does is to perform the previously specified function application along each dimension of C-structure.

¹⁷ This unconstrained value for the CONSTITUTIVE captures that *fake* does not impose any additional conditions on the stuff or parts which a fake ϕ should have. For example, a fake gun can be made of any stuff and internal parts, as long as it is made to pass for a gun. If we adopt the entry for *fake_m* specified in (15) above, the CONSTITUTIVE entry would have to be modified to specify that a fake_m ϕ does not have the CONSTITUTIVE of ϕ .

(17) *fake gun*

E-structure: $\lambda x. \neg Q_E(\llbracket \textit{gun} \rrbracket)(x) \wedge \neg Q_A(\llbracket \textit{gun} \rrbracket)(x) \wedge$
 $\exists e_2[\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(\llbracket \textit{gun} \rrbracket)(x))]$

C-structure:

CONSTITUTIVE: $Q_C(\llbracket \textit{gun} \rrbracket)$ FORMAL: $Q_F(\llbracket \textit{gun} \rrbracket)$ TELIC: $\neg Q_T(\llbracket \textit{gun} \rrbracket)$ AGENTIVE: $\lambda x. \exists e_2[\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(\llbracket \textit{gun} \rrbracket)(x))]$

We can now compute $\llbracket \textit{typical fake gun} \rrbracket_E$, which has the structure $[\textit{typical} \llbracket \textit{fake gun} \rrbracket]$, hence, by FA:

(18) $\llbracket \textit{typical} \rrbracket_E(\llbracket \textit{fake gun} \rrbracket)$

$$= \lambda x. Q_E(\llbracket \textit{fake gun} \rrbracket)(x) \wedge Q_C(\llbracket \textit{fake gun} \rrbracket)(x) \wedge$$

$$Q_F(\llbracket \textit{fake gun} \rrbracket)(x) \wedge Q_T(\llbracket \textit{fake gun} \rrbracket)(x) \wedge Q_A(\llbracket \textit{fake gun} \rrbracket)(x)$$

Comparing (10) with (18), we can see that there are important differences between *fake gun* and *typical fake gun*. According to (10), a *fake gun* is not a gun, was not made to function for shooting, and was made to seem like a gun. However, there are malfunctioning fake guns and badly made fake guns, which is part of the reason why we did not say anything about whether or not fake guns can actually shoot or can look different than real guns. We can imagine a badly made fake gun that can shoot or that can be easily identified as a fake gun. However, a *typical* fake gun, as correctly captured in (18), is an artifact that looks like a gun and that cannot be used to shoot.

Just like *typical*, *fake* can modify complex NPs of the same type as *gun* such as *red gun* and *heavy gun*. So just like privative NPs, intersective and subsective NPs must also have an E-structure and a C-structure. There are various options for how to determine the C-structure of intersective, subsective and other types of complex NPs. Indeed, there is plenty of research on how to determine the meaning, including the C-structure, of complex expressions when their constituents also have complex meanings (McNally 2006, Partee & Kamp 1995, Prinz 2012, Pustejovsky 1995, 2012, Hampton & Jonsson 2012). Although some look very promising, most of the proposals are preliminary. Hence we cannot yet provide a full account of how to determine, for each type of complex NP, their C-structure from the meanings of their parts. However, we can show that, at least for paradigmatic intersective and subsective NPs, this is not an in-principle problem. Furthermore, some of

the most plausible suggestions can be implemented in the DUAL CONTENT framework.

Consider the following lexical entry for *plastic* (for brevity, I simplify the representations of the values of C-structure):

(19) *plastic*

E-structure: $\lambda x.\text{PLASTIC}(x)$

C-structure:

CONSTITUTIVE: $\lambda x.\text{PLASTIC-MASS}(x)$

FORMAL: $\lambda x.\text{PERCEPTUAL-PLASTIC}(x)$

TELIC:

AGENTIVE: $\lambda x.\text{PLASTIC-COMP}(x)$ ¹⁸

We can easily come up with a modified rule of *PM* that can be used to combine (19) with *gun* and get the appropriate C-structure for *red gun*. However, so that we do not have to further modify the general compositional operations we have introduced so far to account for the E-structure and C-structure of complex privative NPs, we can assume that when used attributively, *plastic* and other intersective Adjs are type-shifted into semantic restructuring operators, similar to the case of privative Adjs, except that they are simpler.

¹⁸ This entry follows the suggestion—made above in the discussion of *gold*, *fake gold*, etc.—that the CONSTITUTIVE of mass terms only specifies that they do not have identifiable parts. The information about their molecular composition (e.g., that they come from a certain combination of elements), when available, goes into the AGENTIVE. Now, it might seem strange to distinguish, as we do in this entry, between the E-structure and the AGENTIVE of *plastic*. Various theorists argue that the essence of substances is just the stuff they are made of. Although we could implement this view within DUAL CONTENT, it is worth noting that it is not uncontroversial (see, e.g., Bloom 2007), and it is arguably implausible for the case of *plastic*. In addition, there are cases of intersective Adjs that seem to have an E-structure that cannot be identified with any particular dimension of their C-structure. For example, color Adjs might have E-structures with functions of type $\langle e, t \rangle$ that are uncommitted to the particular quantity/quality of the way in which entities that fall under them have to be colored. Still, these quantity/quality dimensions are arguably part of the C-structure of colors, i.e., part of the typical way in which we conceive of objects being colored.

(20) *plastic*E-structure: $\lambda D_C. [\lambda x. \text{PLASTIC}(x) \wedge Q_E(D_C)(x)]$

C-structure:

CONSTITUTIVE: $\lambda D_C. [\lambda x. \text{PLASTIC-MASS}(x) \wedge Q_C(D_C)(x)]$ FORMAL: $\lambda D_C. [\lambda x. \text{PERCEPTUAL-PLASTIC}(x) \wedge Q_F(D_C)(x)]$ TELIC: $\lambda D_C. Q_T(D_C)$ AGENTIVE: $\lambda D_C. [\lambda x. \text{PLASTIC-COMP}(x) \wedge Q_A(D_C)(x)]$

By FA^{DC} , if (20) is combined with *gun*, as specified in (8), we get the following result:

(21) *plastic gun*E-structure: $\lambda x. \text{PLASTIC}(x) \wedge \text{GUN}(x)$

C-structure:

CONSTITUTIVE: $\lambda x. \text{PLASTIC-MASS}(x) \wedge \text{PARTS-GUN}(x)$ FORMAL: $\lambda x. \text{PERCEPTUAL-PLASTIC}(x) \wedge \text{PERCEPTUAL-GUN}(x)$ TELIC: $\lambda x. \text{GEN } e [\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)]$ AGENTIVE: $\lambda x. \text{PLASTIC-COMP}(x) \wedge$ $\exists e_1 [\text{MAKING}(e_1) \wedge \text{GOAL}(e_1,$ $\text{GEN } e (\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x))]$

It is easy to see that *fake* can apply to *plastic gun*, following the procedure used in the previous cases: a fake plastic gun is, in the most salient reading, a plastic artifact that is a fake gun, in the sense explained above. Also available, on this view, is the reading according to which a *fake plastic gun* is a real gun that is made of plastic. Now, note that sometimes *fake* combines first with another N used as a modifier, and only then with the head N. In this case, a *fake plastic gun* or *fake steel gun* can only mean a real gun that is made of fake plastic or fake steel. This is often the intended meaning in cases such as *fake gold ring* and *fake diamond necklace*. Given the correct underlying structure (e.g., $[[\text{fake plastic}] \text{ gun}]$, $[[\text{fake gold}] \text{ ring}]$) and lexical entries, DUAL CONTENT easily accounts for these cases.

This example is only meant to illustrate a simple case of intersective composition, but it does suggest a general strategy for calculating the E-structure and C-structure of intersective NPs, namely, to assume that the lexical entries of modifiers provide, in their C-structure, information about

how and which qualia of the N they combine with should be modified, and which qualia should be directly inherited from the modified N. Of course, within this strategy, different types of modifiers demand slightly different treatments. Some modifiers might raise special challenges for the composition of the C-structure of their corresponding complex NPs.¹⁹ Constructing a full compositional account to determine the C-structure of all types of complex NPs is an ongoing project — but one which is constantly advancing and, despite some early challenges, most notably by Fodor (1998) and Fodor & Lepore (2002), does not seem to face any deep in-principle problems.²⁰

To further develop and defend DUAL CONTENT, I will compare it with another account of privative NPs, one which can be combined with atomistic lexical semantic theories (Sections 4–6). I then conclude by addressing some objections to DUAL CONTENT (Section 7).

4 Contextualist account of privative NPs

Partee (2007, 2010) presents a contextualist account of privative NPs which can also be implemented in a standard type-driven semantic framework. In addition, her account, which has recently been gaining approval (Recanati 2010, Morzycki forthcoming), can be paired with an atomic account of the internal semantic structure of head Ns. So before we can confidently hold that privative NPs motivate DUAL CONTENT’s account of lexical decomposition, we have to evaluate the prospects of Partee’s account.

Partee’s account rests on two simple proposals:

- i. eliminate the class of privative Adjs, and
- ii. treat privatives as subsecutive Adjs that coerce their N arguments into systematic meaning expansions.

Take the following examples:

-
- ¹⁹ For example, in some cases the modifier might contribute a quale value which is inconsistent with the value of the head N. This might happen in cases such as *purple dollar bill*, since it is plausible to assume that the colors specified in the FORMAL for *dollar bill* are inconsistent with *purple*. In these cases, there must be a mechanism which resolves these inconsistencies, favoring the value of the modifier. This kind of solution is compatible with DUAL CONTENT.
- ²⁰ For an interesting response to the compositionality challenge to prototype theory, see Prinz 2012. See also Part V of *The Oxford handbook of compositionality* (Hinzen, Machery & Werning 2012), which has various essays that address issues related to the combinatorial properties of lexical items with complex semantic structures.

(22) Is that gun real or fake?

(23) A fake gun is not a gun

Partee argues that when interpreting (22) we expand the meaning of *gun* to include both fake guns and real guns. In (23), the same is true of the first occurrence of *gun*, modified by *fake*, but not of the second. In short, Partee's account amounts to the following readings of (22) and (23):

(24) Is that gun_2 real or fake?

(25) A fake gun_2 is not a gun_1

Subscript 1 marks literal interpretations and 2 marks loose interpretations. For example, the extension of gun_2 includes not only real guns but also things that look like guns, such as toy guns and fake guns.

Cases like (22) are neutral between DUAL CONTENT and Partee's account. In many ordinary contexts we use artifact Ns such as *gun* in the loose sense of gun_2 — even when we use them in NPs that don't have modifiers such as *fake* or *real*. Some examples are when *gun* appears in demonstrative expressions such as (22), and phrases like *that red gun*, *that shiny gun*, and so on. In general, when such demonstratives are used to pick out an object that loosely looks like a gun, they are taken as properly used, whether or not the object is a real gun. In these and many other cases, artifact Ns are interpreted in the loose sense. DUAL CONTENT is compatible with this fact.²¹

Consider next examples like (23). According to Partee, we should treat privative Adjs as subjectives that coerce certain meaning-shifts in the Ns which they modify, so we need an account of what determines the direction of such meaning-shifts. Partee appeals to two plausible principles of interpretation (Partee & Kamp 1995):

(NON-VACUITY) In any given context, try to interpret any predicate so that both its positive and negative extension are non-empty.

(HEAD-PRIMACY) In a modifier-head structure, the head is interpreted relative to the context of the whole constituent, and the modifier is interpreted

²¹ As we will discuss in detail in Section 7, DUAL CONTENT is compatible with contextualist compositional accounts — such as that defended in (Recanati 2010) — which hold that the literal meaning of the constituents of complex expressions is usually affected, to some degree, by meaning modulations. However, the issue in this discussion is about the correct account of the default interpretation of privative NPs, and in this particular case DUAL CONTENT, unlike SHIFTING HEADS, does not invoke meaning modulations.

relative to the local context created from the former context by the interpretation of the head.

NON-VACUITY applies not only to simple predicates, but also to complex predicates, e.g., to predicates obtained via a Adj+N combination. HEAD-PRIMACY ensures that the interpretation of a head N in a given context restricts the local domain for the interpretation of the Adj to the positive extension of the head N in that context.

Partee is surely right that NON-VACUITY and HEAD-PRIMACY help explain certain patterns of context-induced meaning-shifts. Evidence for HEAD-PRIMACY comes from examples like these:

- (26) a. giant midget
 b. midget giant

In most contexts, we take (26a) to mean a midget, but a very large one, and (26b) to mean a giant, but a very small one. Such examples suggest that in Adj+N constructions, one first interprets the head N against the general context of the utterance, and then shifts the meaning of the modifying Adj so that — to respect NON-VACUITY — it subsects the extension of the head N.

There is also plenty of evidence for NON-VACUITY. NON-VACUITY is a particular manifestation of the rational pursuit of the basic goals of communication, as conceived in the Gricean tradition. In some contexts, NON-VACUITY can be taken as a consequence of general principles of cognitive processing, such as the Relevance Principle (Sperber & Noveck 2004). Overall, NON-VACUITY is a descriptively adequate information maximization principle of interpretation, which might follow from certain general cognitive biases and/or basic goals of rational communication. We will follow Partee and assume that something close to NON-VACUITY is a pragmatic principle of interpretation.

The next step in Partee's account of the direction of the meaning-shifts of the heads of privative NPs is a proposal about how NON-VACUITY and HEAD-PRIMACY interact. NON-VACUITY and HEAD-PRIMACY usually work together to partly determine the pragmatic interpretation of complex expressions, but sometimes satisfying one principle requires violating the other. In cases of conflict, Partee argues that NON-VACUITY ranks higher than HEAD-PRIMACY.

There is a productive class of complex expressions that supports Partee's proposed ranking, called the *constitutive material NPs*, such as *stone lion*, *wooden horse*, and *rubber duck*. When interpreting these expressions, we au-

tomatically and effortlessly loosen the meaning of the head Ns. For example, we generally interpret *stone lion* not as a real lion made of stone, but as a lion representation which is made of stone (e.g., a lion stone-sculpture). This suggests that in *stone lion* we interpret *lion* as *lion*₂ — that is, as including lions and things that only look like lions. In these cases, the direction of the meaning-shift of the head Ns seems to be determined by NON-VACUITY. For most mature speakers strongly believe that, for instance, in most contexts the set $\{x: x \text{ is a lion and } x \text{ is made of stone}\}$ is empty, whereas the set $\{x: x \text{ is a lion or representation of a lion and } x \text{ is made of stone}\}$ includes some objects and excludes others.

In short, Partee proposes that the default interpretation of certain constitutive material NPs involves the following processing sequence:

- i. triggered by a violation of NON-VACUITY, the interpretation proceeds to
- ii. coerce the meaning of the head N such that when composed with the literally taken modifying A,
- iii. the resulting expression is believed to respect NON-VACUITY, and
- iv. this meaning shift occurs despite the fact that HEAD-PRIMACY is thereby violated.

This is a persuasive account of constitutive material NPs; we will assume in what follows that it is essentially correct.

Partee's key suggestion is that the principles involved in the default interpretation of privative NPs are the same as those involved in the default interpretation of constitutive material NPs. To illustrate, consider the process which, according to Partee, leads to the default interpretation of *fake gun*. Assume that $\llbracket gun_1 \rrbracket$ is a function of type $\langle e, t \rangle$, $\llbracket fake \rrbracket$ is a function of type $\langle \langle e, t \rangle, \langle e, t \rangle \rangle$, and that the extension of *fake gun*₁ is $\{x: x \text{ is a fake (thing) and } x \text{ is a gun}\}$. According to Partee, in most contexts mature speakers (firmly) believe that this set is empty, or in other words, that there are no objects which are both real guns and fake things. If so, *fake gun*₁ — the literal interpretation of *fake gun* — violates NON-VACUITY, which triggers a pragmatic meaning-shifting process. Partee implicitly assumes that there is no easy way of respecting both NON-VACUITY and HEAD-PRIMACY by

shifting the meaning of *fake*.²² Since NON-VACUITY ranks higher than HEAD PRIMACY, the next option is to shift the meaning of *gun* to the expanded *gun*₂. Most speakers believe that the set { x : x is a fake (thing) and x is a gun or looks like a gun} includes some objects and excludes others. So by interpreting *fake gun* as *fake gun*₂ we satisfy NON-VACUITY.

Call Partee’s account of privative NPs “SHIFTING HEADS”. We will now examine some cases in which DUAL CONTENT and SHIFTING HEADS make different predictions about the interpretation of privative NPs, and argue that DUAL CONTENT makes the correct predictions.

5 DUAL CONTENT VS. SHIFTING HEADS

The main difference between DUAL CONTENT and SHIFTING HEADS is this:

- According to SHIFTING HEADS, pragmatically-induced meaning-shifts of the head Ns, guided by principles such as NON-VACUITY, are an essential part of the process leading to the default interpretation of privative NPs. These are the default interpretations because in most contexts we (explicitly or implicitly) follow NON-VACUITY and believe that *fake gun*₁ and the like have an empty extension.
- According to DUAL CONTENT, the default interpretation of privative NPs does not involve a violation of NON-VACUITY: for instance, $[[fake]]$ takes *gun*₁ and returns the predicate specified in (10); a predicate which most mature speakers (correctly) believe has a positive extension (i.e., fake guns) and a negative extension. According to DUAL CONTENT, the default interpretation of *fake gun* is based on the literal interpretation of the parts and rests only on linguistic competence.

This contrast suggests a way of comparing these accounts. NON-VACUITY is a pragmatic principle of interpretation. Like other such principles, we can relax or suspend it in certain contexts. In such contexts, SHIFTING HEADS predicts that the interpretation of privative NPs should differ from their

²² It is not clear why Partee does not consider this route (i.e., loosening the meaning of the privative Adj instead of that of the head N), especially since this would satisfy both NON-VACUITY and HEAD-PRIMACY. In many Adj+N constructions, we modify the meaning of the Adj rather than that of the head N. Maybe Partee does not consider this option because she thinks that — unlike the case of, say, color Adjs — it is not easy to come up with a loose interpretation of *fake* and other privative Adjs. However, given certain plausible proposals for lexical entries for *fake*, this is a very problematic assumption, as we will see in Section 6.

default interpretation. In contrast, DUAL CONTENT doesn't predict this, for according to it the default interpretation of privative NPs is roughly the literal interpretation. Furthermore, according to SHIFTING HEADS the process that determines the default interpretation of privative NPs is analogous to the process that determines the default interpretation of constitutive material NPs. So SHIFTING HEADS predicts that, when NON-VACUITY is relaxed, the interpretations of privative NPs and of constitutive material NPs should be affected in similar ways.²³

This prediction, however, is incorrect. When NON-VACUITY is suspended or relaxed, the interpretation of privative NPs is affected in a completely different way from the interpretation of constitutive material NPs, as predicted by DUAL CONTENT.

Consider first the interpretation of constitutive material NPs. The default interpretation of expressions like *stone lion* and *plastic rabbit* is *stone lion*₂ and *plastic rabbit*₂. The standard explanation of this is that, in general,

- i. speakers believe that there are no stone lions₁ and plastic rabbits₁,
- ii. they accept something like NON-VACUITY, and
- iii. to satisfy NON-VACUITY they loosen the meaning of *lion* and *rabbit*.

However, although accepting NON-VACUITY is the default stance, interpreters can and sometimes in fact do suspend or relax NON-VACUITY. They do this when they are instructed by the context — including hedges used by the speaker — to take the expressions literally, even if they believe that this would result in empty predicates. Consider the following examples:

²³ Whether we want to say that NON-VACUITY can be “suspended” or just “relaxed” depends on the way in which the principle is fleshed out. Either option is consistent with the position we will present. The basic point can be illustrated with a simple example. In most contexts we would interpret an utterance of *many people believe that politicians are snakes* metaphorically. But there are cases in which, based on certain cues, we can interpret it as meaning that many people really believe that politicians are, literally, snakes. This is so even if we in fact believe that this assertion is false — that is, that it cannot possibly be the case that many people believe such a thing. It is in this sense that NON-VACUITY can be suspended or relaxed: namely, that in some contexts, we can interpret expressions in ways that, relative to our beliefs, result in highly implausible assertions. This suggests another way one might flesh out the notion of relaxing NON-VACUITY, which some theorists might find more congenial: we always follow NON-VACUITY but in certain special contexts we can adopt specially unconstrained sets of beliefs (or different ones, for instance in a fictional discourse) relative to which assertions are evaluated. For our purposes, any of these notions of suspending or relaxing NON-VACUITY will do.

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

In cases like (27) and (28), the use of *literally* and the reinforcing context have a unified effect: to instruct listeners that the head Ns should be taken literally, even if they believe that it is unlikely that the resulting expressions have a positive extension, in this or any likely future world. Part of the effect of contexts such as these, is to instruct interpreters to withhold NON-VACUITY-guided meaning modulations. This is why (27) and (28) are interpreted as:

- (29) ... stone lions₁ and rubber rabbits₁.
- (30) ... silicon cows₁.

As (27) and (28) illustrate, relaxing NON-VACUITY directly affects the interpretation of constitutive material NPs in the expected ways, since there is no doubt that NON-VACUITY determines the direction of the meaning shifts of the head Ns in the default interpretations of constitutive material NPs.²⁴

At this point, we should present some new terminology and brief clarifications regarding literal contexts. Let *L-constitutive material NPs* stand for constitutive material NPs in contexts that call for strict literal interpretations. Similarly, let *L-privative NPs* stand for privative NPs in contexts that call for strict literal interpretations. We can induce literal interpretations by general contextual cues. So, slightly reinforced, examples like (27) and (28) call for a literal reading even if we delete the tokens of *literally*. Still, the use of *literally* in these cases falls squarely within what [Israel 2002](#) argues is its orthodox usage, namely, as a metalinguistic regulator or instruction to take expressions literally, to add nothing or change nothing to the meaning of the words. According to Israel, this orthodox use of *literally* has four important characteristics:

²⁴ To be sure, one could claim instead that this suggests that NON-VACUITY is not really responsible for the interpretation of constitutive material NPs. For don't (27) and (28) show that, in some sense, the extension of constitutive material NPs are not really empty, at least relative to some possible worlds? However, the crucial point is that in most contexts the relevant possible worlds are such that only the loose interpretations of constitutive material NPs have a positive and negative extension. This is all that is presupposed by Partee's account of constitutive material NPs. (For an elaboration of Partee's account, see [Hogeweg 2012](#).)

- i. It is used with expressions that lend themselves, in context, to multiple interpretations.
- ii. At least one of the interpretations is pragmatic.
- iii. The pragmatic interpretation is likely to be seen as more plausible than the literal interpretation.
- iv. The speaker is committed to the truth of the literal interpretation, even if it is remarkable or might seem highly implausible.

This account of the orthodox metalinguistic use of *literally* fits perfectly with the observed patterns in the case of L-constitutive material modifiers.²⁵

Recall the key move in Partee's account: in the relevant processing respects, she treats the default interpretation of privative NPs just like the interpretation of constitutive material NPs. It follows, *prima facie*, that the account predicts that when NON-VACUITY is suspended or relaxed the interpretation of constitutive material NPs and of privative NPs should be affected in analogous ways. In other words, this account predicts that, in L-privative NPs, we interpret *fake gun* as *fake gun*₁, not as *fake gun*₂.

These predictions are incorrect. In most contexts, the constituent privative NPs in L-privative NPs are interpreted in essentially the same way as the original, unmodified privative NPs, unlike the cases of L-constitutive material NPs. To see this, compare (27) and (28) with the following examples of L-privative NPs:

²⁵ A reviewer suggested that although, as stated, both DUAL CONTENT and SHIFTING HEADS assume that the standing or default meaning of, say, *lion* and *gun* is *lion*₁ and *gun*₁, this assumption might be incorrect. The idea seems to be that we might give a simpler account of privatives if we just assume that the default meaning Ns is the loose meaning, namely, *lion*₂, *gun*₂, and so on. This suggestion seems intuitively incorrect, but more importantly, it fails to account for the compositional behavior of L-constitutive material NPs, under the plausible assumption that *literally* is a metalinguistic operator of the sort proposed by Israel 2002. These cases show that, when instructed to take things literally, speakers interpret *stone lion* and *silicon rabbit* as *stone lion*₁ and *silicon rabbit*₁, and not as *stone lion*₂ and *silicon rabbit*₂. This strongly suggests that the literal, default meaning of common Ns is not the loose meaning. This same conclusion is also supported by expressions which use *strictly speaking*. For example, *Strictly speaking, penguins are birds* and *Strictly speaking, tomatoes are fruits* are clearly true. However, *Strictly speaking, lion sculptures are lions* and *Strictly speaking, rubber toy rabbits are rabbits* seem clearly false, although they should come out as obvious truths under the assumption that the default meanings of *lion* and *rabbit* are *lion*₂ and *rabbit*₂.

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

Examples (31) and (32), just like (27) and (28), instruct interpreters to take the expressions literally; however, this does not affect the interpretation of privative NPs in the way it affects the interpretation of constitutive material NPs. In (31) and (32) the interpretations of *fake gun* are just their corresponding default interpretations, namely, *fake gun*, which SHIFTING HEADS represents as *fake gun*₂. In other words, relaxing/suspending NON-VACUITY doesn't change the default interpretation of privative NPs, at least not in the way it changes the default interpretation of constitutive material NPs, which are paradigmatic cases in which the default interpretation is partly determined by NON-VACUITY.

This result is problematic if, following SHIFTING HEADS, you hold that NON-VACUITY is responsible for determining the default interpretations of privative NPs. But it is expected if you accept DUAL CONTENT, according to which the default interpretation of privative NPs is neither triggered nor affected by NON-VACUITY. In this respect, privative NPs behave like complex NPs involving genuine modifiers, such as *dark red paint* and *perfect gun*, which, in most contexts that call for literal interpretations, can only be interpreted as referring to (good examples of) dark red paint and perfect guns. In short, in contexts that call for literal interpretations, privative NPs behave more like paradigmatic cases of NPs involving genuine modifiers than like paradigmatic cases of NPs involving pragmatic meaning shifts of the head Ns. This is exactly what is predicted by DUAL CONTENT. Hence the case of L-privative NPs strongly supports DUAL CONTENT OVER SHIFTING HEADS.²⁶

²⁶ This criticism of SHIFTING HEADS extends, *mutatis mutandis*, to other contextualist accounts of privative NPs. For example, one might propose replacing NON-VACUITY with other, more traditional Gricean maxims of interpretation, such as BE TRUTHFUL. However, essentially the same criticism applies to Gricean reconstructions of SHIFTING HEADS, since they still depend on the idea that the default interpretation of privative NPs involves a pragmatically induced meaning shift of the head Ns. This assumption is precisely what is in conflict with the behavior of L-privative NPs, especially when compared with the behavior of L-constitutive material NPs.

6 An attempt to save SHIFTING HEADS

To defend SHIFTING HEADS, one might suggest that the lack of analogy between the interpretation of L-privative NPs and L-constitutive material NPs springs from an important difference in the status of the common-ground beliefs relevant to each case.

Here is the idea. Most speakers believe that there are no stone lions₁ and fake guns₁. The first belief is dropped from the common ground when we are instructed to take things literally, even if the resulting assertion might seem unbelievable. This explains cases like (27) and (28). By parallel reasoning, it would seem that the belief that there are no fake guns₁ should be dropped from the common ground in cases such as (31) and (32). However, maybe the belief that there are no fake guns₁ is a more entrenched belief than the belief that there are no stone lions₁. There are possible worlds in which there are stone lions₁; but there are no possible worlds in which there are fake guns₁. We can't believe that there are fake guns₁ in something like the sense that we can't believe that there are married bachelors. So whatever we do to make sense of expressions such as (31) and (32), what we cannot do is drop from the common ground the belief that there are no fake guns₁. This explains why expressions like (31) and (32) feel genuinely redundant: the contextual instruction to take the privative NP literally does not change its meaning relative to the default cases.

However, this response is based on a fundamental mistake: interpreted in the way suggested by SHIFTING HEADS, *fake gun*₁ is not an empty predicate, nor is there reason to think that mature speakers believe that it is.

To evaluate this response, we need to consider possible lexical entries for *fake*. Partee herself does not pronounce on this, but others who accept her account of privatives assume that they are first order predicates.²⁷ What could the meaning of *fake*, taken as a first-order predicate, be? Initially, let's assume a simple atomic account:

$$(33) \quad \llbracket fake \rrbracket = \lambda x.FAKE(x)$$

The point of (33) is just to capture the idea, which seems correct, that when we use *a fake* predicatively, as in *everything that they sell at that store is a fake*, we mean something like 'a fake thing of some sort'. As we said before, DUAL CONTENT is compatible with the fact that we sometimes use *fake* in this way;

²⁷ For example, this is how Recanati (2010, Chapter 2) approvingly interprets Partee's account of privative NPs.

but unlike the contextualist accounts, it does not assume that this is how we use it in the default interpretation of privative NPs. This proposed response in defense of SHIFTING HEADS assumes that when *fake* is interpreted roughly as (33), the intersection between *fake* and *gun*₁ is empty, or at least that most mature speakers strongly believe this, and are unwilling to drop this belief from the common ground.

Now, here's the crucial question: is it really the case that, so interpreted, there are no fake guns₁? In other words, is it really the case that no guns₁ are also fake things? Consider the following scenarios:

(EVIL PARK OWNER) Mark is an evil paintball gun park owner. He really hates John. Unaware of this secret hate, John goes to Mark's park to play. Mark comes up with the plan of secretly replacing John's paintball gun with a fake paintball gun that is actually a real gun, with the intention that John kill someone. Fortunately, when John handles the fake paintball gun he notices something suspiciously off with its weight and refuses to use it.

(EVIL PARTYGOER) Peter is a clever terrorist who wants to cause havoc at a Halloween party. To evade the security measures, he designs a gun that is intended to look like a toy gun but is actually a very powerful gun. Fortunately the security personnel are suspicious of Peter and his fake toy gun, so he gets caught and his plot is averted.

EVIL PARK OWNER presents an object that we appropriately called a *fake paintball gun*: it is an object that was made to look like a paintball gun and feel like a paintball gun, but is not really one, since it was made for and can be used to kill by shooting. Similarly, EVIL PARTYGOER presents an object that we appropriately called a *fake toy gun*: it is an object that was made to look like a toy gun, but is not really a toy gun, since it was made for and can be used to kill by shooting. So these are examples of guns₁ that are also fake objects. Furthermore, these scenarios are not far fetched. Indeed, they are closer to our world than extraordinary scenarios in which MIT scientists create real animals from non-organic materials. If, under instruction to take things literally, we have no hesitation to widen the common ground so that it includes worlds in which there are stone lions₁ and rubber rabbits₁, then under the same instruction, we should have even less hesitation to widen the common ground so that it includes worlds, such as our own, in which there are guns₁ which are also fake things (of some sort).

But this prediction is incorrect: we do not interpret L-privative NPs in the literal sense (e.g., as *fake gun*₁ taken in the sense it would have according to this reconstruction of SHIFTING HEADS) in which we interpret L-constitutive material NPs (e.g., as *stone lion*₁). In other words, these accounts fail to explain why in L-privative NPs we never use, say, *fake gun* to denote things like the subset of fake toy guns or fake paintball guns that are real guns, even in scenarios that call for this reading (such as cases like (31) and (32)).²⁸ This, as we just saw, should be quite easy to do if *fake* when used attributively were really a first-order predicate. Furthermore, assuming that *fake* is a gappy or context-sensitive first-order predicate, as in (34), won't solve this problem.

(34) $\llbracket \textit{fake} \rrbracket = \lambda x. \text{FAKE-AS-P}(x)$, where the value of **P** has to be provided by the context.

In some cases, gappy expressions use the immediate linguistic context to determine the meaning of the open parameter. For example, *small elephant* usually means an elephant that is small for an elephant, but as is well known, in some contexts, *small elephant* can mean an elephant that is of a normal size for an elephant but is small compared with other kinds of giant creatures. If *fake*, when used attributively, was a gappy first-order predicate, then when combined with *gun* it would usually mean 'fake for a gun₁'; but in certain contexts, such as some of the ones mentioned above, it could also mean 'fake for a toy gun' or 'fake for a paintball gun', even when combined with gun₁. So we should get the reading according to which *fake gun* can mean 'fake toy gun that is a real gun', 'fake paintball gun that is a real gun', and so on; but we do not.²⁹

²⁸ Note that this argument does *not* assume that *x*'s being a fake paintball gun entails that it is a fake gun. What it assumes is that *some* fake paintball guns could be real guns. The reason is simple: a fake paintball gun is something that was made to look like a paintball gun but is not really a paintball gun. As in EVIL PARK OWNER, some objects could be deceptively made to look like paintball guns but be in fact real guns.

²⁹ To close this discussion of SHIFTING HEADS, I should add two brief but important points. First, I take this criticism of SHIFTING HEADS to undermine one way of trying to save lexical atomism. Now, although SHIFTING HEADS is one of the most plausible and influential accounts, a full reply to lexical atomism would require that we consider other options. This latter task is beyond the scope of the present work. Second, there is good reason to think that the debate between theories of modification could soon be informed by direct psycholinguistic and neurolinguistic evidence. This is because theories such as DUAL CONTENT and SHIFTING HEADS can be distinguished by the way they classify, in terms of their coercion patterns, privative NPs and constitutive material NPs (and other subclasses). According to DUAL CONTENT, the default interpretations of expressions like *fake gun* do not involve

7 DUAL CONTENT: Objections and clarifications

To conclude this discussion of DUAL CONTENT, I want to make some clarifications and address some important objections. I will first discuss the relation between DUAL CONTENT and contextualism, and then the notions of E-structure, C-structure, and their relation.

7.1 IS DUAL CONTENT compatible with contextualism?

Contemporary formal semantics has moved in the direction of contextualism. To mention a few examples, there are influential accounts of quantifiers (Szabó & Stanley 2000), modals (Kratzer 2012), predicative and attributive adjectives (Rothschild & Segal 2009, Kennedy & McNally 2010, Recanati 2010), and noun-noun compounds (Weiskopf 2007) that posit context-sensitive parameters in logical form. There are also various arguments for the view that the meaning of expressions in general, and not only those that have context-sensitive parameters, is usually modulated during online processing (Carston 1997, Recanati 2004, 2010). To accommodate these accounts, recent discussions of compositionality suggest that in modeling our linguistic competence we should think of compositionality as the principle that the meaning of token expressions is determined from the contextually saturated and modulated meaning of their constituents (Pagin & Pelletier 2007, Szabó 2010, Recanati 2010).

It is important to clarify the place of DUAL CONTENT in this general debate, especially since we criticized the main contextualist account of privative NPs. To do so, it is useful to introduce some new terminology, essentially following Recanati 2010. First, since lexical items can have free parameters, we will say that, for any expression e ,

$$(35) \llbracket e \rrbracket_c = f_e(c),$$

where f_e is the character of e and $f_e(c)$ is the occasion meaning of e in context c . If e has no free parameters, then for all c 's, $f_e(c) = m$, where m is the standing meaning of e . Secondly, since the meaning of tokens is

coercion, but those for *stone lion* and *fake lion* do (in the latter case, because a TELIC has to be artificially provided). According to SHIFTING HEADS, the default interpretations of *fake gun*, *fake lion*, and *stone lion* all involve coercion. Recent efforts to find the neural correlates and signatures of operations of coerced meaning modulation have been quite successful (e.g., Pylkkanen & McElree (2006) and Pylkkanen, Brennan & Bemis (2013)). This justifies some optimism that these accounts could soon be evaluated using direct neurolinguistic data.

determined from the meanings of the modulated (instead of the standing) meanings of their immediate constituents, we can represent the modulated meanings of an expression e , $\llbracket e \rrbracket_{M,c}$, as

$$(36) \quad \llbracket e \rrbracket_{M,c} = \text{mod}(e, c)(\llbracket e \rrbracket_c),$$

where *mod* takes as an argument an expression e and context c in which e occurs and returns as value the modulation function $f_{M,e}$, which takes $\llbracket e \rrbracket_c$ and returns the meaning that is salient/relevant/appropriate for e in c .³⁰

For simplicity, assume that traditional *FA* is the main compositional operation. We can incorporate these two forms of context-sensitivity by reformulating *TN* and *FA* as follows (assuming that all non-branching nodes are terminal nodes):

(*TN*) If α is a terminal node, then $\llbracket \alpha \rrbracket_{M,c} = \text{mod}(\alpha, c)(\llbracket \alpha \rrbracket_c)$, where $\llbracket \alpha \rrbracket_c$ is specified in the lexicon.

(*FA*) If α is a branching node, $\{\beta, \gamma\}$ is the set of α 's daughters, and $\llbracket \beta \rrbracket_{M,c_1}$ is a function whose domain contains $\llbracket \gamma \rrbracket_{M,c_2}$, then $\llbracket \alpha \rrbracket_{M,c} = \llbracket \beta \rrbracket_{M,c_1}(\llbracket \gamma \rrbracket_{M,c_2})$

According to this implementation, *mod* does not operate on the outputs of *FA*, but only on terminal nodes or lexical items. This allows a constrained form of meaning modulation.³¹

DUAL CONTENT is compatible with this general contextualist framework. The revised version of *FA* we introduced in Section 3 — FA^{DC} — can be easily reformulated to allow for these two forms of context-sensitivity. But what is important to note, for our purposes, is that unlike other accounts of privative

³⁰ Different contextualist accounts can further restrict *mod* in different ways. That some such restriction is required seems to me uncontroversial; otherwise this sort of account will over-predict “reasonable” meaning shifts that are actually not available. The over-generation problem faced by unrestricted contextualist accounts is illustrated with several interesting examples in Asher 2011 and, focusing on color words, in Kennedy & McNally 2010. For our purposes, we can ignore further details of how *mod* should be constrained.

³¹ There are at least two ways of implementing the context-sensitive *mod* function to get compositional rules that determine the meanings of complex expressions in terms of the modulated meanings of their constituents. On the unconstrained implementation, *mod* is generalized to apply at every level of interpretation. On the constrained implementation — which is the one we adopt here — *mod* applies only to lexical items (i.e., terminal nodes). Recanati 2010 and Pagin & Pelletier 2007 favor the more unconstrained version. For our purposes, this choice is open, but for detailed reasons to favor the constrained implementation, see Del Pinal 2014.

NPs, DUAL CONTENT does not assume that either privative Adjs or common Ns have context-sensitive free parameters. It also does not assume that, when tokens of privative NPs are interpreted in the default way, the standing and modulated meanings of the token Adjs and Ns differ significantly. In this sense, according to DUAL CONTENT, the default interpretation of privative NPs is not context-sensitive, or rather, does not involve significant meaning modulations of the head Ns.

That does *not* mean, however, that there is *never* any meaning modulation involved in the interpretation of privative and other NPs involving modifiers that are treated as semantic restructuring operators. Consider the following example:

(BAD HUSBAND) It was recently revealed that Smith, a supposedly ideal husband, had an illicit love affair. Ann, a famous town gossip, tells Sam, “this does not surprise me at all. Smith is a typical man.”

To properly understand Ann’s claim that Smith is a *typical man* we have to modulate the C-structure of *man*—specifically the TELIC—to include cheating as one of the things that a man does. For most of us, this is a sufficiently salient cultural stereotype and the required modulation is easy. Once the C-structure of *man* is enriched in this way, we can continue to treat *typical* as a semantic restructuring operator, along the lines of (13) above. This, in turn, would explain why in BAD HUSBAND Ann is asserting about Smith that he is a cheat. Examples like this are common. What is important to note, however, is that allowing the general possibility of meaning modulation does not affect the basic idea of DUAL CONTENT, namely, that some Adjs partly operate on the C-structure of Ns.

Indeed, any theory which admits something like C-structure, such as DUAL CONTENT, should be made compatible with some form of contextualism. As is well known, there is plenty of evidence showing that the prototypes associated with terms are modulated as a function of features such as preceding discourse and task demands (Roth & Shoben 1983, Barsalou 1987). For example, when *animal* is processed in the context of *milk*, there is a prototype effect (e.g., increased ease of categorization) for *cow* and *goat* relative to *horse* and *mule*, but when *animal* is processed in the context of *riding* the effect goes in the other direction, favoring *horse* and *mule* (Roth & Shoben 1983).³² This sort of finding has been so often replicated, that it is fair to say

³² Importantly, the effects of context on prototypes can be long-lasting. For example, in a classic experiment by Barclay et al. (1974), one group of subjects was asked to encode sentences

that the relative instability of lexical prototypes is by now a well established tenet in cognitive psychology (Roth & Shoben 1983, Barsalou 1987, Murphy 2002). Insofar as the C-structure posited by DUAL CONTENT incorporates something like enriched prototypes, we should also expect that it undergoes some modulation across contexts. These considerations should be kept in mind when judging the descriptive adequacy of proposed lexical entries for privative and subsective Adjs and common Ns.

7.2 Why not include functional and genealogical information in the E-structure of lexical artifact terms?

Consider the following objection to DUAL CONTENT. As we argued in Section 2, Putnam and Kripke gave us good reasons to reject accounts of natural kinds that assume that their extension is determined by associated descriptions, including definitions and stereotype information. So if we need something like enriched stereotype information to fully account for our linguistic competence, as we argued above, it follows that we need to make a distinction between two components of meaning, E-structure and C-structure. However, one could argue that the reasons to reject the idea that definitions determine the extension of lexical terms do not apply to the case of artifacts. Indeed, some linguists and philosophers have argued that, in general, artifact kinds can be defined in terms of their functional and/or genealogical properties (Thomasson 2003). If this is correct, it might seem reasonable to incorporate the TELIC or AGENTIVE of artifact terms into their E-structure. In turn, this would mean that we might be able to account for privative NPs without assuming that common Ns have a dual semantic content which includes a non-extension-determining C-structure.

This is an interesting objection, but it is not hard to see that lexical artifact terms cannot, in general, be defined in the way proposed. For example, we know that there are malfunctioning guns and antique or display guns which cannot shoot. There are computers which have ceased to work and cars which do not even have an engine. Some erasers do not erase well, and some dining tables are used as computer desks but do not thereby change their identity. Most chairs are used for sitting but not all items used for sitting are chairs,

such as *the man lifted the piano*, and another group was asked to encode *the man tuned the piano*. On a subsequent test, *something heavy* was a more effective cue for recalling *piano* than *something with a nice sound* for the first group of subjects, and the opposite pattern was observed for the second group.

just as not all things that can be used to shoot are guns. In all these cases and countless more, we distinguish between artifacts and their typical function, so the TELIC cannot in general be part of E-structure. These intuitions are supported by some experimental results (Keleman & Carey 2007). In one representative example, subjects were shown teapots used to water plants, and lampshades used as umbrellas. Most subjects still choose *not* to name these objects based on their current use (Rips 1989, Matan & Carey 2001).

We can also easily distinguish between artifacts and the function for which they were created. Suppose that, in some remarkable situation, you witness the spontaneous emergence of a gun or a pencil. You handle both objects and they have their usual functions: to shoot and write. It is appropriate to call these objects *gun* and *pencil*, despite their lacking the typical AGENTIVE of guns and pencils. Indeed, talking about the spontaneous generation of artifacts does not seem to involve any kind of incoherence, unlike talking about, say, unperceived sensations. Furthermore, not all artifacts that are made for shooting are guns, just as not all artifacts that are made for sitting are chairs. That the typical AGENTIVE of simple artifact terms is not part of E-structure is also suggested by experimental studies (Chaigneau & Barsalou 2001). For example, Chaigneau (2002) conducted an experiment in which subjects were given information about the intention for which an object was created. Subjects were also given information about the object's physical structure and current use. In some scenarios, the creator's intention and the actual use of the object differed. For many kinds of artifacts, most subjects choose to individuate them based on actual use. For example, objects created for a non-mopping purpose but which were successfully used for mopping, were categorized by most subjects on the basis of their actual, as opposed to the creator's intended, use.

Now, to say that the TELIC and AGENTIVE are not components of the E-structure of lexical artifact terms is not to say that they are not part of the E-structure of *complex* artifact terms in certain constructions. That they are inherited into the E-structure of certain complex terms such as *perfect gun* and *counterfeit Rolex* is precisely the point of DUAL CONTENT.³³ In addition, we

³³ A reviewer asks, if we concede that, say, guns can be created *ex nihilo*, then why not say the same thing about *fake guns*, and so deny that AGENTIVE information is represented in its E-structure? In response, note, first, that the point is not obviously generalizable. For example, assume that, as we argued, *gun* does not include TELIC information in its E-structure; it clearly does not follow that the same holds for *perfect gun*—as we said above, intuitively, a perfect gun is one that, among other things, shoots perfectly. Likewise, we can agree that a tiger might not have the usual striped coat, but *typical tiger* arguably

are not denying that in most cases the typical TELIC and AGENTIVE associated with lexical artifact terms plays a crucial role in categorization. Indeed, some prominent developmental psychologists argue that, in the case of artifacts, beliefs about the creator's functional intention is the most important categorization criterion (Bloom 1996, 1998, Keleman & Carey 2007). Other cognitive psychologists argue that the most important feature is instead the TELIC (Chaigneau & Barsalou 2001). Perhaps most persuasively, Malt & Sloman (2007) and Sloman & Malt (2003) propose, partly on the basis of an extensive evaluation of the literature, that which features are most important for artifact categorization (among form, function and genealogy) depends on aspects of the relevant context, including task demands. These proposals are compatible with DUAL CONTENT as long as they accept that features of lexical terms that are important for categorization are, in most cases, only part of C-structure, as most cognitive psychologists would accept in any case.³⁴

7.3 What then is E-structure? Why not just drop that notion altogether?

At this point, one might wonder why we should maintain that lexical artifact terms have an E-structure at all. The reason for this, however, is basically the same as for why we should maintain the distinction between E-structure and associated descriptions in the case of natural kind terms. The standard arguments are well known, so I will only briefly rehearse them to show that

includes that information in its E-structure. According to DUAL CONTENT, certain complex expressions bring into their E-structure information that was only part of the C-structure of the constituents. So a testable prediction of this view is that, in the case of privatives such as *fake* and *counterfeit*, speakers will exhibit stronger and less flexible intuitions about the AGENTIVE requirements for something to qualify as a fake or counterfeit artifact than they would for artifacts in general.

³⁴ To properly account for categorization patterns, the representations of the C-structure of lexical terms have to be supplemented with weights on each quale, reflecting their relative importance. This is compatible with DUAL CONTENT, but it would introduce technical complications which are orthogonal to the main aim of this paper. I have ignored information about the relative weights of the dimensions of C-structure because it does not seem to play a key compositional role in privative NPs. For example, to capture the difference between *fake*, *counterfeit*, and *artificial*, we had to refer, in their E-structure, to qualia functions that directly access different dimensions of the C-structure of the N they modify (and there was no need to further operate on their weights). One might think that the compositional behavior of privatives can be captured by operating just on the weights of the modified Ns (e.g., by inverting them), without direct reference to their quale. However, it seems impossible to capture the difference between *fake*, *counterfeit*, and *artificial* by operating just on the weights of the modified Ns.

they also apply to artifact terms. The aim here is only to show that the case for assuming an E-structure independent of associated beliefs is at least as strong for artifacts as it is for natural kind terms.³⁵

The main reason for rejecting descriptive theories and keeping a directly referential component stem from the *problem of ignorance* and the *problem of error* (Putnam 1970, Burge 2007, Kripke 1980, Kobes 1989, Weiskopf 2009). The problem of ignorance is that speakers can refer to a kind even if they do not know any descriptions which uniquely pick out the kind. The problem of error is that speakers can refer to a kind even if most of their beliefs about the kind are incorrect. We can talk about elms and beeches even if we do not know much about their differences, and we can discuss the basic properties of molybdenum even if we are radically mistaken about what they are. Reference to artifactual kinds exhibits the same patterns (Kornblith 2007). We can talk about amulets and nuclear reactors even if we only have a vague idea of their form and function, one which might not distinguish them from similar artifacts. We can talk about rheostats and spandrels even if we are ignorant of or entirely mistaken about their basic function. Just like in the case of natural kind terms, we may acquire these terms from other speakers and use them with the intention to co-refer. This division of linguistic labor explains why the ability to refer is compatible with such ignorance and error. In our terms, this is just to say that we need a notion of E-structure, independent of their C-structure, to account for our referential competence with lexical artifact terms.³⁶

Now, this story faces the same challenge faced in the case of natural kind terms, namely, that we do not yet have an accepted account of what

³⁵ For a more detailed response along these lines, to which the following discussion is indebted, see Kornblith 2007.

³⁶ One could insist that there is an important difference between natural kind and artifact terms which might bear on whether we need to posit something like E-structure in the latter case. Plausibly, the rationale for preserving an E-structure for natural kinds presupposes that they have an extension which is independent of our current beliefs. The parallel presupposition might not seem plausible in the case of artifacts. As one anonymous reviewer puts it, “we may lack definitions for artifacts, but it seems that it cannot be that most of our beliefs about a given artifact are mistaken”. However, I would dispute that there is a relevant distinction in this respect between artifact and natural kinds. Consider, for example, historical artifacts. Archaeologists might discover some objects and classify them as ancient grain storerooms, only to then discover that they are underground temples. This situation is in fact quite common: we discover an artifact, refer to the class via samples, are massively confused about their TELIC and AGENTIVE, and try to discover new facts to improve and perhaps radically change our current beliefs about the kind.

ultimately determines the extension of the kinds. Combined with the idea that descriptions or C-structure do not in general determine reference, this is part of the reason why we took the E-structure of common Ns as semantically atomic, and their relation to their extension as primitive. We can still ask, of course, how the extension of terms is determined. As we mentioned above, some authors suggest that, for artifacts, it might be something like the intentions of their creators or producers. But, again, it is crucial to note that whatever story we give along these lines about the role of intended function in reference determination, it is going to be much more complicated and abstract than the information most speakers actually include in the AGENTIVE of artifact terms, which is usually just the *typical* reason or way in which the artifact came into being. So even if something like the producers' intentions fixes the correctness conditions of artifact kinds, this is no reason for suggesting that, for simple lexical artifactual terms, we should include the AGENTIVE or TELIC as part of E-structure. This is so even in cases where radical ignorance and error are rare, as in the case of common artifacts such as guns and pencils.³⁷

³⁷ For example, Thomasson (2003, 2007) argues that the essence of artifacts lies in the intentions of their makers:

Necessarily, for all x and all artifactual kinds K , x is a K , only if x is the product of a largely successful intention that Kx , where one intends Kx only if one has a substantive concept of the nature of K s that largely matches that of some group of prior makers of K s (if there are any) and intends to realize that concept by imposing K -relevant features on the object. (Thomasson 2003)

A consequence of this constitutive condition is that “reference to artifactual kinds... cannot proceed without someone (namely, those responsible for the production and reproduction of these artifacts) having a substantive concept of the nature of the kind”, where the notion of the “nature of the kind” is fleshed out in terms of its intended function (Thomasson 2007).

We might grant that there is some essential or strong connection between producers' or conceivers' intentions and the extensions of artifact terms, but it seems clear that Thomasson's proposal is too strong. To illustrate, note that it entails, counterintuitively, that in the following scenarios the objects we are talking about are not “really” guns.

(COPY) Little Mark is a precocious infant with remarkable craftsmanship skills. He finds some very detailed instructions to make an object he has never seen before, which might have some function that is unspecified in the instructions. He somehow gets the materials and follows the instructions perfectly. His intention is only to execute the instructions; he does not really care or even wonder much about what the function of the object might be. Little Mark's father then enters the room only to find, much to his dismay, that Mark constructed a gun.

7.4 What is the relation between E-structure and C-structure?

To talk about the relation between E-structure and C-structure, we should consider the extension of each component, where we can take the C-structure of a term as a complex first order predicate that is the conjunction of all its quale values.³⁸ There are five ways in which the extension of the E-structure (call this set E) and that of the C-structure (C) of artifact and natural kind lexical terms could be related (see Weiskopf 2009):

- i. $E = C$
- ii. $E \subset C$
- iii. $C \subset E$
- iv. $[E \cap C \neq \emptyset] \wedge \neg[(E \subset C) \vee (C \subset E)]$
- v. $E \cap C = \emptyset$

The theories that first introduced a distinction between cognitive and referential content assumed that the former fully determines the latter. These theories conform to case (i). Prominent examples are the Fregean theory of sense and the Millian theory of names. Recent dual content theories, especially in cognitive psychology, tend to hold that cognitive content does not in general determine referential content (Prinz 2002, Murphy 2002, Prinz 2012, Weiskopf 2009). DUAL CONTENT belongs to this class. For many artifact

(ACCIDENT) Strange accidents and remarkable things happen in the magical world of HH. One day, the toy factory production line somehow popped out a series of guns. After a series of unusual accidents, everyone in HH was frightened and demanded an investigation. After arduous work by the top detectives of HH, it was discovered that no one intended that this be so. It was merely a very strange accident due to some unintended glitch in the production line.

In COPY and ACCIDENT we can talk about the relevant objects as being real guns, despite their failing to satisfy Thomasson's condition, since their proximal creators did not intend that they be objects of the gun kind. One might reformulate Thomasson's story to deal with these sorts of cases. But the only point I am pressing is that these reference fixing conditions are quite complicated and independent of speakers' beliefs. In this sense, they are different from the information typically incorporated in the AGENTIVE of the C-structure of artifact terms.

³⁸ We can determine the extension of C-structure in other ways, e.g., by considering disjunctions involving each quale, or by combining each quale with weights and appealing to a threshold function. The points I make below also apply to these alternative ways of determining the extension of the C-structure of terms.

terms, the set of objects that satisfies their C-structure is usually a subset of the set of objects that satisfies the E-structure (case (iii)). For example, the C-structure assigned to *gun* in (8) above picks out guns that look and function like typical guns, but, again, some guns are not typical in this way. C-structure can also sometimes overextend (case (ii)). For example, many speakers might have a C-structure for *couch* that extends to sofas. In short, for most artifacts we can only hold that C and E-structure have some extensional overlap (case (iv)). However, we saw above that there are important exceptions, including cases in which we are entirely wrong about the properties of a particular artifact term (case (v)). Following Weiskopf (2009), we can call terms whose C- and E-structures satisfy (ii), (iii) or (iv) *adequate*. *Fully adequate* terms satisfy (i); and terms that only satisfy (v) are *inadequate*.

For most artifact and natural kind terms, then, the relation between their E-structure and C-structure is relatively loose. This does not mean, however, that the two components are independent. There are important constraints that govern their relation. I cannot discuss this issue in detail here, but I will mention two plausible candidates that should be further explored: these are constraints regarding *accurate* and *efficient* representation (Weiskopf 2009). *Accurate representation* captures the idea that one aim of our conceptual system is to achieve full adequacy of C-structures with respect to E-structures — that is, we aim to acquire beliefs about the properties of kinds which are accurate. *Efficient representation* captures the idea that this has to be done under certain limitations of time, information, working memory and processing power. The interplay between these two constraints means that C-structure is often allowed a certain degree of inaccuracy, especially in cases when the cost of the inaccuracies is low and the cost of achieving full adequacy is high.

Having emphasized the distinction, and the sometimes loose relation, between E-structure and C-structure, it is important to recall that the reason why C-structure is part of the *linguistic* meaning of lexical items is that, if DUAL CONTENT is correct, it plays a crucial compositional role in certain common constructions, including privative NPs such as *fake gun* and *counterfeit document* and subsecutive modifications such as *powerful gun*, *perfect knife*, and *typical chair*. In all these cases, aspects of the C-structure of the constituents enter into the E-structure of the resulting complex NPs via purely compositional operations. It is because of this combinatorial dynamic that DUAL CONTENT can claim that C-structure is part of linguistic meaning.

This highlights an important contrast between DUAL CONTENT and its predecessors: most previous theories do not show in detail how the C-structure of lexical terms can affect the E-structure of complex terms of which they are constituents. Indeed, dual content theorists tend to assume that the composition of E-structure and C-structure proceeds along parallel but completely independent routes (see e.g. Prinz 2002 and Weiskopf 2009):

When concepts combine, they contribute both their cognitive and referential contents. . . . The familiar semantic rules of composition in each case produce, in parallel, complex contents from the contents of the constituents. R-contents of concepts combine to form R-contents of judgments, and similarly for C-contents, although neither C-contents nor R-contents determine the other. (Weiskopf 2009, p. 145)³⁹

Weiskopf's notions of "referential" and "cognitive" content correspond, roughly, to our E-structure and C-structure. From the perspective of DUAL CONTENT, holding to the view of parallel and independent composition of E-structure and C-structure blocks one of the most useful tools available to dual content theories, namely, semantic restructuring operators and the corresponding dynamic view of compositionality. If DUAL CONTENT is correct, certain constructions use the C-structure of constituents to determine the E-structure of complexes. It follows that we cannot dismiss the notion of C-structure on the grounds that, although it might be needed to account for certain psychological facts such as typicality patterns, it is not needed to account for purely linguistic facts.

8 Conclusion

According to DUAL CONTENT, common Ns have internal semantic structure, which consists of E-structure and C-structure. I have shown how we can use this structure to provide a satisfactory compositional account of privative NPs (a traditionally problematic class of constructions) as well as of other interesting subsective and intersective NPs. DUAL CONTENT is not the *only* plausible account of privative NPs. As we saw, Partee's (2007, 2010) SHIFTING HEADS is compatible with atomic lexical semantic theories. Theorists keen on

³⁹ I should note that, aside from this claim about compositionality, there is much that I agree with in Weiskopf's (2009) excellent defense of dual content theories.

rejecting the view that lexical terms have internal semantic structure, or that compositional operations have access to this structure, might be tempted to adopt SHIFTING HEADS. I showed, however, that SHIFTING HEADS and similar contextualist accounts are fundamentally mistaken. Still, given the potential variety of accounts of privative NPs compatible with lexical atomism, the crucial question is this: Why insist on constructing accounts of complex NPs that don't require that Ns have a complex semantic structure to which the compositional operations are sensitive?

If the representations and operations introduced by DUAL CONTENT could only be used to deal with privative NPs, it might be wise to try to explain these NPs in a more traditional way. But this is clearly not the case. On the one hand, DUAL CONTENT employs representations of the meaning of common Ns which are very close to those used to represent lexical concepts in recent influential psychological theories. These representations play a crucial role in accounts of language acquisition and fast and slow categorization. On the other hand, powerful accounts in the family of DUAL CONTENT have been presented to deal with many types of problematic complex expressions (Busa & Bouillon 2001, Vikner & Jensen 2002, Leslie 2014, Pustejovsky 1995, 2012, McNally 2006, Pinker 2007, van Tiel 2014, Sassoon 2011, Knobe, Prasada & Newman 2013). All in all, I think it is time to expand our basic formal semantic framework to include compositional operations which are sensitive to the internal semantic structure of lexical terms, including their non-extension-determining C-structure. In this way, we will begin to integrate compositional formal semantics with a psychologically motivated and realistic lexical semantics.

References

- Asher, Nicolas. 2011. *Lexical meaning in context: A web of words*. Cambridge, UK: Cambridge University Press.
- Barclay, J. R., John D. Bransford, Jeffery J. Franks, Nancy S. McCarrell & Kathy Nitsch. 1974. Comprehension and semantic flexibility. *Journal of Verbal Learning and Verbal Behavior* 13(4). 471-481. [http://dx.doi.org/10.1016/S0022-5371\(74\)80024-1](http://dx.doi.org/10.1016/S0022-5371(74)80024-1).
- Barsalou, Lawrence W. 1987. The instability of graded structure: Implications for the nature of concepts. In U. Neisser (ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization*, 101-140. Cambridge, UK: Cambridge University Press.

- Bloom, Paul. 1996. Intention, history, and artifact concepts. *Cognition* 60(1). 1-29. [http://dx.doi.org/10.1016/0010-0277\(95\)00699-0](http://dx.doi.org/10.1016/0010-0277(95)00699-0).
- Bloom, Paul. 1998. Theories of artifact categorization. *Cognition* 66(1). 87-93. [http://dx.doi.org/10.1016/S0010-0277\(98\)00003-1](http://dx.doi.org/10.1016/S0010-0277(98)00003-1).
- Bloom, Paul. 2002. *How children learn the meaning of words*. Cambridge, MA: MIT Press.
- Bloom, Paul. 2007. Water as an artifact kind. In Eric Margolis & Stephen Laurence (eds.), *Creations of the mind: Theories of artifacts and their representation*, chap. 9, 150-156. Oxford: Oxford University Press.
- Burge, Tyler. 2007. *Foundations of mind*. Vol. 2. Oxford: Oxford University Press.
- Busa, Federica & Pierrette Bouillon (eds.). 2001. *The language of word meaning*. Cambridge: Cambridge University Press.
- Carston, Robyn. 1997. Enrichment and loosening: Complementary processes in deriving the proposition expressed. *Linguistische Berichte* 8. 103-127. http://dx.doi.org/10.1007/978-3-663-11116-0_7.
- Chaigneau, Sergio E. 2002. *Studies in the conceptual structure of object function*. Atlanta, GA: Emory University PhD thesis.
- Chaigneau, Sergio E. & Lawrence W. Barsalou. 2001. The role of function in categories. *Theoria et Historia Scientiarum*.
- Coulson, Seana & Gilles Fauconnier. 1999. Fake guns and stone lions: Conceptual blending and privative adjectives. In B. Fox, D. Jurafsky & L. Michaels (eds.), *Cognition and function in language*. Palo Alto, CA: CSLI.
- Del Pinal, Guillermo. 2014. The structure of semantic competence: Compositionality as an innate constraint of the Faculty of Language. *Mind & Language* forthcoming.
- Fodor, Jerry. 1998. *Concepts: Where cognitive science went wrong*. Oxford: Oxford University Press.
- Fodor, Jerry. 2008. *LOT 2: The language of thought revisited*. Oxford: Oxford University Press.
- Fodor, Jerry & Ernest Lepore. 2002. *The compositionality papers*. Oxford: Oxford University Press.
- Franks, Bradley. 1995. Sense generation: A quasi-classical approach to concepts and concept combination. *Cognitive Science* 19(4). 441-505. http://dx.doi.org/10.1207/s15516709cog1904_2.
- Hampton, James A. & Martin L. Jonsson. 2012. Typicality and compositionality: The logic of combining vague concepts. In Marcus Werning, Wolfram

- Hinzen & Edouard Machery (eds.), *The Oxford handbook of compositionality*, chap. 18, 385–402. New York, NY: Oxford University Press.
- Harley, Heidi. 2012. Lexical decomposition in modern syntactic theory. In Markus Werning, Wolfram Hinzen & Edouard Machery (eds.), *The Oxford handbook of compositionality*, chap. 15, 328–350. Oxford: Oxford University Press.
- Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Oxford: Blackwell Publishers Ltd.
- Hinzen, Wolfram. 2012. Syntax in the atom. In Markus Werning, Wolfram Hinzen & Edouard Machery (eds.), *The Oxford handbook of compositionality*, chap. 16, 351–370. Oxford: Oxford University Press.
- Hinzen, Wolfram, Edouard Machery & Markus Werning (eds.). 2012. *The Oxford handbook of compositionality*. Oxford: Oxford University Press.
- Hogeweg, Lotte. 2012. Rich lexical representations and conflicting features. *International Review of Pragmatics* 4. 209–231. <http://dx.doi.org/10.1163/18773109-00040205>.
- Israel, Michael. 2002. Literally speaking. *Journal of Pragmatics* 34(4). 423–432. [http://dx.doi.org/10.1016/S0378-2166\(01\)00048-0](http://dx.doi.org/10.1016/S0378-2166(01)00048-0).
- Keil, Frank C., Marissa L. Greif & Rebekkah S. Kerner. 2007. A world apart: How concepts of the constructed world are different in representation and in development. In Eric Margolis & Stephen Laurence (eds.), *Creations of the mind: Theories of artifacts and their representation*, 231–248. Oxford: Oxford University Press.
- Keleman, Deborah & Susan Carey. 2007. The essence of artifacts: Developing the design stance. In Eric Margolis & Stephen Laurence (eds.), *Creations of the mind: Theories of artifacts and their representation*, chap. 12, 212–230. Oxford: Oxford University Press.
- Kennedy, Chistopher & Louise McNally. 2010. Color, context, and compositionality. *Synthese* 174. 79–98. <http://dx.doi.org/10.1007/s11229-009-9685-7>.
- Knobe, Joshua, Sandeep Prasada & George E. Newman. 2013. Dual character concepts and the normative dimension of conceptual representation. *Cognition* 127. 242–257. <http://dx.doi.org/10.1016/j.cognition.2013.01.005>.
- Kobes, Bernard W. 1989. Semantics and psychological prototypes. *Pacific Philosophical Quarterly* 70(1). 1–18.
- Kornblith, Hilary. 2007. How to refer to artifacts. In Eric Margolis & Stephen Laurence (eds.), *Creations of the mind: Theories of artifacts and their representation*, chap. 8, 138–149. Oxford: Oxford University Press.

- Kratzer, Angelika. 2012. *Modals and conditionals*. Oxford: Oxford University Press.
- Krifka, Manfred, Francis Jeffrey Pelletier, Gregory N. Carlson, Alice ter Muelen, Gennaro Chierchia & Godehard Link. 1995. Genericity: An introduction. In Gregory N. Carlson & Francis J. Pelletier (eds.), *The generic book*, chap. 1, 1-124. Chicago, IL: The University of Chicago Press.
- Kripke, Saul. 1980. *Naming and necessity*. Cambridge, MA: Harvard University Press.
- Lakoff, George & Mark Johnson. 1980. *Metaphors we live by*. Chicago, IL: The University of Chicago Press.
- Leslie, Sarah-Jane. 2014. "Hillary Clinton is the only man in the Obama administration": Dual character concepts, generics and gender. *Analytic Philosophy* forthcoming.
- Malt, Barbara C. & Steven A. Sloman. 2007. Artifact categorization: The good, the bad, and the ugly. In Eric Margolis & Stephen Laurence (eds.), *Creations of the mind: Theories of artifacts and their representation*, chap. 6, 85-123. Oxford: Oxford University Press.
- Margolis, Eric. 1998. How to acquire a concept. *Mind & Language* 13(3). 347-369. <http://dx.doi.org/http://dx.doi.org/10.1111/1468-0017.00081>.
- Margolis, Eric & Stephen Laurence. 2007. *Creations of the mind: Theories of artifacts and their representation*. Oxford: Oxford University Press.
- Matan, Adee & Susan Carey. 2001. Developmental changes within the core of artifact concepts. *Cognition* 78(1). 1-26. [http://dx.doi.org/10.1016/S0010-0277\(00\)00094-9](http://dx.doi.org/10.1016/S0010-0277(00)00094-9).
- McNally, Louise. 2006. Lexical representation and modification within the noun phrase. *Recherches Linguistiques de Vincennes* 34. 191-206. <http://dx.doi.org/http://rlv.revues.org/1395>.
- Moravcsik, Julius. 1998. *Meaning, creativity, and the partial inscrutibility of the human mind*. Stanford, CA: CSLI Publications.
- Morzycki, Marcin. Forthcoming. *Modification*. Cambridge, UK: Cambridge University Press.
- Murphy, Gregory L. 2002. *The big book of concepts*. Cambridge, MA: MIT Press.
- Pagin, Peter & Jeffrey Pelletier. 2007. Content, context, and composition. In Gerhard Preyer & Georg Peter (eds.), *Context-sensitivity and semantic minimalism: New essays on semantics and pragmatics*, 25-62. Oxford: Oxford University Press.
- Partee, Barbara. 2007. Compositionality and coercion in semantics: The dynamics of adjective meaning. In Gerlouw Bouma, Irene Kraemer & Joost

- Zwarts (eds.), *Proceedings of the colloquium "Cognitive Foundations of Interpretation"*. Amsterdam: Royal Netherlands Academy of Arts & Sciences.
- Partee, Barbara. 2010. Privative: Subjectives plus coercion? In R. Baule, U. Reyle & T.E. Zimmerman (eds.), *Presupposition and discourse: Essays offered to Hans Kamp*. Amsterdam: Elsevier.
- Partee, Barbara & Hans Kamp. 1995. Prototype theory and compositionality. *Cognition* 57(2). 129–191. [http://dx.doi.org/10.1016/0010-0277\(94\)00659-9](http://dx.doi.org/10.1016/0010-0277(94)00659-9).
- Pinker, Steven. 2007. *The stuff of thought: Language as a window into human nature*. New York, NY: the Penguin Group.
- Prinz, Jesse. 2002. *Furnishing the mind*. Cambridge, MA: MIT Press.
- Prinz, Jesse. 2012. Regaining composure: A defense of prototype compositionality. In Markus Werning, Wolfram Hinzen & Edouard Machery (eds.), *The Oxford handbook of compositionality*, chap. 21, 437–453. Oxford: Oxford University Press.
- Pustejovsky, James. 1995. *The generative lexicon*. Cambridge, MA: MIT Press.
- Pustejovsky, James. 2012. Co-compositionality in grammar. In Markus Werning, Wolfram Hinzen & Edouard Machery (eds.), *The Oxford handbook of compositionality*, chap. 17, 371–384. Oxford: Oxford University Press.
- Putnam, Hillary. 1970. Is semantics possible? *Metaphilosophy* 1(3). 187–201. <http://dx.doi.org/10.1111/j.1467-9973.1970.tb00602.x>.
- Pylkkanen, Liina, Jonathan Brennan & Douglas K. Bemis. 2013. Grounding the cognitive neuroscience of semantics in linguistic theory. *Language and Cognitive Processes* 26(9). 1317–1337. <http://dx.doi.org/10.1080/01690965.2010.527490>.
- Pylkkanen, Liina & Brian McElree. 2006. The syntax-semantic interface: Online composition of sentence meaning. In Mathew J. Traxler & Morton Ann Bersbacher (eds.), *Handbook of psycholinguistics*, 537–577. New York, NY: Elsevier.
- Recanati, François. 2004. *Literal meaning*. Cambridge, UK: Cambridge University Press.
- Recanati, François. 2010. *Truth-conditional pragmatics*. Oxford: Oxford University Press.
- Rey, Georges. 1983. Concepts and stereotypes. *Cognition* 15(1). 237–262.
- Rips, Lance J. 1989. Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (eds.), *Similarity and analogical reasoning*, 21–59. Cambridge, UK: Cambridge University Press.

- Roth, Emilie M. & Edward J. Shoben. 1983. The effect of context on the structure of categories. *Cognitive Psychology* 15(3). 346-378. [http://dx.doi.org/10.1016/0010-0285\(83\)90012-9](http://dx.doi.org/10.1016/0010-0285(83)90012-9).
- Rothschild, Daniel & Gabriel Segal. 2009. Indexical predicates. *Mind & Language* 24(4). 467-493. <http://dx.doi.org/10.1111/j.1468-0017.2009.01371.x>.
- Sassoon, Galit W. 2011. Adjectival vs. nominal categorization processes. *Belgian Journal of Linguistics* 25. 104-147. <http://dx.doi.org/10.1075/bjl.25.06sas>.
- Sloman, Steven A. & Barbara C. Malt. 2003. Artifacts are not ascribed essences, nor are they treated as belonging to kinds. *Language and Cognitive Processes* 18(5-6). 1-20. <http://dx.doi.org/10.1080/01690960344000035>.
- Sperber, Dan & Ira A. Noveck (eds.). 2004. *Experimental pragmatics*. New York, NY: Palgrave Macmillan.
- Szabó, Zoltan. 2010. The determination of content. *Philosophical Studies* 148(2). 253-272. <http://dx.doi.org/10.1007/s11098-008-9296-z>.
- Szabó, Zoltan & Jason Stanley. 2000. On quantifier domain restriction. *Mind & Language* 15(2-3). 219-61. <http://dx.doi.org/10.1111/1468-0017.00130>.
- Thomasson, Amie L. 2003. Realism and human kinds. *Philosophy and Phenomenological Research* 67(3). 580-609. <http://dx.doi.org/10.1111/j.1933-1592.2003.tb00309.x>.
- Thomasson, Amie L. 2007. Artifacts and human concepts. In Eric Margolis & Stephen Laurence (eds.), *Creations of the mind: Theories of artifacts and their representation*, chap. 4, 52-73. Oxford: Oxford University Press.
- van Tiel, Bob. 2014. Embedded scalars and typicality. *Journal of Semantics* 31. 147-177. <http://dx.doi.org/10.1093/jos/fft002>.
- Vikner, Carl & Per Anker Jensen. 2002. A semantic analysis of the English genitive: Interaction of lexical and formal semantics. *Studia Linguistica* 56(2). 191-226. <http://dx.doi.org/10.1111/1467-9582.00092>.
- Weiskopf, Daniel. 2007. Compound nominals, context, and compositionality. *Synthese* 156(1). 161-204. <http://dx.doi.org/10.1007/s11229-005-3489-110.1007/s11229-005-3489-110.1007/s11229-005-3489-110.1007/s11229-005-3489-1>.
- Weiskopf, Daniel. 2009. Atomism, pluralism, and conceptual content. *Philosophy and Phenomenological Research* 79(1). 131-163. <http://dx.doi.org/10.1111/j.1933-1592.2009.00269.x>. <http://www3.interscience.wiley.com/journal/122464216/abstract>.

Dual Content Semantics

Wunderlich, Dieter. 2012. Lexical decomposition in grammar. In Marcus Werning, Wolfram Hinzen & Edouard Machery (eds.), *The Oxford handbook of compositionality*, chap. 14, 307–327. Oxford: Oxford University Press.

Guillermo Del Pinal
Columbia University
Department of Philosophy
708 Philosophy Hall MC4871
1150 Amsterdam Avenue
New York, NY 10027
ged2102@columbia.edu