

Meaning, modulation, and context: a multidimensional semantics for truth-conditional pragmatics

Guillermo Del Pinal^{1,2}

Published online: 17 January 2018
© Springer Science+Business Media B.V. 2018

Abstract The meaning that expressions take on particular occasions often depends on the context in ways which seem to transcend its direct effect on context-sensitive parameters. ‘Truth-conditional pragmatics’ is the project of trying to model such semantic flexibility within a compositional truth-conditional framework. Most proposals proceed by radically ‘freeing up’ the compositional operations of language. I argue, however, that the resulting theories are too unconstrained, and predict flexibility in cases where it is not observed. These accounts fall into this position because they rarely, if ever, take advantage of the rich information made available by lexical items. I hold, instead, that lexical items encode both extension and non-extension determining information. Under certain conditions, the non-extension determining information of an expression *e* can enter into the compositional processes that determine the meaning of more complex expressions which contain *e*. This paper presents and motivates a set of type-driven compositional operations that can access non-extension determining information and introduce bits of it into the meaning of complex expressions. The resulting multidimensional semantics has the tools to deal with key cases of semantic flexibility in appropriately constrained ways, making it a promising framework to pursue the project of truth-conditional pragmatics.

Keywords Contextualism · Modulation · Lexical semantics · Multidimensional semantics · Truth conditional pragmatics · Adjectives · Nouns · Modification

This research was supported by the Alexander von Humboldt foundation.

✉ Guillermo Del Pinal
delpinal@umich.edu

¹ Leibniz Zentrum für Allgemeine Sprachwissenschaft, Schützenstr. 18, 10117 Berlin, Germany

² University of Michigan, Weinberg Institute for Cognitive Science, 914 Weiser Hall, Ann Arbor, MI 48109-1042, USA

1 Introduction

Our linguistic competence can be accurately described as a productive and systematic representational system: we can understand an unbounded number of expressions, including novel ones, and there is a systematic pattern to that capacity. Specifically, we can understand most expressions formed out of familiar items arranged in accordance with the syntactic rules of our natural language. This systematic productivity allows us, among other things, to produce and understand meanings that we find wildly implausible. To many theorists, this suggests that our linguistic competence can be modeled using compositional truth-conditional theories.

At the same time, our linguistic competence can also be aptly described as a context-sensitive and highly flexible representational system. In particular, the meaning that expressions take on particular occasions often depends on the context, and there seem to be many different ways in which context is involved in these processes. Consider these simple examples familiar from the contextualist literature:

- (1) Mary cut the grass.
- (2) John cut the cake.

It is natural to understand Mary's cut in (1) as a horizontal cut, and John's cut in (2) as a vertical cut into pieces. However, it is easy to imagine contexts in which *cut the grass* means 'cut into pieces' (e.g., when the topic is about selling plots of grass), and *cut the cake* means 'cut horizontally' (e.g., in a fantastic case where cakes are growing out of control). This kind of context-sensitive flexibility suggests that our linguistic competence should not be modelled with the compositional theories that seem so well suited to account for its systematic productivity.

'Truth-conditional pragmatics', as I understand it here, is the project of trying to formally reconcile both aspects of our linguistic competence. The guiding assumption is that we can account, within a broadly compositional truth-conditional theory of meaning, for the systematic productivity of our linguistic competence without ignoring, or abstracting away from, its inherent semantic flexibility. Proponents accept that, to achieve that goal, semantic theories must be substantially revised; but they insist that we should not abandon the compositional truth-conditional project. Accordingly, this approach should be distinguished from contextualist views which also emphasise the semantic flexibility of language, but which abandon the search for anything like a formal theory of our linguistic competence.¹

The aim of this paper is to present and motivate a set of tools—specifically, a set of lexical representations and combinatorial operations—that can be used to account for

¹ 'Truth-conditional pragmatics', so understood, includes positions such as that of Recanati (2010), Szabó (2010), Rothschild and Segal (2009) and Lasersohn (2012), but not neo-Wittgenstenian positions such as that of Travis (1994), and some of the versions of the approach outlined in Bezuidenhout (2002). Truth-conditional pragmatics should also be distinguished from positions that, for one reason or another, deny that semantic theories should account for cases of semantic flexibility, such as the version of 'minimalism' defended in Fodor and Lepore (2002) and Cappelen and Lepore (2005).

semantic flexibility within a compositional truth-conditional framework.² I begin by examining, in Sect. 2, previous proposals in this tradition, and arguing that they face a basic problem: they try to account for semantic flexibility by radically ‘freeing up’ the compositional operations of language (Pagin and Pelletier 2007; Recanati 2010; Szabó 2010; Lasersohn 2012). The resulting theories are invariably too unconstrained: they predict flexibility in cases where it is clearly not observed. The reason why most accounts are forced to free the compositional operations—i.e., to increase their expressive power—is that they rarely, if ever, take advantage of the rich information encoded in lexical items. Building on Putnam’s early work on lexical competence, among others, I propose, in Sect. 3, a way of representing rich lexical items and a set of compositional operations that can access that information and introduce bits of it into the meanings of complex expressions. In Sect. 4, I show that this multidimensional semantics can deal with the main cases of semantic flexibility without over-generating unavailable meanings.

This strategy will surely raise some eyebrows. Many philosophers think that non-atomic accounts of the lexicon invariably result in either definitional or full-blown prototype theories. To be sure, psychologists and cognitive scientists have generally worked with rich representations of lexical meaning and concepts (for reviews, see Murphy 2002; Machery 2009). Furthermore, Putnam (1970, 1988) and Chomsky (2000, 2012) have argued that, to begin to model our semantic competence in any serious way, we need to make our theories sensitive to the rich and subtle information encoded in lexical items. The account of the lexicon I defend incorporates one of their basic insights: namely, that having full compositional competence with lexical items often requires that we grasp more information than that which determines their extension. At the same time, I will show that we can implement that insight in a truth-conditional framework. Specifically, I will show that we can introduce, into a compositional truth-conditional theory, informationally rich lexical items without falling into definitional/descriptivist or other objectionable views of the lexicon. The main task of this paper is to explain how compositional operations can interact with these kinds of lexical items, and illustrate the advantages this provides for our models of semantic flexibility.

2 Compositional truth-conditional pragmatics

To set the stage for this discussion, I begin by presenting and motivating the basic framework of truth-conditional pragmatics (Sect. 2.1). I then argue that, although the basic framework can successfully account for some key cases of semantic flexibility,

² In this paper, I will not directly engage with views, such as some versions of minimalism, that reject the desideratum that we should account for semantic flexibility in our compositional semantic theories. The current task is to explore the ways in which we should modify the basic truth-conditional framework, if we aim to model cases of semantic flexibility such as (1) and (2). Ultimately, I agree with Rothschild and Segal (2009) that, if we sever the connection between what is intuitively said by utterances and the data that our semantic theories should account for (as often results from accounts which reject the view that we should model semantic flexibility), it becomes extremely hard to see what type of data should constrain our semantic theories.

it vastly over-generates non-available readings in minimal variants of those key cases (Sect. 2.2). This mainly critical section will end by briefly suggesting why I think that, to appropriately constrain truth-conditional pragmatics, we need to develop a more complex and realistic model of lexical semantics than extant accounts have so far worked with. The rest of the paper presents, motivates and defends one such proposal (Sects. 3–5).

2.1 Motivating the basic framework

A main task of truth-conditional theories is to assign semantic values to lexical items and specify rules for combining them to determine the semantic values of complex expressions. For example, we define an interpretation function ‘ $\llbracket \cdot \rrbracket$ ’ over lexical entries such as:

$$(3) \quad \llbracket \textit{John} \rrbracket = \text{JOHN}$$

$$(4) \quad \llbracket \textit{green} \rrbracket = \lambda x. \text{GREEN}(x)$$

$$(5) \quad \llbracket \textit{cut} \rrbracket = \lambda x. \lambda y. \text{CUT}(y, x)$$

And in terms of compositional rules that determine the meaning of phrases from their structure and the meanings of their parts (assume for simplicity that all non-branching nodes are terminal nodes):

(TN) If α is a terminal node, then $\llbracket \alpha \rrbracket$ is specified in the lexicon.

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

As formulated, there is no place where the effects of context can enter truth-conditional interpretation. To allow for the most basic kind of context sensitivity—exhibited by indexicals and demonstratives such as *I* and *that*—we can add a context parameter to the interpretation function, and assume that the meanings of some expressions are characters. For any expression e :

$$(6) \quad \llbracket e \rrbracket^c = f_e(c)$$

where f_e is the character of e and $f_e(c)$ is the occasion meaning of e in context c . For example, we can revise the entry for *green* in (4), as follows:

$$(7) \quad \llbracket \textit{green} \rrbracket^c = \lambda x. \text{GREEN}(C)(x)$$

‘ C ’ is a comparison class provided by the context. If an expression e has no free or context sensitive parameters, we can say that for all c ’s, $f_e(c) = m$, where m is the standing meaning of e . Despite some technical challenges (e.g., what precisely are contexts? In entries like (7), should we lambda abstract over C before ‘ λx ’ to force covert saturation early in composition?), most would agree that there is no deep tension between truth-conditional semantics and character-based context sensitivity (see, e.g., Stanley 2007; Rothschild and Segal 2009; Recanati 2010).

However, a basic tenet of truth-conditional pragmatics is that characters are not the only source of context sensitivity. For example, in cases like (1)–(2) it does not

seem that the flexibility of *cut* can be traced to an empty parameter, such that the term cannot be part of a truth-evaluable content until that is specified. Since such cases are ubiquitous, most truth-conditional pragmatists hold that there is an additional source of context sensitivity: free modulation. Following [Pagin and Pelletier \(2007\)](#) and [Recanati \(2010\)](#), we can represent the *modulated* meaning of an expression *e*, $\llbracket e \rrbracket_M^c$, as follows:

$$(8) \quad \llbracket e \rrbracket_M^c = mod(e, c)(\llbracket e \rrbracket^c)$$

mod takes as an argument an expression *e* and context *c* in which *e* occurs and returns as value a modulation function, call it ‘*f_{M,e}*’, which takes $\llbracket e \rrbracket^c$ and returns the meaning that is salient, relevant or appropriate for *e* in *c*. To illustrate, consider (9): in a normal cake cutting context, *mod*(*cut*, *c*) returns the modulation function specified in (9-b), which can take the standing meaning of *cut*, as specified in (5) above, and return the enriched ‘cut into pieces’ meaning in (9-c):

$$(9) \quad \begin{aligned} \text{a. } & \llbracket cut \rrbracket_M^c = mod(cut, c)(\llbracket cut \rrbracket^c) \\ \text{b. } & = [\lambda P_{\langle e, \langle e, t \rangle \rangle} . \lambda x . \lambda y . P(x)(y) \wedge IN \text{ PIECES}(x)](\llbracket cut \rrbracket^c) \\ \text{c. } & = \lambda x . \lambda y . CUT(x, y) \wedge IN \text{ PIECES}(x) \end{aligned}$$

What is crucial to note about the overall effect of *mod* is that it has enough expressive power to add descriptive content to the meaning of expressions. In what follows, I call processes with the expressive power illustrated in (9), and which operate in addition to the saturation of explicit context sensitive parameters, ‘free modulation’.³

Critics have objected that flexible interpretation procedures such as (8) (i.e., free modulation) are in tension with the compositionality of language ([Fodor 2001, 2003](#)). In response, truth-conditional pragmatists such as [Pagin and Pelletier \(2007\)](#) and [Recanati \(2010\)](#) argue that we can allow for widespread modulation by simply assuming that the compositional combinatorial rules take the modulated (instead of the occasion) meanings of the immediate constituents of complex expressions. To incorporate this suggestion, we can formulate $\llbracket \cdot \rrbracket_M^c$ as follows:

- (*TN_M*) If α is a terminal node, then $\llbracket \alpha \rrbracket_M^c = mod(\alpha, c)(\llbracket \alpha \rrbracket^c)$, where $\llbracket \alpha \rrbracket^c$ is the character of *e*, as specified in the lexicon, applied to *c*.
- (*FA_M*) 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 $\llbracket \alpha \rrbracket_M^c = mod(\alpha, c)(\llbracket \beta \rrbracket_M^{c_1}(\llbracket \gamma \rrbracket_M^{c_2}))$ ⁴

According to this formulation of the interpretation function, modulation can take as input the occasion meaning of lexical items (see *TN_M*), and also the output of any

³ That this descriptive content has to be somehow added from ‘outside’ the modified expression is undeniable if we also assume lexical semantic atomism. To be clear, some truth-conditional pragmatists think that these operations are somehow constrained, and we will discuss specific proposals below. However, their actual definition of modulation tends to have the expressive power of free modulation (e.g., [Recanati 2010](#)). Ultimately, my multidimensional semantics presents one way in which modulation could be constrained, which I hope some truth-conditional pragmatists will find convincing, and broadly complementary to similar efforts such as [Pagin \(2014\)](#).

⁴ *c*₁, *c*₂ stand for the subparts of context *c* in which the respective constituents occur.

combinatorial operation (see FA_M). In addition, each instance of modulation has as much expressive power as allowed by the *mod* function as defined in (8).⁵

This proposal has some *prima facie* virtues. It can be used to account for examples such as (1) and (2), in which the manner of cutting is determined not only by the argument of *cut* but also by the wider context, as illustrated in (9-a)–(9-c). In addition, it can also deal with other famous examples of semantic flexibility that have been used to attack compositional truth-conditional theories, in particular colour and privative adjectives (Lakoff and Johnson 1980; Lahav 1989; Coulson and Fauconnier 1999; Travis 1994).

Consider first a simple version of Travis' famous 'green leaves' cases. The basic observation is that in a context like (10), the assertion of (10-a) is true, whereas in a context like (11), the assertion of (11-a) is false. Assuming the relevant parts of (10-a) and (11-a) have the same meaning, this raises a challenge to the compositionality of 'what is said'. However, if free modulation can enter the interpretation process (i.e., if it can affect what is said), we can hold that, in each context, the effect of *mod* is different, so that we get reading (10-b) in context (10) (the subscript 'qual' marks a 'qualitative' sense of green), and (11-b) in context (11) (the subscript 'non-grad' marks a sense in which what matters is the way in which the leaves came to be green).

- (10) Pia paints the leaves green to prepare them to paint a still life.
- a. The leaves are green
 - b. The leaves are green_{qual}
- (11) Pia responds to a Botanist friend seeking to study green leaf chemistry.
- a. The leaves are green
 - b. The leaves are green_{nongrad}

As a second example, consider privative adjectives. How can we give an account of cases like (12) without making cases like (13) trivially true?

- (12) A fake gun is not a gun.
- (13) That giraffe is a fake gun.

Partee (2007, 2010) defends an account of privatives such as *fake gun* that can be interpreted as using something like free modulation, and has for this reason been endorsed by theorists such as Recanati (2010, ch. 2). Partee argues that, in the relevant cases, *gun* involves a shift (to a 'looks like' reading) similar to that observed in the case of constitutive material modifiers:

- (14) ...stone lion ...
- a. ...stone lion₁ ...
 - b. ...stone lion₂ ...

⁵ Theorists that allow essentially this way of reconciling compositionality with semantic flexibility include Recanati (2010), Lasnik (2012), Szabó (2010) and Pagan and Pelletier (2007). There are, of course, some differences between these accounts. When discussing details additional to the basic framework presented so far, the account we follow most closely here—as a standard example of truth-conditional pragmatics—is that of Recanati (2010).

- (15) ...fake gun ...
 - a. ...fake gun₁ ...
 - b. ...fake gun₂ ...

In (14)–(15), the subscript ‘2’ marks a shift from the literal meaning of these terms to a loose (‘looks like’) reading. In these cases, the effect of free modulation is to turn the normally empty predicates in (14-a) and (15-a) into the more informative predicates in (14-b) and (15-b), which in most contexts have a positive and negative extension (there are stone things which do/do not look like lions, and fake things which do/do not look like guns).

In short, incorporating free modulation into our accounts of interpretation by adopting something like $\llbracket \cdot \rrbracket_M^c$ seems like a promising way to deal with the sorts of examples of semantic flexibility that have been raised against compositional truth-conditional models of linguistic competence.

2.2 Over-generation problems

Despite its initial plausibility, the enriched interpretation function, $\llbracket \cdot \rrbracket_M^c$, raises an immediate worry, namely, that it over-generates non-available meanings. Consider the following examples (see Asher 2011):

- (16) a. John’s sister was hit in the fender by a truck and will cost a lot to repair.
- b. John’s sister’s car was hit in the fender by a truck and will cost a lot to repair.

Note that (16-a) cannot be read as (16-b), which is puzzling given the assumption that our compositional competence includes free modulation. Why can’t a modulation function take $\llbracket \text{John’s sister} \rrbracket^c$ and shift it to a meaning of the same semantic type that is salient and relevant in the context, namely, ‘John’s sister’s car’?

At first glance, there seems to be a straightforward response. The problem is just that free modulation, as currently implemented in $\llbracket \cdot \rrbracket_M^c$, can apply at any level of syntactic structure. However, most of the examples of semantic flexibility discussed in Sect. 2.1 are traceable to the modulation of lexical items. So a natural response to over-generation objections based on cases like (16-a)–(16-b) is to formulate a constrained version of $\llbracket \cdot \rrbracket_M^c$ in which free modulation applies only to lexical items. Consider the following proposal:

- (TN_M) If α is a terminal node, then $\llbracket \alpha \rrbracket_M^c = mod(\alpha, c)(\llbracket \alpha \rrbracket^c)$, where $\llbracket \alpha \rrbracket^c$ is the character of e , as specified in the lexicon, applied to c .
- (FA_M) 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 $\llbracket \alpha \rrbracket_M^c = \llbracket \beta \rrbracket_M^{c_1}(\llbracket \gamma \rrbracket_M^{c_2})$

In this implementation, *mod* does not operate on the outputs of FA_M , but only on terminal nodes/lexical items. This captures a more constrained form of meaning modulation.

Since the compositional step amounts to FA_M , we can say that, on this implementation, free modulation is pre-compositional.⁶

However, even when free modulation is confined to non-branching nodes, $\llbracket \rrbracket_M^c$ is still too unconstrained. The main problem, we will see, is not related to its structural location. Rather, free modulation itself has too much expressive power, and still over-generates non-available readings. Consider the following example, due to Asher (2011):

- (17) a. Mary stopped the apple.
b. Mary stopped eating the apple.

In most contexts, the meaning expressed by (17-b) would be the most relevant and salient interpretation of (17-a). However, that reading is unavailable. If a relevance seeking free modulation function could apply to any lexical item, these reading would, it seems, be easily obtained. Indeed, even adding a priming context does not result in the desired modulations. Consider:

- (18) John was busy, but is now ready to go for lunch.
a. He finally stopped the garden.
b. He finally stopped mowing/fixing the garden.

In this case, *stopped* in (18-a) cannot mean ‘stopped mowing/fixing’, as in (18-b), even though that would be the easiest and most relevant modulation in the context (indeed, it might even be more informative than ‘finished’, if the speaker does not want to suggest that the task was completed). Here’s another example:

- (19) John and Mary want to hang some paintings.
a. John began the nails.
b. John began hammering the nails.

In this case, *began* in (19-a) cannot mean ‘began hammering’, as in (19-b), although this would result in a relevant modulation and could be achieved by a simple, contextually salient enrichment of *began*. Such examples (which can be easily multiplied) suggest that the best approach is not to confine the syntactic location of *mod* to the leafs (terminal nodes), but rather to somehow constrain its operations or expressive power.⁷

Indeed, variations of the examples involving color terms and privative adjectives—cases often taken to directly support the view that interpretation involves something

⁶ As before, we assume for brevity that functional application is the only combinatorial rule. I should also add that, as far as I know, most truth-conditional pragmatists don’t address the issue of whether modulation should be confined to particular levels of syntactic structure. For related discussions, see Lewis (2014) and Del Pinal (2015b, 2016).

⁷ Stanley (2002) and others have raised over-generation objections against truth-conditional pragmatic accounts which adopt free modulation. Readers familiar with that debate will note that the kinds of examples that I mentioned above, which focus on content words and open class expressions, are different from those commonly discussed in the philosophical literature, which tend to focus on logical and functional terms. Here I can remain neutral with respect to whether non-available cases of quantifier domain restriction present a serious challenge to free *mod*. For recent discussion and responses to Stanley’s original argument, see Hall (2014) and Pupa (2015).

like free *mod*—also suggest that this formulation of $\llbracket \cdot \rrbracket_M^c$ over-generates readings. Take the following scenario, due to Kennedy and McNally (2010): consider an assertion of (20-a) in a context in which leaf A and leaf B are objectively the same rich shade of green, but leaf A has been painted green, while B is naturally green. The robust intuition here is that (20-a) is false.

- (20) a. Leaf A is green but leaf B isn't.
b. Leaf A is painted green but leaf B isn't.

Although *green* is often associated with *painted green*, there is no reading of (20-a) that has the meaning of (20-b), despite the fact that this reading would make (20-a) true and informative in the assumed context (e.g., to say which leaf is painted green).

The free modulation account of privative adjectives, inspired by Partee (2007, 2010), faces a similar problem. Consider (21-a). What is striking about this example is that it does not have a reading roughly like that in (21-b). That is, *fake guns* in (21-a) is not modulated to something like *fake toy guns*. Since in many cases we often associate *guns* with *toy guns*, this would seem to be a rather easy and salient free modulation.

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

Again, the claim in (21-a) would make perfect sense if *gun* could be modulated so as to include all sorts of things that look loosely like guns, including as model and toy guns. In the supposed context, some artifacts that look like toy guns are in fact real guns, and this would be the items denoted by *fake guns* in (21-a), to get a reading like (21-b). However, that reading is not available, even though it would allow us to understand the assertion in (21-a) in a contextually relevant and coherent way (namely, as saying that the terrorists made toy gun look-a-likes that were in fact real guns).

It seems clear, then, that the view that the interpretation function involves free modulation massively over-generates.⁸ A tempting response, at this point, is to invoke Chomsky's distinction between competence and performance (1986, 2000). The basic idea is well-known: our syntactic competence has the capacity to generate sentences of unbounded complexity, but this is constrained by performance factors such as limits on working memory. Why not appeal to a similar argument on behalf of truth-conditional pragmatics? Our linguistic competence involves an unconstrained *mod*, which can freely modify the meanings of expressions to increase their relevance. Still, which meanings we can actually generate given the input depends on factors such as how working memory interacts with the speed and complexity of incoming input. We can only make predictions about the available meanings of particular expressions in

⁸ Recanati (2010) suggests that one possible constrain on free *mod* is to assume that it cannot perform type shifting of its argument. This move might help with examples like (17)–(19), but has no effect on cases like (20-a)–(21-a), where the unobserved modulation amounts to a simple intersective operation. For further discussion, see Del Pinal (2015a, 2016)

conjunction with such performance factors. In some cases, the available meanings might not be the most relevant or salient, relative to those suggested by slow and explicit deliberation.

Does this appeal to performance help explain why examples such as (18-a)–(21-a) do not shift to the intuitively most relevant meanings? This seems unlikely. Those cases begin with an explicit context that tries to prime the unavailable readings. The unavailable readings can be easily recovered from the immediate linguistic context, and would require just one enrichment operation. The only other available meaning is the literal meaning, hence there do not seem to be many intervening/interfering options. To be sure, the prospects of this proposal depend on which specific properties of working memory and online processing are invoked. Still, currently known limits—e.g., holding too many non-unitized items in mind, or performing too many simultaneous operations on them—do not seem to be violated by the unavailable readings in (18-a)–(21-a). Until such details are provided, it seems reasonable to hold that performance considerations, although interesting and important, are not sufficient to explain the unavailability of the explicitly primed meanings in cases such as (18-a)–(21-a).

The most plausible suggestion, at this point, is that there are direct, inherent limits on the modulation operations available to the interpretation function.

2.3 Diagnosis

Why is it so tempting to hold that the only or best way to deal with semantic flexibility is to radically free the compositional operations of language? Many philosophers and semanticists work under the assumption that lexical items provide minimal information to compositional linguistic processes. Dominant theories treat most lexical items as either semantically atomic, or as having a very simple internal structure. Some think of that assumption as a useful simplification; others, I suspect, simply don't see how we can add semantic complexity into lexical items without falling into problematic 'definitional' theories. Also influential is Fodor's argument that lexical atomism is the only theory that is compositional in the way required to account for the systematic productivity of language/thought (Fodor 1998; Fodor and Lepore 2002). Independently of the justification for such austere views of the lexicon, the result is that we are left with a considerable gap between the kind of information that can be computed from the lexicon and the intuitive meaning of utterances (what is said).

Given those assumptions and results, a natural move is to hold that the interpretation function is intertwined with free modulation operations that have almost unbounded expressive power. However, I have argued that this move seriously over-generates non-available interpretations. Still, I think that truth-conditional pragmatists can overcome this problem; but to tackle semantic flexibility within a compositional framework, we must adopt richer models of the information encoded in the lexicon than have hitherto been explored, while avoiding the pitfalls of definitional theories of meaning. In addition, we need to adjust the compositional operations so that they have access to that information. In the rest of this paper, I present a proposal for how to do that (Sect. 3), and show that it can deal with semantic flexibility in a relatively constrained way (Sect. 4). I stress that this proposal falls within the project of truth-conditional

pragmatics, and is intended to complement other proposals in this tradition for how to constrain the modulation operations which are part of compositional semantics (e.g., [Pagin 2014](#); [Recanati 2004, 2010](#)).

3 Multidimensional semantics: the basic framework

If we take seriously the idea that lexical items have a complex or multidimensional semantic structure to which compositional processes are sensitive, two key questions arise. First, how should we represent the information encoded in semantically complex lexical items? Second, how should we revise the interpretation function so that the compositional operations can appropriately interact with those kinds of representations? We tackle the first question in Sect. 3.1, and the second one in Sect. 3.2.

3.1 Lexical items

The proposal I'll defend can be thought of as a multidimensional 'Putnam-style' semantics. Here's the basic idea. Most lexical entries have an 'E-structure' and 'C-structure'. The E-structure directly represents or determines their extension, and the C-structure consists of what [Putnam \(1970\)](#) describes as a set of beliefs about the extension. Putnam sometimes calls this second component a 'stereotype/prototype'; but at least for terms such as natural kinds, what he means is something like a simple encoding of a folk theory about the extension. To capture this suggestion, I assume that the C-structure of lexical items encodes what, according to our best psychological theories, are the basic components of the corresponding kinds of concepts. There are relevant open debates, some of which we discuss below. Still, we can adopt an account of C-structure that is general enough to accommodate various proposals, and yet sufficiently concrete to allow us to see how these kinds of complex structures can be integrated into compositional theories, and what advantages this might bring to accounts of semantic flexibility.

The basic features of our Putnam-style lexicon are the following. (i) The C-structures associated with most classes of concepts are multidimensional. In the case of common nouns, these include dimensions like perceptual (e.g., what typical members look like), geneological (the usual way in which members come into being) and functional (e.g., the function, if any, of the members). (ii) C-structures encode the weight of each dimension as a function of its importance/usefulness for categorization and projection/induction. (iii) C-structures encode at least some basic relations between these dimensions, including their relative centrality.⁹ (iv) Different *kinds* of concepts can be represented by differences in the dimensions encoded and in the weights and

⁹ To say that a dimension/feature *d* is central in concept *C* is to say that other dimensions/features depend on *d* more than *d* depends on them. The idea that conceptual structures encode dependency relations (hence relative centrality) was originally defended by [Keil \(1989\)](#), who was directly influenced by the criticisms by [Putnam \(1970, 1975\)](#) and [Kripke \(1980\)](#) of cluster/definitional theories. Currently, even proponents of prototype theory accept that, for many domains, concepts encode dependency relations (e.g., [Hampton 2006](#)).

relations between dimensions.¹⁰ (v) Satisfying the dimensions in its C-structure is not necessary for entities to fall under a concept, and our linguistic competence reveals an implicit understanding of this. Hence we should not, in general, represent C-structures as if they provided (conjunctive or even disjunctive) definitions, i.e., as if they determined E-structure. This last point has been emphasized by many philosophers, at least since Putnam (1970), but tends to be ignored by scientists who defend non-atomic lexical theories.¹¹

To illustrate some of these points, consider the simplified entries for *lion* in (22) and for *gun* in (23) below. The E-structure of each entry stands for the core/atomic component which determines its extension. We take the relation between the E-structure and the extension of lexical terms as primitive (i.e., as given by the relevant model). The C-structure of each entry can be thought of as a restricted set of general beliefs associated with (or about) the extension. In the case of common nouns, the dimensions of C-structure include, at least, information about how entities in the class tend to look, feel or taste ('P' for 'PERCEPTUAL'), what materials or parts they are made of ('C' for 'CONSTITUTIVE'), how they came to being or for what purpose they were created ('A' for 'AGENTIVE'), and what their intended and typical function is, if any ('T' for 'TELIC') (cf. Pustejovsky 1995; Moravcsik 1998; Del Pinal 2015a). Different classes of terms may differ in the dimensions they represent. For example, it is arguable that many natural kinds do not have a default value for the TELIC, although this is an important dimension for artifact kinds, a contrast that is captured in (22) and (23). Note, also, that the C-structure of a term is informationally richer than its 'associated prototype' (as this notion is used in philosophy) and is closer to encoding a 'conception', in the sense of Burge (1993).¹²

$$(22) \quad \llbracket \textit{lion} \rrbracket_M^c =$$

E-structure: $\lambda x. \text{LION}(x)$
C-structure:
 C: $\lambda x. \text{SUBSTANCE_LION}(x)$
 P: $\lambda x. \text{PERCEPTUAL_LION}(x)$
 T:
 A: $\lambda x. \exists e_1 [\text{BIOLOGICAL_BIRTH_LION}(e_1, x)]$

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

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

¹⁰ For example, a common view in cognitive science is that artifact concepts encode a dimension which represents their basic function, although there are ongoing debates about its relative centrality (see, e.g., Margolis and Laurence 2007). At the same time, everyone agrees that most perceptual concepts do not encode a function dimension.

¹¹ For a detailed discussion of why C-structure does not, in general, determine E-structure (including a discussion of 'the problem of ignorance' and 'the problem of error'), and why both components should nevertheless be kept as part of our lexical competence, see Del Pinal (2015a), § 7.2–7.4.

¹² When philosophers talk about 'prototypes' they often mean things like 'the representation of the perceptual average of a class'. However, Putnam's 'stereotypes/prototypes' are closer to what are now called 'theory-based' prototypes (Hampton 2006), structures which encode, at least, the dependency relations between dimensions/features.

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

These basic Putnam-style lexical entries can be further refined.¹³ In addition, we will see later how this framework can be extended to other classes of expressions. Still, we can already see why assuming that the interpretation function has access to Putnam-style entries could help account for semantic flexibility. In typical cases, ‘enrichments’ look like operations which take dimensions of the C-structure and upload them into and conjoin them with the value of the E-structure (e.g. when the at-issue meaning of *lion* is ‘typical lion’). ‘Loosenings’ look like operations which take a dimension of the C-structure and use it to replace the value of the E-structure (e.g. when the at-issue meaning of *lion* is ‘lion look alike’). Furthermore, these operations could be triggered by neighboring linguistic material (e.g., to accommodate the semantic types of sister nodes so that functional application can proceed), but they could also be triggered by more general aspects of the context such as intonation (e.g., as in *Wow! That’s a lion!*). At the same time, we don’t want to say that, in general, dimensions of the C-structure are part of E-structure, i.e., of core meaning: in the default case, *Pete is a lion* can be true even if Pete does not, say, look or behave like a lion.

3.2 Compositional interpretation

The main challenge for integrating a Putnam-style semantics with a truth-conditional theory is to determine how, precisely, the compositional operations should interact with semantically complex lexical entries such as (22) and (23). Clearly, we have to modify the interpretation function used in standard type-driven semantic theories so that it can properly interact with lexical entries which have a dual semantic structure. The first obvious modification is that the interpretation function will involve two computations. One computation, which we designate ‘ $\llbracket \cdot \rrbracket_{ME}^c$ ’, determines the E-structure of expressions. The other computation, which we designate ‘ $\llbracket \cdot \rrbracket_{MC}^c$ ’, determines their C-structure. In addition, we will assume that the original ‘ $\llbracket \cdot \rrbracket_M^c$ ’ retrieves the ‘full meaning’ of expressions as a tuple of their E-structure and C-structure:

$$- \llbracket \cdot \rrbracket_M^c = \left\langle \llbracket \cdot \rrbracket_{ME}^c, \llbracket \cdot \rrbracket_{MC}^c \right\rangle$$

Thus far, it would seem that E-structure and C-structure are computed in parallel, i.e., without interacting. However, the point of introducing C-structure into our model of semantic competence is to allow it to play a (constrained) role in the determination of

¹³ In particular, these entries do not yet encode information about the weights of, and dependency relations between, dimensions. In addition, the values provided for each dimension are for illustration only. To mention one case, the value of the PERCEPTUAL in (22) is likely something closer to the following predicate: $\lambda x. \text{SIM}(\text{LION}_p)(x) > \text{STND}_{\text{sim},c}(\text{LION}_p)$, where an entity e satisfies this predicate if its similarity to the perceptual prototype of ‘lion’, i.e., ‘ $\text{SIM}(\text{LION}_p)(e)$ ’, passes the similarity standard relevant in c , i.e., ‘ $\text{STND}_{\text{sim},c}(\text{LION}_p)$ ’.

aspects of E-structure. How do we let the compositional operations which determine the E-structure of complex terms have access, under certain conditions, to the C-structure of the constituents? To achieve this, we introduce the following tools:

- T1. ‘Dimension operators’: partial functions from the full meaning of terms into their specific C-structure denotations. For example, common nouns are associated with operators such as Q_A , which returns the value of the AGENTIVE, and Q_T , which returns the value of the TELIC:

$$(24) \quad Q_A(\llbracket lion \rrbracket_M^c) = \lambda x. \exists e_1[\text{BIOLOGICAL_BIRTH_LION}(e_1, x)]$$

$$(25) \quad Q_T(\llbracket gun \rrbracket_M^c) = \lambda x. \text{GEN } e[\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)]$$

In addition, we include an operator E which takes the full meaning of expressions and returns the value of its E-structure.

- T2. ‘Core enrichment operators’: partial functions from the full meaning of terms into combinations of their E-structure and C-structure. For example, common Ns are associated with operator A , which returns a conjunction of the E-structure and the value of the AGENTIVE:

$$(26) \quad A(\llbracket lion \rrbracket_M^c) = \lambda x. E(\llbracket lion \rrbracket_M^c)(x) \wedge Q_A(\llbracket lion \rrbracket_M^c)(x) \\ = \lambda x. \text{LION}(x) \wedge \exists e_1[\text{BIOLOGICAL_BIRTH_LION}(e_1, x)]$$

In addition, we include a function ALL which is an identity operator on full meanings.¹⁴

The basic function of dimension and core-enrichment operators is simple. The former take full meanings and return the value of a particular dimension. The latter take full meanings and return combinations of the E-structure and one or more of the dimensions of C-structure.¹⁵

The key move in our compositional implementation of a Putnam-style semantics is to introduce dimension and core enrichment operators into the specification of $\llbracket \cdot \rrbracket_{M_E}^c$, i.e., into the determination of E-structure. To do this, we can piggy back on the structural insights of the original idea of modulated interpretation (cf. [Pain and Pelletier 2007](#); [Recanati 2010](#)), as in (27):

$$(27) \quad \llbracket e \rrbracket_{M_E}^c = \text{mod}_I(e, c)(\llbracket e \rrbracket_M^c)$$

¹⁴ There are accounts in the literature which use tools similar to dimension and core enrichment operators, e.g., [Vikner and Jensen \(2002\)](#)’s account of genitives and [Pustejovsky \(1995, 2012\)](#)’s theory of the Generative Lexicon. However, these accounts do not postulate a division between E-structure and C-structure, hence are not strictly ‘Putnam-style’. In addition, the way in which these operators are integrated with the interpretation function is quite different from the technique I specify below. The account presented in [Del Pinal \(2015a\)](#) also uses similar operators, and does postulate an E/C-structure, but the compositional implementation is different and less flexible than the account presented here.

¹⁵ Terminological note. I use ‘ Q_i ’ as the name of dimension operators, where the subscript stands for the name of the dimension picked out: e.g., Q_P returns the PERCEPTUAL and Q_T returns the TELIC. For core enrichment operators, I just use the abbreviated name of the dimensions: e.g., P returns the E-structure conjoined with the PERCEPTUAL. We could define core enrichment operators in terms of combinations of dimension operators, so we need not assume that these are two distinct classes of operators. However, for reference in the informal descriptions—in particular, to see whether the effect is an intuitive ‘enrichment’ or ‘loosening’, relative to the E-structure—it is useful to keep using these two names.

Note that (27) uses mod_l instead of free mod , the original pragmatist function introduced in (8) to define the notion of modulated interpretation. mod_l also takes as an argument an expression e and context c in which e occurs. However, it returns as value a lexical modulation function l_i drawn from the set of dimension and core enrichment operators available for e . Call this set ' $\mathbf{M}_{(e)}$ ', where the term in parentheses is used to identify the relevant expression.¹⁶ In the case of common nouns, the l_i 's include the examples used in T1 and T2, and more generally, dimension operators into any dimension of the C-structure, and core enrichment operators that combine the E-structure with any dimension/s of the C-structure. However, as we will see, not all expressions have a C-structure. In such cases, the only available function will be E , which selects the original E-structure. This holds also for some classes of simple expressions, such as certain modifiers of predicates (e.g., privative adjectives).

As a first approximation, we can say that the job of mod_l is to select an l_i that takes the full meaning of e and returns an E-structure that is relevant for e in c . However, we can be more specific. As a default, the selected function will be E , i.e., the operator that returns the E-structure of full meanings. If there is a type clash between sister nodes such that functional application (or whatever combinatorial operation we postulate) cannot proceed, then mod_l will act as a type shifter by selecting the operator that returns a dimension's value that eliminates the type clash. This is only possible if there is such value available in the C-structure. Finally, when there is support from context (intonation, task demands, etc), mod_l can select an l_i which enriches or loosens the original E-structured of the term to which it applies.¹⁷

We are now in a position to reformulate the interpretation function $\llbracket \cdot \rrbracket_M^c$ so that it can interact with Putnam-style lexical items. For simplicity, we continue to focus just on TN_M and FA_M , defined for our multidimensional semantics below. As a first pass, note that TN_M simply calls the full meaning tuples, i.e., the E/C-structures, as specified in the lexicon. FA_M determines the E-structure of a complex expression by functional application after each constituent has been filtered by mod_l , and it determines its C-structure via f , which takes the full meaning of each constituent and applies functional application along each dimension of their C-structure.

(TN_M) If α is a terminal node, then $\llbracket \alpha \rrbracket_{ME}^c$ is the E-structure of α as specified in the lexicon, relative to c , $\llbracket \alpha \rrbracket_{MC}^c$ is the C-structure of α as specified in the lexicon, relative to c , and $\llbracket \alpha \rrbracket_M^c = \left\langle \llbracket \alpha \rrbracket_{ME}^c, \llbracket \alpha \rrbracket_{MC}^c \right\rangle$.

¹⁶ If α is a complex expression, the set $\mathbf{M}_{(\alpha)}$ is usually determined by the head of α . For example, the set of dimension and core enrichment operators available to *red box* is the same as that available to *box*.

¹⁷ This last case is the one closest to the original pragmatist notion of 'free' modulation. Indeed, a full account of this process could be obtained by combining our multidimensional semantics with previous proposals for how to predict both when enrichment occurs and what its function is, such as Pagin (2014)'s account of free enrichment as coherence raising. The additional constraint, when implemented in our framework, is that mod_l can, say, raise coherence but only by selecting an appropriate l_i . In addition, we could also hold that these 'free' processes are at least partly determined by certain conventions. For example, intonation often determines which operator is selected: for certain classes of expressions (e.g., social role nouns such as *man* and *woman*), we learn that certain intonations reliably call for particular dimension/core enrichment operators.

(FA_M) If α is a branching node, $\{\beta, \gamma\}$ is the set of α 's daughters in c , and $mod_l(\beta, c_1)(\llbracket \beta \rrbracket_M^{c_1})$ is a function whose domain contains $mod_l(\gamma, c_2)(\llbracket \gamma \rrbracket_M^{c_2})$, then¹⁸:

- (i) $\llbracket \alpha \rrbracket_{ME}^c = mod_l(\beta, c_1)(\llbracket \beta \rrbracket_M^{c_1})(mod_l(\gamma, c_2)(\llbracket \gamma \rrbracket_M^{c_2}))$
- (ii) $\llbracket \alpha \rrbracket_{MC}^c = f(\llbracket \beta \rrbracket_M^{c_1}, \llbracket \gamma \rrbracket_M^{c_2})$
- (iii) $\llbracket \alpha \rrbracket_M^c = \left\langle \llbracket \alpha \rrbracket_{ME}^c, \llbracket \alpha \rrbracket_{MC}^c \right\rangle$

To illustrate TN_M and FA_M , we briefly consider two examples. These examples are substantially simplified; their point is only to highlight the basic properties of our new $\llbracket \cdot \rrbracket_M^c$. We will then be in a position to tackle, in Sect. 4, more realistic examples, including our original target cases of semantic flexibility.

Consider first $\llbracket steel\ gun \rrbracket_M^c$, and assume that the full meaning of *steel*, when taken as a nominal modifier, is as in (28) below, and that of *gun* is as in (23) above. By TN_M , $\llbracket steel \rrbracket_M^{c_1} = \left\langle \llbracket steel \rrbracket_{ME}^{c_1}, \llbracket steel \rrbracket_{MC}^{c_1} \right\rangle$ and $\llbracket gun \rrbracket_M^{c_2} = \left\langle \llbracket gun \rrbracket_{ME}^{c_2}, \llbracket gun \rrbracket_{MC}^{c_2} \right\rangle$. The next step is to apply FN_M . Assume that mod_l operates in default mode, hence selects the dimension operator E in each sub-context, i.e., $E(\llbracket steel \rrbracket_M^{c_1})$ and $E(\llbracket gun \rrbracket_M^{c_2})$. $\llbracket steel\ gun \rrbracket_{ME}^c$ will then be as specified in the E-structure of (29), and $\llbracket steel\ gun \rrbracket_{MC}^c$, which is obtained by applying functional application along each dimension, will be as specified in the C-structure of (29). As illustrated in (29), the resulting E-structure of *steel gun* includes less information than its C-structure.

(28) $\llbracket steel \rrbracket_M^c =$
E-structure: $\lambda P.\lambda x. P(x) \wedge STEEL(x)$
C-structure:
 C: $\lambda P.\lambda x. P(x) \wedge STEEL(x)$
 P: $\lambda P.\lambda x. P(x) \wedge STEEL_PERCEPTUAL(x)$
 T: $\lambda P.P$
 A: $\lambda P.P$

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

Terms such as *steel* simply add descriptive content of their own, along one or more dimensions. More interesting for illustrating the flexibility of $\llbracket \cdot \rrbracket_M^c$ are modifiers which plausibly use the C-structure of their arguments to generate a modified predicate.

¹⁸ As in the original formulation, c_1, c_2 stand for the subparts of context c in which the respective constituents occur.

Consider *typical gun*. Suppose the entry for *typical* is (30), where the variable ‘ D_C ’ ranges over full meaning tuples, and dimension operators are used to specify the effect of this modifier along each dimension. To apply FA_M to generate the meaning of *typical gun*, mod_l now has to return a full meaning tuple for *gun*, which can be achieved by selecting the *ALL* core enrichment operator, i.e., $ALL(\llbracket gun \rrbracket_M^c)$. We can then apply functional application along each dimension, resulting in (31). Note that the E-structure of *typical gun* requires that an entity have the perceptual features of guns, although this requirement is not part of the E-structure of the unmodified *gun*, as can be seen by comparing (23) and (31). So in this case application of FA_M has the result that parts of the C-structure of a constituent are uploaded into the E-structure of the complex expression.

- (30) $\llbracket typical \rrbracket_M^c =$
E-structure: $\lambda D_C \lambda x. Q_E(D_C)(x) \wedge Q_P(D_C)(x)$
C-structure:
 C: $\lambda D_C. Q_C(D_C)$
 P: $\lambda D_C. Q_P(D_C)$
 T: $\lambda D_C. Q_T(D_C)$
 A: $\lambda D_C. Q_A(D_C)$
- (31) $\llbracket typical\ gun \rrbracket_M^c =$
E-structure: $\lambda x. GUN(x) \wedge PERCEPTUAL_GUN(x)$
C-structure:
 C: $\lambda x. PARTS_GUN(x)$
 P: $\lambda x. PERCEPTUAL_GUN(x)$
 T: $\lambda x. GEN\ e[SHOOTING(e) \wedge INSTRUMENT(e, x)]$
 A: $\lambda x. \exists e_1[MAKING(e_1) \wedge GOAL(e_1, GEN\ e[SHOOTING(e) \wedge INSTRUMENT(e, x)])]$

At this point, I can highlight five key features of $\llbracket \rrbracket_M^c$. First, as should be clear from the formulation of FA_M , mod_l is confined to E-structure. In the default case, illustrated by *steel gun* in (29), it selects E and so returns the E-structure of each constituent. But as illustrated by *typical gun* in (31), mod_l can also return an l_i that enriches or alters the default E-structure (e.g., to satisfy the type requirements for functional application). At this point, it is open precisely how the ‘expressions’ and ‘contexts’ which are taken as arguments by mod_l are represented (e.g., do expressions encode grammatical class and prosody/intonation?), and how this determines the selected l_i (esp., in the relatively ‘free’ cases). We will be in a better position to address this issue after discussing some case studies in Sect. 4.

Second, the C-structure of *complex* expressions is (generally) computed by $\llbracket \rrbracket_M^c$. This operation is specified in terms of the function f in FA_M , which takes the full meanings of constituents and returns the C-structure of the combination. In the cases we consider in this paper, f just performs pointwise functional application at each dimension of C-structure, as illustrated in (29) and (31). Why hold that generating the C-structure of complex expressions is part of our semantic competence? The main

reason is that, in some cases, the C-structure of a complex expression is involved in determining the E-structure of more complex expressions of which it is a constituent.¹⁹ Some of the examples examined in Sect. 4 have this property.²⁰

Third, although C-structure is generally computed, E-structure is the engine of composition, which can continue even when there is no information in C-structure. When that is the case, all the dimension operators which return values of C-structure will be useless (i.e., undefined), and mod_l will simply select E and return the E-structure of its argument. This entails that this multidimensional semantics is compatible with the view that, sometimes, C-structures for complex expressions cannot be compositionally derived from the C-structures of their immediate constituents (see Fodor 1998; Fodor and Lepore 2002; for a critical response see Prinz 2002; Hampton and Jonsson 2012; Del Pinal 2016). Now, when a complex expression has no C-structure, no modulations will be linguistically available. Still, interpretation can continue: for on this view, what is said by a complex expression S is determined by $\llbracket S \rrbracket_{M_E}^c$, i.e., by the E-structure of S .²¹

Fourth, although I have advertised this multidimensional semantics as broadly compositional, it is important to note that what is strictly compositional is the total meaning function $\llbracket \rrbracket_M^c$. If we focus just on the E-dimension of a complex expression α , we can see that it need not be strictly compositional relative to the E-dimension of the constituents of α . For example, in cases such as *typical gun* in (31), mod_l can enrich the E-dimension of a constituent of α with one of its C-dimensions, which is then inherited

¹⁹ There are additional reasons for holding that the C-structure of complex expressions is often computed. C-structure plays a key role once linguistically encoded meanings interface with other cognitive systems. For example, the perceptual dimension of C-structure is often used for categorization: in the usual case, when asked to, say, bring a *red shirt*, one has to compute (part of) the C-structure $f(\llbracket red \rrbracket_M^{c_1}, \llbracket shirt \rrbracket_M^{c_2})$, which includes the perceptual prototype. Still, information such as that red shirts are normally red on the outside is not usually part of the E-meaning of *red shirt*.

²⁰ As a reviewer pointed out to me, my multidimensional semantics for $\llbracket \rrbracket_M^c$ is structurally similar to some multidimensional systems for use-conditional (incl., expressive) content (esp., Portner 2007; McCready 2010; Gutzmann 2015). The work in this tradition developed mainly as a response to the strict constraints imposed by Potts (2005)'s seminal work regarding how non at-issue meaning can interact with surrounding linguistic material. In future work, I will compare these multidimensional systems in more detail. Still, I would like to note here that, in my view, we can treat C-structure as a species of use-conditional content, in the sense of Gutzmann (2015). If so, the arguments and system developed for $\llbracket \rrbracket_M^c$ support Gutzmann (2015)'s claim that we must enrich the compositional semantics of Potts (2005)'s original multidimensional system to allow (i) mixed lexical items defined for all dimensions of meaning, (ii) compositionality at the use-conditional level (as in the computation of C-structures for complex expressions), and (iii) a more dynamic interaction between the various dimensions of meaning (as in cases where the C-structure of a constituent affects the E-structure of a complex expression).

²¹ Interestingly, this implementation of a multidimensional semantics, according to which complex expressions are often assigned a C-structure, allows us to incorporate an account of 'loose speech' which is somewhat similar to Lasnik (1999)'s influential 'pragmatic halos' account. We cannot pursue this topic in detail, but just to briefly advertise this application of the theory, consider the following rule:

- An assertion of S of type $\langle t \rangle$ (or $\langle s, t \rangle$, were s ranges over worlds) is loosely-speaking-true if any dimension in $\llbracket S \rrbracket_{M_C}^c = 1$.

For example, *John is a lion* is 'loosely-speaking-true' if John looks like a lion, behaves like a lion, or, more generally, has some of the properties represented in the C-structure of *lion*.

into the E-structure of α .²² As we will see in Sect. 4, this dynamic mechanism is precisely what will allow us to model key cases of semantic flexibility, and can in the end be taken as an argument for the view that, in natural languages, E-structure satisfies only ‘general compositionality’, as defined by [Pagin and Westerstahl \(2010a, b, c\)](#).

Fifth, since $\llbracket \cdot \rrbracket_M^C$ is not strictly compositional along the E-dimension, the computation of E-structure for a complex expression can turn out to be redundant at a subsequent stage of composition. To see this, consider again *steel gun* in (29), and assume that it is modified by *typical*, defined as in (30). The resulting E-structure of *typical steel gun* is a kind of enrichment: we keep the original E-structure of *steel gun* but also upload its PERCEPTUAL dimension. However, there seem to be contexts which drop the E-structure of *steel gun*, and replace it with something in its C-structure. Consider *steel gun look-alike*. In our system, it is natural to assume that what *look-alike* does along the E-dimension is replace the E-structure of its argument with the PERCEPTUAL of its C-structure. If so, bottom-up interpretation would in this case calculate the E-structure of *steel gun* only to drop it at a later stage of composition. This unusual property of our multidimensional semantics is justified to the extent that the framework as a whole, and this property in particular, is descriptively valuable. Establishing that is the aim of Sect. 4.

4 Semantic flexibility in a multidimensional semantics

The framework we have introduced to let the compositional operations which determine E-structure interact with C-structure can be used to deal with various examples of enrichment (via core enrichment operations) and loosening (via dimension operators). The task now is to see how some of the phenomena of semantic flexibility which motivated theorists to adopt free modulation—in particular, expressions involving privative and colour adjectives—can be modeled in our multidimensional Putnam-style semantics. To be clear, the aim here is not to present final accounts in each case; indeed, somewhat different accounts are compatible with the overall framework. The aim is rather to present some detailed and plausible accounts of semantic flexibility, show how they overcome the problems faced by less constrained truth-conditional pragmatic theories, and thereby illustrate the advantages of the kinds of tools made uniquely available by a multidimensional Putnam-style semantics.

4.1 Privative adjectives

Consider first privatives and certain special adjectives which seem to operate on the internal semantics of common nouns. We have seen that, to deal with privatives, truth-conditional pragmatists appeal to free *mod*, which systematically generates unavailable readings ([Partee 2007, 2010](#); [Recanati 2010](#)). In contrast, a Putnam-style

²² In other words, this multidimensional semantics has the following property. In a context such as *typical gun*, we couldn’t substitute for *gun* a term such as *gun*₁—which we assume has the same E-structure of *gun* but a different C-structure—and guarantee the preservation of the E-structure of the phrase. Preservation of E-structure in a context like *typical gun* also requires that *gun* and *gun*₁ have the same C-structure.

semantics can be used to give promising accounts of privatives which do not over-generate meanings.

Paradigmatic examples of privative noun phrases are expressions like *fake gun*, *counterfeit dollar* and *artificial heart*. In most contexts, a normal gun would not count as a *fake gun*. At the same time, a phone is normally not a fake gun—in other words, a fake gun is an artifact that looks, handles, etc. like a gun but does not function like one. Furthermore, we don’t want to say that a malfunctioning gun is a fake gun, or that all fake guns succeed in really looking like guns. Intuitively, a fake gun is an artifact that was made to look like (or more generally, to have the perceptual features of) a gun, but not to function like one. How can we compositionally derive this behavior?

In a Putnam-style semantics, many common nouns encode, in their C-structure, the information we invoke in intuitive descriptions of what privative adjectives seem to do to their arguments. For this reason, we can treat terms such as *fake* as what they intuitively seem: modifiers of the semantic structure of predicates. As a first approximation, consider the entry for *fake* in (32), were ‘ D_C ’ ranges over full meanings, i.e., tuples of the E-structure and each dimension of C-structure. Note that each dimension takes the full meaning of the argument N of *fake*, and uses dimension operators to specify its effect. Hence for functional application to proceed, the mod_i which ranges over the sister N has to return its full meaning ($= \llbracket N \rrbracket_M^c$). This can be done via the *ALL* core enrichment operator.

$$(32) \quad \llbracket fake \rrbracket_M^c =$$

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

C-structure:

C: $\lambda D_C. Q_C(D_C)$

P: $\lambda D_C. Q_P(D_C)$

T: $\lambda D_C. [\lambda x. \neg Q_T(D_C)(x) \wedge Q_P(D_C)(x)]$

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

Given (32) and entry (23) above for *gun*, $\llbracket fake\ gun \rrbracket_M^c$ results, by FAM , in (33). Its E-structure encodes the following condition: it is satisfied by entities that are not guns, were not made to be guns, and were made to have the perceptual features of guns. This allows, crucially, that a fake gun can be badly made and not look like a gun. As in other cases, the C-structure of *fake gun* is more informative than its E-structure: e.g., it says that fake guns look like guns but don’t shoot. This is because, according to (32), a *fake* N will be ‘conceived’ as having the perceptual features but not the function of N, but this is not part of its E-structure. We return to the advantage of having this information in C-structure when we explore, in Sect. 4.3, how expressions like *fake gun* interact with modifiers such as *typical* and *perfect*.

$$(33) \quad \llbracket fake\ gun \rrbracket_M^c =$$

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

C-structure:

C: $Q_C(\llbracket gun \rrbracket_M^c)$

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

Crucially, this account does not generate the unavailable interpretation mentioned in (21-a), which we used to criticize Partee (2007, 2010)’s account. Given the multi-dimensional type of the argument of *fake*, the head noun must provide a full meaning for functional application to proceed. The privative *fake* then modifies that meaning in the way specified in (32). Hence the default (at-issue) interpretation of *fake gun* is not based on modulating the meaning of *gun* as a function of the wider context; it is just based on providing the privative with the type required for functional application. Consequently, this account does not predict that there should be a different reading in contexts that call for the default reading and in special contexts such as the terrorists at the halloween party in (21-a).²³

How does entry (32) fare when combined with other lexical items? Del Pinal (2015a) argues in detail that, given independently motivated entries for the head nouns, an entry close to (32) gives the intuitively correct result for a wide range of cases, including *fake pen* and *fake gold* and less obvious examples such as *fake smile* and *fake lawyer*. For current purposes, however, it is not crucial that we agree on every detail of (32).²⁴ The aim here is only to motivate the kinds of representations and operations that a Putnam-style semantics makes available by providing promising accounts of privative adjectives which need not appeal to something like free *mod*.

Given this aim, it is important to show that this basic account can be extended to other privatives. Indeed, the tools made available by a Putnam-style semantics become especially useful when the differences are subtle, as in the case of *counterfeit* and *artificial*, which contrast in interesting ways with each other and with *fake*. An important difference between many paradigmatic uses of *counterfeit* and *fake* is that, unlike a *fake*, a *counterfeit* N is usually made to look and function like an N. For example, a *counterfeit* Rolex is made both to look and function like a Rolex. *Counterfeit*, in its 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. This is captured in the following partial entry:

²³ To be clear, in a Putnam-style framework, $\llbracket \text{fake gun} \rrbracket_M^c$ can itself be the argument to *mod*_l (e.g., when it is combined with other expressions). When the C-structure of *fake gun* is computed, as in (33), all the dimensions of C-structure are in principle available to *mod*_l. Crucially, none of these dimensions encodes a predicate which must be satisfied by entities with the function of a gun, i.e., of shooting (e.g., consider the TELIC in (33)). This explains the unavailable reading in (21-a). At the same time, this mechanism also explains why, once we consider *fake gun* in wider contexts, we can observe additional flexibility. To illustrate, suppose that *mod*_l filters $\llbracket \text{fake gun} \rrbracket_M^c$, specified as in (33), using the dimension operator *Q*_T. In that case, its E-meaning will be equal to *Q*_T($\llbracket \text{fake gun} \rrbracket_M^c$). Hence, anything that passes some perceptual similarity threshold (and doesn’t have the relevant function of guns, e.g., shooting in the relevant way) may count as a *fake gun* (e.g., a lump of wood that, for whatever reason, looks like a gun). For further discussion, see Sect. 4.3 below.

²⁴ For example, one could argue that we should drop the condition that, in its E-structure, *fake* negates the E-structure of its nominal argument. If so, the inference that ‘a fake N is not an N’ would result from the intuition that, if an artifact *x* was created to, say, look but explicitly not to function like an N, it is, at best, extremely unlikely that it is an N.

$$(34) \quad \llbracket \text{counterfeit} \rrbracket_M^c =$$

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

C-structure:

...

If we apply (34) to a plausible lexical entry for *Rolex*, we get as its E-structure the result 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 $Q_A(\llbracket \text{Rolex} \rrbracket_M^c)$, and is an artifact that was made to look *and* function like a Rolex, as specified in $Q_P(\llbracket \text{Rolex} \rrbracket_M^c)$ and $Q_T(\llbracket \text{Rolex} \rrbracket_M^c)$ respectively.²⁵ Consider, in contrast, 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 is made to function like one. The following partial entry captures the contrast between *artificial*, on the one hand, and *fake/counterfeit* on the other:

$$(35) \quad \llbracket \text{artificial} \rrbracket_M^c =$$

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

C-structure:

...

Perhaps more clearly than in *fake* and *counterfeit*, in this case we might want to eliminate ‘ $\neg Q_E(D_C)(x)$ ’ from (35). Speakers I have informally queried about whether, say, artificial hands and legs are really arms and legs report mixed and unstable intuitions. This is ultimately an empirical question, and both options can be implemented in a Putnam-style semantics.

The Putnam-style tools which we used to model privatives can also be used to model modifiers which, although not strictly privative, also seem to access the C-structure of their arguments. Consider certain attributive uses of *true* and *real*, recently examined by Knobe et al. (2013), Leslie (2015) and Del Pinal and Reuter (2017). For example, Knobe et al. (2013) argue that *true*, when used as in (36-a), seems to operate on a ‘normative’ dimension which they argue is part of our conceptual representations of certain classes, esp., of social roles. In particular, *true* displays an acceptability pattern different from that of expressions which operate on the function (i.e., TELIC) of social roles, such as *good* in (36-b).

- (36) a. Mary is a true scientist.
 b. Mary is a good scientist.

Studies show that there are contexts in which participants agree that Mary is a *good* scientist but not a *true* scientist. For example, a scenario where Mary is a skilled

²⁵ On this account, one difference between fakes and counterfeits is that the latter are made to function like the artifacts they are counterfeits of, whereas this is not required in the case of fakes, although in some cases it is possible to so use them.

scientist but is ultimately not committed to the relevant scientific values (e.g., she works mainly to be famous). There are also contexts in which participants agree that Mary is a *true* scientist even if she is not a particularly good or skillful scientist. For example, a scenario where Mary is committed to satisfy the values associated with being a scientist, although she is not (yet) very skilled as a scientist (see [Knobe et al. 2013](#); [Leslie 2015](#); [Del Pinal and Reuter 2017](#)). These acceptability patterns suggest that some social role terms have a C-structure which includes a normative dimension that represents something like the basic values/commitments associated with the role. This dimension is uploaded into the E-structure of complex noun phrases when social roles are modified by *true*. [Leslie \(2015\)](#) argues that slurs and social recriminations such as *real men don't cry!* also often involve similar operations on the normative dimension.

These kinds of cases can be easily modeled in our Putnam-style semantics. Consider the partial entries in (37) and (38). When used in the relevant way, *true* and *real* take the full meaning of their arguments, and modify them such that the value of the normative dimension replaces the value of the original E-structure. Entries (37) and (38) include a dimension operator, ' Q_N ', which returns the value of the normative dimension (if there is one, as in the case of social role terms, otherwise the modification is undefined). The entry for *real* is a bit more complex because, arguably, x is a *real scientist* requires that x satisfies both the normative and functional dimensions of *scientist*.

(37) $[[true]]_M^c =$
E-structure: $\lambda DC. [\lambda x. Q_N(DC)(x)]$
C-structure:
 ...

(38) $[[real]]_M^c =$
E-structure: $\lambda DC. [\lambda x. Q_N(DC)(x) \wedge Q_T(DC)(x)]$
C-structure:
 ...

Compared to privatives, the E-structures of these modifiers is simple. Still, we seem to approximate the desired result. To illustrate, we have seen that native speakers accept that Mary can be a *true scientist* even if she is not, in the ordinary sense, a scientist. Entry (37) captures this: the E-structure of *true* N does not include the E-structure of N.

Interestingly, since what *true/real* do to the E-structure of its argument is to replace it with dimensions of its C-structure via dimension operators, this Putnam-style account predicts that, in some contexts, their effect can be achieved solely by *mod_l*. This prediction is formulated in (40), and is illustrated by derogatory remarks such as (39-a), in which the target social role term is not explicitly modified by *true* or *real* ([Leslie 2015](#)). Given (40), we can see why (39-a) and (39-b) often have indistinguishable at-issue contents.

- (39) a. Hillary Clinton is the only man in the Obama administration!
- b. Hillary Clinton is the only real man in the Obama administration!

- (40) Let α be a modifier of β in context c . If α is ‘simple’ (i.e., its effect on the E-structure of β is to replace it with a dimension of the C-structure of β), then there is a context c' such that $\llbracket \alpha \beta \rrbracket_{ME}^c = \text{mod}_I(\beta, c') \llbracket \beta \rrbracket_M^{c'}$.

In short, the effect of ‘simple’ modifiers on the E-structure of the head can be achieved, in some contexts, simply via mod_I . ‘Intersective’ modifiers which add descriptive content (e.g., *plastic* in *plastic gun*), and paradigmatic privative adjectives (e.g. *fake* and *counterfeit*) are not ‘simple’, according to this definition. Hence there is no expectation that we can find a context which induces, solely via mod_I , their effect on the E-structure of the head.

Finally, some modifiers seem to alter the centrality ordering of the C-structure of their arguments. Indeed, *true* and *real*, in the sense discussed above, seem like paradigmatic members of this class (see the studies reported in Del Pinal and Reuter 2017). I will not discuss these effects in detail here, but just note that they can be easily modeled in a Putnam-style semantics. Here is one option. Suppose that the full meaning structure of a social role term is as in (41), where the degree of centrality of each dimension of C-structure is represented by its relative position in the tuple, such that the AGENTIVE is the most central dimension, followed by the TELIC, and so on. Note that the C-structure effect of the entry in (42), when applied to (41), is just to change the centrality orderings amongst the dimensions. In the output, the TELIC is now the most central dimension, followed by the NORMATIVE, and so on. In this way, we can model changes that certain modifiers effect on our conceptions, even when these changes do not directly add any descriptive content. We explore one reason for modeling centrality in Sect. 5.2

- (41) $\llbracket \text{scientist} \rrbracket_M^c =$
E-structure: $\lambda x. \text{SCIENTIST}(x)$
C-structure: $\langle A, T, N, P \rangle$
- (42) $\llbracket \text{real} \rrbracket_M^c =$
E-structure: $\lambda DC. [\lambda x. Q_N(DC)(x) \wedge Q_T(DC)(x)]$
C-structure:
 $\langle \lambda DC. Q_T(DC), \lambda DC. Q_N(DC), \lambda DC. Q_A(DC), \lambda DC. Q_P(DC) \rangle$

4.2 Colour adjectives

We have seen that a multidimensional Putnam-style semantics can be used to give adequate, constrained accounts of the compositional behavior of privative and subjective modifiers whose flexibility seems to depend on having access to the C-structure of their arguments. Since those kinds of modifiers are arguably special, and our multidimensional semantics does add some complexity to the standard truth conditional theories, we might reasonably ask what the advantages are of adopting this framework relative to more standard examples of semantic flexibility. In this respect, colour terms are an excellent case study. For although often taken as paradigmatic ‘intersective’ adjectives, we have seen that—as captured in the Travis color cases in (10)–(11)—their

compositional behavior exhibits the kind of semantic flexibility which has led some theorists to adopt free *mod*.

There are various promising accounts of color adjectives that address the Travis cases without appealing to anything like free *mod*. One of the most convincing accounts, using a standard compositional truth-conditional theory, is presented by Kennedy and McNally (2010). Based on the way in which colour adjectives interact with degree (e.g., *very*, *completely*) and comparative (e.g., *more*, *less*) morphology, they argue that colour adjectives are ambiguous between two gradable and a non-gradable, classificatory sense. To illustrate this, consider the gradable senses of *green* presented in (43) and (44):

$$(43) \quad \llbracket green_{quant} \rrbracket = \lambda x. \text{QUANT}(\text{GREEN})(x)$$

$$(44) \quad \llbracket green_{qual} \rrbracket = \lambda x. \text{QUAL}(\text{GREEN})(x)$$

These gradable senses are functions of type $\langle e, d \rangle$, i.e., from entities to degrees. Colour adjectives are associated with two measure functions. One measures how much of the entity manifests the colour, as in (43). The other measures how closely the entity’s manifestation of the colour approximates the relevant prototype, as in (44). These functions of type $\langle e, d \rangle$ are converted into functions of type $\langle e, t \rangle$ when they compose with degree morphemes. For example, the null degree morpheme *pos*, presented in (45), introduces a relation to a contextual standard of comparison. When *pos* combines with (44), as in (46), it results in a predicate of type $\langle e, t \rangle$ which is satisfied by entities whose degree of greenness_{qual} in *c* passes the standard in *c*.

$$(45) \quad \llbracket pos \rrbracket^c = \lambda g_{\langle e, d \rangle} \lambda x. g(x) > \text{STND}_c(g)$$

$$(46) \quad \llbracket pos green_{qual} \rrbracket^c = \lambda x. \text{QUAL}(\text{GREEN})(x) > \text{STND}_c(\text{QUAL}(\text{GREEN}))$$

The non-gradable, classificatory sense of *green* is presented in (47):

$$(47) \quad \llbracket green_{nongr} \rrbracket^c = \lambda x. P_i(x) \wedge \text{COR}(P_i, \text{GREEN})$$

Unlike the gradable senses in (43)–(44), the sense in (47) is of type $\langle e, t \rangle$. Its satisfaction condition is only that the entity to which it applies falls under a property that is correlated with greenness. This property is represented by the free variable P_i , and Kennedy and McNally suggest that it is typically determined by features of the argument to the colour adjective. Finally, to capture the intuition that senses (43), (44), and (47) are related, each entry is specified in terms of ‘GREEN’, a mass noun denotation of type $\langle e \rangle$.

Kennedy and McNally’s account can be directly applied to the Pia green leaves cases (10)–(11). The interpretation of (10-a) as (10-b) is obtained by disambiguating *green* to (44), given context (10). And the interpretation of (11-a) as (11-b) is obtained by disambiguating *green* to (47), given context (11). Crucially, on this account we need not assume that there is a free *mod* operation that takes the literal meaning of *green* and outputs the relevant refined or enriched readings. The main readings are provided via disambiguation, and each of the disambiguated senses has parameters of

context-sensitivity—the relevant standards for the gradable readings and the correlated property for the non-gradable reading.²⁶

Since Kennedy and McNally's account can deal with Travis cases without postulating anything like free *mod*, what do we gain by implementing it in a multidimensional semantics? To answer this question, consider first whether an ambiguity-based implementation of the various senses of colour terms seems intuitively plausible. If we reflect on how we think and talk about colours, is it clear that we are always primarily thinking along *one* of these senses? Or is it more accurate to hold that we have a relatively unified conception of 'being coloured' that has these senses as dimensions, various *combinations* of which can be at-issue in a given colour-involving assertion? The latter view would be supported by examples where the at-issue content of colour predications involves conjunctions of two or more of these dimensions. Interestingly, such cases abound.

Suppose you want a green car. In the usual case, you would have concrete expectations with respect to both the quality *and* quantity of its greenness, and even with respect to the way in which it became green (e.g., that it be factory painted). If, in some context, you believe that your interlocutors might miss the full range of your expectations, you could communicate using modified predications of *green*, as in (48):

(48) I want a car that is factory green and perfectly so.

Consider another example. Although in some cases a *naturally green leaf* need not be green, a *typical naturally green leaf* usually refers to leaves that are both *green_{qual/quant}* and *naturally green_{nongrad}*. However, in Kennedy and McNally's implementation, colour words, once in use, have to take one of the three senses. Once we are talking about *naturally green* leaves, the gradable senses of *green* are no longer available. In general, then, it is not clear how to get the default readings of expressions like those in (49):

- (49) a. Pia wants leaves that are perfectly, completely and naturally green.
 b. Pia wants completely/perfectly, naturally green leaves.
 c. Pia wants typical naturally green leaves.

The problem remains even if some of these cases involve ellipsis. For example, suppose, as seems reasonable, that the LF of (49-a) is something like '*...completely green₊ and naturally green₁*'. Following standard accounts of ellipsis, the elided *green* should have the same meaning as *green₁*. Still, this meaning would have to correspond to one of the disambiguated senses, and so on this view it could not shift between the gradable and non-gradable senses, as seems to be required to get the default reading of (49-a).

Is there an easy fix? Kennedy and McNally point out that one could implement the two gradable senses as a binary choice of a gradable-type parameter. This suggestion does not solve the present problem. For example, it doesn't address readings that call

²⁶ Hansen (2011) argues that the gradable and non-gradable senses of colour terms have additional context-sensitive parameters. The refinements proposed by Hansen (2011) can be implemented in a Putnam-style semantics, but doing so here would add unnecessary complexity to the presentation.

for mixed gradable and non-gradable interpretations, as in (48) and (49-b). Also, it doesn't help deal with modifiers that seem to simultaneously operate on various senses, as in (49-c). Another possible response is to drop the constraint on elipsis according to which elided terms are copied from disambiguated logical forms. However, as is well known, this would massively over-generate readings. For example, (50) would be incorrectly predicted to have a reading where John went to a financial bank and Mary went to a river bank.

- (50) John went to the bank to cash a check, and Mary went to the bank to get some aesthetic enjoyment.

Finally, we cannot deal with these examples by conjoining all dimensions into one long 'definition', since there are many cases in which colour predications are minimal, i.e., in which only one of the senses is at-issue. Taken together, these problems point to the difficulty of encoding 'conceptions' in a standard one-dimensional semantics: Kennedy and McNally avoid a 'definitional' view, but only by adopting a problematic ambiguity-based implementation.

These problems can be resolved, however, if we implement Kennedy and McNally's account in a multidimensional Putnam-style semantics. The proposal I present here proceeds from the observation that each of the gradable and non-gradable senses of colour terms is closely related to particular dimensions of C-structure. This suggests a way of representing the various senses of a colour term without having to posit lexical ambiguity. This is illustrated in (51) for *green*, where each sense is a dimension of its C-structure. For simplicity, we assume that the CONSTITUTIVE dimension encodes the quantitative gradable sense, the PERCEPTUAL encodes the qualitative gradable sense, and the AGENTIVE encodes the nongradable reading.

- (51) $\llbracket green \rrbracket_M^C =$
E-structure: $\lambda x. GREEN(x)$
C-structure:
 C: $\lambda x. QUANT(GREEN)(x)$
 P: $\lambda x. QUAL(GREEN)(x)$
 T:
 A: $\lambda x. P_i(x) \wedge COR(P_i, GREEN)$

A key difference between the entries for nouns like *lion* and those for colour terms is that, in the latter case, the CONSTITUTIVE and PERCEPTUAL are not of type $\langle e, t \rangle$.²⁷ Still, modifiers of colours, including degree morphemes, can be formulated analogously to semantic restructuring operators. They take full meanings (which in this case include dimensions of type $\langle e, d \rangle$), select dimensions of C-structure, and add the relevant standards to turn them into functions of type $\langle e, t \rangle$. For example, *par-*

²⁷ According to this entry, colour terms can be used in a 'minimal' way, since they have an atomic non-gradable E-meaning. This will be the at-issue content when *mod_l* simply returns the E-structure. However, it is compatible with our Putnam-style semantics that there are classes of words which don't have a default E-structure—hence on each use *mod_l* would have to select, for composition to proceed, some dimension/s of the conception and upload it into the E-structure—and that colour words belong to this class.

tially adds a low standard to the CONSTITUTIVE and uploads it into the E-structure of its argument; *very* adds a high standard to the PERCEPTUAL and uploads it into the E-structure of its argument; *naturally* adds a constraint to the AGENTIVE non-gradable reading and uploads it into the E-structure of its argument.

As mentioned before, colour adjectives can be modified by null degree morphemes. Following Kennedy and McNally, we will also assume that the relevant null degree morphemes are semantically like overt degree morphemes. Adjusted to our framework, this means that they take full meanings, except that in the null cases the relevant standard is contextually determined. Consider the E-structure of the entry for *posqual* in (52): it takes full meanings ($\cdot D_C^{grad}$, ranges over ordered tuples of the E-structure and dimensions of C-structure, where some of the dimensions are of type $\langle e, d \rangle$), selects the relevant dimension of type $\langle e, d \rangle$, and turns it into a predicate of type $\langle e, t \rangle$.

$$(52) \quad \llbracket posqual \rrbracket_M^c =$$

E-structure:
 $\lambda D_C^{gd} \lambda x. Q_E(D_C^{gd})(x) \wedge [Q_P(D_C^{gd})(x) \succ \text{STND}_c(Q_P(D_C^{gd}))]$

C-structure:
 C: $\lambda D_C^{gd} \cdot Q_C(D_C^{gd})$
 P: $\lambda D_C^{gd} \lambda x. Q_P(D_C^{gd})(x) \succ \text{STND}_c(Q_P(D_C^{gd}))$
 T:
 A: $\lambda D_C^{gd} \cdot Q_A(D_C^{gd})$

If *posqual* is combined with entry (51) for *green*, we get the result in (53), where its E-structure is satisfied by entities whose qualitative degree of greenness passes a certain contextually determined standard.

$$(53) \quad \llbracket posqual \ green \rrbracket_M^c =$$

E-structure:
 $\lambda x. \text{GREEN}(x) \wedge [\text{QUAL}(\text{GREEN})(x) \succ \text{STND}_c(\text{QUAL}(\text{GREEN}))]$

C-structure:
 C: $\lambda x. \text{QUANT}(\text{GREEN})(x)$
 P: $\lambda x. \text{QUAL}(\text{GREEN})(x) \succ \text{STND}_c(\text{QUAL}(\text{GREEN}))$
 T:
 A: $\lambda x. P_i(x) \wedge \text{COR}(P_i, \text{GREEN})$

In general, any of the dimensions of the C-structure of *green* can be uploaded into the E-structure, depending on the context and the overt/covert modifiers. However, regardless of which dimension is uploaded, if any, the other dimensions in the C-structure are available for further composition, as illustrated in (53). As a result, we can now account for cases such as *Pia wants leafs that are completely, naturally green*. When *naturally* composes with *green*, the gradable readings are still available in C-structure, so this expression can be further modified by intensifiers of the gradable senses such as *perfectly/completely*, in which case we get a rich interpretation such that the at-issue meaning refers to leafs that are prototypically/completely, and naturally green.

What about cases like (49-a), which arguably involve ellipsis? Suppose expressions like (54-a) have LFs roughly like (54-b):

- (54) a. ... leaves that are completely and naturally green
 b. ... leaves that are completely ~~green~~_T and naturally green₁

Following standard views, we assume, again, that in (54-a) the elided term must be a copy of *green*₁, as captured in (54-b). Since this is specified at a level of representation where ambiguities are resolved, this means that the copied term has the same meaning as the source term. Given the combinatorial properties of our Putnam-style semantics, this condition has to be formulated as saying that by ‘meaning’ we technically mean $\llbracket \cdot \rrbracket_M^c$, i.e., full meanings. Assuming an entry for $\llbracket \textit{green} \rrbracket_M^c$ as in (51), we get the target readings. For that entry can be combined with modifiers such as *completely*, *perfectly*, and *naturally*, returning in each case a slightly different E-structure.

We just said that, given the combinatorial rules of a Putnam-style semantics, standard constraints on ellipsis entail that full meanings, and not just E-meanings, are copied from the source unto the sites of the elided expressions. This implication for ellipsis, which results from unique features of our multidimensional semantics, has additional advantages. One is that it can be used to account for Chomsky’s famous examples of sense modulation across elided/anaphoric material (2000). This phenomena is illustrated in (55):

- (55) This book₁ is very interesting, but it₁ is simply too heavy to take on the plane.

Briefly, the key observation here is that *book*₁, given its full multidimensional meaning, can be modulated in different ways at each site. At the same time, this framework doesn’t overgenerate elided readings in cases like (50), since none of the E/C-dimensions of the concept of a river bank are E/C-dimensions of the concept of a financial bank, and vice versa.

4.3 Interaction between privatives, colours, and subjectives

To conclude these case studies, we examine combinations of privatives, colours terms, and subjective modifiers. As far as I know, most of these interactions have not been explored in the relevant literature. Showing that our account makes adequate predictions in these cases further supports the view that a Putnam-style multidimensional semantics issues in unique and promising tools for compositional semantics.

Let us begin by examining the predictions of our account for combinations of privatives and colour terms. As our target case, take [*fake [green leaf]*], as used in examples like (56). (56) has two main readings, paraphrased (roughly) in (56-a) and (56-b).

- (56) That₁ is a fake green leaf.
 a. x_1 is a green object but a fake leaf.
 b. x_1 is a real leaf but is not naturally green.

To see that our Putnam-style semantics predicts these readings, assume that $\llbracket \textit{fake} \rrbracket_M^c$ is as in (32), and $\llbracket \textit{green leaf} \rrbracket_M^c$ is obtained by turning all dimensions of type $\langle e, d \rangle$,

in (51), into type $\langle e, t \rangle$ and conjoining each dimension of *green* with the corresponding dimension of *leaf*.²⁸ Given those assumptions, $[[fake [green leaf]]]_{ME}^c$ can be satisfied in at least two ways. Recall that $[[fake]]_{ME}^c$ negates the AGENTIVE in the C-structure of the modified expression, and uploads that into the E-structure of the resulting expression. In this case the modified expression is a complex term, and its AGENTIVE includes the condition imposed by *green* conjoined with that imposed by *leaf*. Since $[[fake]]_{ME}^c$ negates that conjunction, then something that falls under it can lack either or both conjuncts. As a result, $[[fake [green leaf]]]_{ME}^c$ is satisfied by an entity that is green but a fake leaf, which corresponds to the reading in (56-a). This is arguably the default. However, it is also satisfied by an entity that is a real leaf but which does not satisfy the correlation between greenness and the contextually relevant property, e.g., a real leaf that is naturally red but was painted green. This corresponds to the reading in (56-b).²⁹

That this second reading is available is easier to see in cases such as *fake red Ferrari*. As before, the prediction is that this can mean either a fake Ferrari that is red (the usual meaning), or a real Ferrari (possibly $red_{qual/quant}$) that is not $red_{nongrad}$. To get the latter reading, consider the following scenario. For car collectors, a classic Ferrari that is originally/factory red is more valuable than one that is later re-painted red. Since this is known, several sketchy dealers repaint red Ferrari's which were not originally red. Now consider the following statement:

- (57) John is a car collector who unknowingly acquired a repainted classic Ferrari. Discovering that he has been swindled, John says:
- a. Damn! This is a fake red Ferrari!

In context (57), (57-a) is intuitively true, even if John's Ferrari is painted red. It seems, then, that the reading entailed by our account of the E-structure of expressions like $[[fake [red Ferrari]]]$ is attested. And just as an entity can be a fake red Ferrari, and be both $red_{qual/quant}$ and a Ferrari, an object can be a fake green leaf, and be both $green_{qual/quant}$ and a leaf.³⁰ In both cases, what has to be lacking for this reading to be satisfied is the relation between the colour and the contextually relevant property of the AGENTIVE.

²⁸ Two caveats. First, in what follows I drop the condition that the E-structure of *fake* negates the E-structure of the modified term. For further discussion and justification of this move, see footnote 24. Secondly, I have not discussed the combinatorial rules for complex predicates whose immediate constituents have E-structures of type $\langle e, t \rangle$. Most type-driven theories include a rule of predicate modification which outputs a complex predicate of type $\langle e, t \rangle$ which is basically a conjunction of the constituents (e.g., Heim and Kratzer 1998). Another option is to use type shift rules, which, in these cases, could output the same result. Both tools can be formulated in a Putnam-style semantics. Either way, in a case like *green leaves*, where both constituents have an E-structure of type $\langle e, t \rangle$, the result will be predicate conjunction along every dimension, as specified above.

²⁹ This second reading is the only available reading if the structure is $[[fake green [leaf]]]$.

³⁰ Consider again the botanist version of the Pia story. We can easily imagine a continuation in which the botanist rightfully protests, upon discovering that s/he has been handed painted green (naturally red) leaves, *Pia, you gave me fake green leaves!*, even though the leaves are not themselves fake and are $green_{qual/quant}$.

We next examine how our multidimensional semantics can be used to model sub-*typical* modifiers such as *typical* and *normal*, and in particular their interaction with privatives and colour terms. We focus on a (common reading) of *typical*. Its first key property is that, in general, *typical* N seems to have the effect of enriching the E-structure of N with (part of) its C-structure. The extent of the enrichment seems to be highly context-sensitive. This is illustrated by the contrast between the response in (58-a), which is somewhat deficient, and the one in (58-b), which is fine.

- (58) a. John: That’s a lion.
David: # No, it’s not. It doesn’t even have a mane.
- b. John: That’s a typical lion.
David: No, it’s not. It doesn’t even have a mane.

The second key property of *typical* is that expressions such as *typical lion* can be modified by covert and overt degree and comparative morphemes, even if the head, in this case *lion*, is not normally a gradable expression. So although in (59)–(60) the (a) cases are deficient, the (b) cases are fine:

- (59) a. *That is a very lion.
 b. That is a very typical lion.
- (60) a. *B is a more lion than C.
 b. B is a more typical lion than C.

(59)–(60) suggest that the output of an expression of the form ‘*typical* N’ needs to have at least some dimensions of type $\langle e, d \rangle$, even if no dimension in the unmodified N is of that type.

Taking a cue from the account of colours, I suggest that what the E-structure of *typical* does is to take full meanings and return a function of type $\langle e, d \rangle$. Crucially, since in a Putnam-style framework modifiers can access the C-structure of their arguments, we can easily construct a procedure that outputs the required function of type $\langle e, d \rangle$. As a first approximation, consider entry (61). In the description of the E-structure, $T(D_C, x) = \{P_{\langle e, t \rangle} \in D_C \mid P(x) = 1\}$ and $|S|$ is a function that returns the cardinality of set S . Accordingly, the former operation takes a full meaning tuple and an entity, and returns the set of predicates in the tuple which are true of that entity, and the latter function takes that set and returns its cardinality.³¹

- (61) $\llbracket \textit{typical} \rrbracket_M^c =$
E-structure: $\lambda D_C \lambda x. |T(D_C, x)|$
C-structure:
 C: $\lambda D_C. Q_C(D_C)$
 P: $\lambda D_C. Q_P(D_C)$
 T: $\lambda D_C. Q_T(D_C)$
 A: $\lambda D_C. Q_A(D_C)$

³¹ For simplicity, we assume here that the C-structure of *typical* is basically inert.

If we combine entry (61) with *lion*, we obtain (62):

$$(62) \quad \llbracket \textit{typical lion} \rrbracket_M^c =$$

E-structure: $\lambda x. |T(\llbracket \textit{lion} \rrbracket_M^c, x)|$

C-structure:

C: $Q_C(\llbracket \textit{lion} \rrbracket_M^c)$

P: $Q_P(\llbracket \textit{lion} \rrbracket_M^c)$

T: $Q_T(\llbracket \textit{lion} \rrbracket_M^c)$

A: $Q_A(\llbracket \textit{lion} \rrbracket_M^c)$

This result accounts for the two key properties of *typical*. The E-structure of *typical lion* is a measure function of type $\langle e, d \rangle$, so it can combine with overt and covert degree and comparative morphology, as in (59-b) and (60-b). For example, *very* can add a high-standard to the E-structure of *typical lion* and output a predicate of type $\langle e, t \rangle$. As in the case of colours, a null degree morpheme, call it '*pos_{typ}*', can add a context-sensitive standard. Assuming the standard is > 1 , at least some of the dimensions of the C-structure of *lion* will have to be satisfied by x for it to count as a *pos_{typ} typical lion*. For example, although x need not satisfy the perceptual features which are part of the conception of *lion* to satisfy the E-structure of *lion*, there are contexts in which x will need to do so to satisfy the E-structure of *pos_{typ} typical lion*. This explains the difference in acceptability between cases like (58-a) and (58-b).

Importantly, this account of *typical* coheres well with the Putnam-style account of privatives and colour adjectives. Suppose that *typical* combines with entry (33) for *fake gun* and with a degree morpheme that determines a high-standard. The result is that an entity will satisfy the E-structure of *typical [fake gun]* only if it looks like a gun and does not shoot, although neither of these conditions have to be satisfied for something to fall under the E-structure of the unmodified *fake gun*, as discussed in Sect. 4.1. Crucially, this result depends on their being a C-structure for complex phrases such as *fake gun*, which our compositional rules, as formulated in Sect. 3.2, allow us to derive. These kinds of examples suggest that the C-structures of complex expressions are often available for further composition. Similarly, the E-structure of *typical [green leaf]*, when provided with a high-contextual standard, will often require that a leaf be green in the gradable *and* the non-gradable senses simultaneously. These results seem adequate, are at least not obviously derivable on standard accounts, and hence further support our multidimensional semantics.

Finally, this account of *typical* solves a challenge faced by extant one-dimensional accounts of similar modifiers. Consider the proposals in (63-a) and (63-b). (63-a), which is loosely inspired by Carlson and Pelletier (2002)'s entry for *average*, takes a predicate and returns a measure function which, given an entity e , returns a cardinality representing the number of properties Q that apply to e , such that each Q is a property of the kind associated with P . (63-b), which is loosely inspired by McCready and Ogata (2007)'s account of Japanese modifiers such as *rashii*, achieves a similar result using instead the set of properties Q which normally (relative to the beliefs of a speaker or common ground) apply to entities that are Ps .

- (63) $\llbracket \textit{typical} \rrbracket^c =$
- a. $\lambda P_{\langle e, t \rangle} . \lambda x . |\{Q : Q \in K_c(P) \wedge Q(x) = 1\}|$
 - b. $\lambda P_{\langle e, t \rangle} . \lambda x . |\{Q : P(x) >_c Q(x)\}|.$

These proposals can arguably account for the two key properties of *typical* illustrated in (58)–(60). However, they face the following challenge, noted by McCready (2015). They both entail that, when determining the degree of typicality of an entity, all properties Q count the same. This has counterintuitive results. Suppose we want to determine which of A or B is more like a typical gun. For simplicity, assume that guns are normally colored gray, made in the USA and can shoot. Object A is gray colored, made in the USA, but does not shoot; object B is brown, made in China, and can shoot. Intuitively, object B is closer to a typical gun, even if A has a higher number of properties of normal guns. This suggests that the properties relevant to determine the typicality of entities as Ps are determined by the conception used to represent Ps , which is precisely what a Putnam-style multidimensional semantics encodes in C-structures. In this example, B is closer than A to a typical gun because only B satisfies the TELIC of guns—a dimension that is central for artifact kinds in general. In contrast, even if most guns have a particular color and country of origin, these properties are arguably not part of any dimension of our conception of guns (see Sect. 5.1). Of course, we could propose, on behalf of one-dimensional views, to add a parameter that constrains the Qs in (63-a) and (63-b) to properties which are part of the conception of Ps . However, this is basically what we have implemented in a multidimensional semantics with access to C-structures. In addition, we have implemented this in a way that allows access to the C-structure of P even when P is itself a novel complex expression.³²

5 Objections and clarifications

The discussion of privative, subsecutive and color adjectives shows that a Putnam-style multidimensional semantics provides us with tools to deal with some important cases of semantic flexibility. The case studies illustrated two tools. The first consists of modifiers which have access to the C-structure of their arguments. This allows us to model privative and other special adjectives without having to appeal to free *mod*. The second tool consists of multidimensional implementations of ‘related’ senses. This allows us to model related senses without falling into either definitional or ambiguity-based implementations. Compared to the latter, multidimensional implementations issue in better accounts of cases where various dimensions are used to enrich the core meaning of expressions. To be clear, although these tools can be used to model a wide range of phenomena, some of which I did not discuss here, I am not sug-

³² This account of *typical* is only intended as a first approximation. Instead of a simple cardinality function, a more complete account will arguably need to employ a degree function that involves a weighted sum of the dimensions. Importantly, the order of the dimensions of C-structure will constrain this function such that the more central a dimension, the higher its weight. Also, some might argue that ‘ x is a *typical N*’, presupposes, rather than asserts, that x is an N. If so, the E-structure of *typical* should be replaced with something like ‘ $\lambda DC \lambda x : Q_E(DC)(x) = 1. |T(\llbracket DC \rrbracket_{M_c}^c, x)|$ ’, with the result that x is presupposed to satisfy the E-structure of the modified predicate, and only the dimensions of C-structure are used to measure typicality. These kinds of refinements can be easily modeled in a Putnam-style semantics.

gesting that all cases of semantic flexibility are due to the operations of mod_l , or can be handled with tools which are unique to a Putnam-style semantics. Indeed, for any class of semantically flexible terms, it is always an open question whether to include context-sensitive parameters in their E-structure. In general, Putnam-style semantics is perfectly compatible with the presence of context-sensitive parameters in any dimension. As illustrated in our account of color terms, there can be good reasons to include context-parameters in dimensions of C-structure. Furthermore, it is often an open question whether flexibility is due to compositional linguistic processes, or to post-compositional, general pragmatic effects. In short, we should keep in mind that ‘vertical’ semantic flexibility—i.e., cases in which the computations of E-meaning fetch information from C-structure—is only one source of flexibility, albeit an important one. Having said that, I now consider some objections (Sects. 5.1–5.2), and end by comparing Putnam-style semantics with prominent philosophical theories of the nature of linguistic meaning (Sect. 5.3).

5.1 Default lexical entries?

In my Putnam-style multidimensional semantics, the expressive power of each application of mod_l is constrained by the semantic information available in the E-structure and C-structure of the relevant expression. Since ‘loosenings’ often result from replacing the E-structure of an expression with a dimension of its C-structure, and ‘enrichments’ often result from adding dimensions of the C-structure to the E-structure, mod_l can perform the basic modulation operations that truth-conditional pragmatists argue are part of our basic linguistic competence (cf. Carston 2012; Recanati 2012). In addition, the lexical limit on the expressive power of mod_l predicts many cases of unavailable enrichments. For example, the unavailable readings in cases like (18-a) and (19-a) are explained as arising from limits on lexically available information. In the default case, neither *stop* nor *garden* have ‘mowing’ as part of their E/C-structure, and neither *began* nor *nails* have ‘hammering’ as part of their E/C-structure. Hence mod_l cannot upload that information into the E-structure of the corresponding expressions. In general, each time we explain an un/available interpretation, we need to appeal not only to constraints on the operations of mod_l , but also to assumptions about the C-structure of the relevant expressions. It follows that each explanation is only as adequate as the assumptions about the conceptual content of the relevant expressions are justified.

Some philosophers might find this last feature of a Putnam-style semantics problematic. For example, when specifying the lexical entry of an artifact term, why assume that it encodes a TELIC dimension? Furthermore, even if we grant that such a dimension is represented in artifact terms, why hold that its value, for a given term, is function x rather than y ? Relatedly, isn’t there a circularity worry here: we will include information in the C-structure of e because we think it might later be used to perform certain modulations, and exclude information merely because it has no conceivable compositional effect? Obviously, we cannot just choose from the armchair which dimensions and values to include in lexical entries. In particular, we cannot just construct the content of a lexical entry to include all the information that we think could in some context be uploaded into the E-structure of host expressions.

These worries reduce to the following question: are there independent, reliable methods for determining the structure and content of lexical entries? The answer is undoubtedly ‘yes’, and this is one of the reasons why we can now begin to integrate a Putnam-style semantics into our compositional truth-conditional theories. The empirical study of concepts, now a mature field of cognitive psychology, provides several ways of testing hypotheses about the structure and content of lexical entries. In particular, there are canonical methods for generating prototypes, which allow us, among other things, to investigate the values and dimensions encoded by different classes of concepts (McRae et al. 2005). The structures generated in this way have proven to be some of the most useful variables discovered by cognitive science (Fodor 1998; Margolis and Laurence 1999; Murphy 2002; Machery 2009, 2011; Rosch 2011; Sassoon 2011). Furthermore, at least since the work of Keil (1989), there’s been substantial improvement in the experimental paradigms used to determine the dependency relations between—hence the relative centrality of—the various dimensions and features of concepts. The resulting dependency structures also constitute uniquely predictive psychological variables (Sloman et al. 1998).

Lexical atomism was a useful theoretical simplification in the development of compositional theories of meaning. However, I have argued in this paper that, if we want to deal with semantic flexibility, that assumption has to be abandoned. We can do this in part because we now have the empirical tools to seriously explore questions about the fine-grained semantic structure of words.³³ From this perspective, a current task of philosophy of language is to determine what is the best general framework to represent these structures, and how to modify our current compositional operations so that they can interact with them in appropriately constrained ways. The Putnam-style semantics presented here is one such proposal: lexical items include an E-structure and C-structure, which encodes weights and centrality relations. The compositional operations which determine the E-meaning of utterances can access C-structure, via the *mod_l* operation of the interpretation function, and are sensitive to centrality, via the ordering of the dimension and core enrichment operations available to each class of expressions.

5.2 Diachronic semantic flexibility?

Another challenge to my Putnam-style semantics comes from the more radical contextualist positions within the camp of truth-conditional pragmatists. Assume that our Putnam-style semantics can deal with basic cases of vertical semantic flexibility, and in a way that respects the observed constraints on available meanings. How then do

³³ Collaborative work between psychologists and philosophers has produced some of the most interesting results in this area. As mentioned above, Knobe et al. (2013) developed a series of experiments which show that certain terms—e.g., social role terms such as *artist*—encode a normative dimension which is distinct from the TELIC dimension. In what is arguably one of the most sophisticated empirical works on the semantic structure of function words, Pietroski et al. (2009) and Lidz et al. (2011) show that, by exploring the interaction between linguistic statements and the verification strategies used in tasks involving well-studied modalities (e.g., vision), we can distinguish which amongst various truth-conditionally equivalent semantic values for particular quantifiers has psychological reality.

we account for the general observation that, with sufficient support from the discourse context, even seemingly ‘unavailable’ meanings can be somehow induced? With sufficient effort and creativity, can we not imagine contexts in which, say, *begin* means ‘begin hammering’, and *gun* means ‘looks loosely like a gun’? If this is correct, doesn’t replacing *mod* with mod_l in our function of modulated interpretation, $[[]_M^c$, result in a model of semantic competence that is too inflexible?

In response, I should first note that, even if we accept the general observation behind this objection, this does not justify adopting free *mod*. As discussed in Sect. 2, *mod* does not just capture some general processes according to which, with sufficient (i.e., enormous!) background changes, almost any meaning change can be induced. Rather, *mod* is assumed to have the expressive power to introduce descriptive content ‘from outside’ to increase relevance in each token application, and this clearly over-generates non-available meanings, especially in the cases and contexts such as the ones we examine above. Having said that, I am not rejecting the general observation behind this objection. Indeed, I think it forces us to make the following clarification: in our accounts of vertical semantic flexibility, we abstracted away from what we can call ‘diachronic semantic flexibility’.

Diachronic flexibility includes changes to the full meanings—i.e., E/C-structure—of expressions due to things like shifting task demands, accumulation of information as conversations develop, and changing beliefs about the fine-grained idiolect used by conversation participants. Anyone who holds that compositional linguistic processes are sensitive to C-structure, will be forced to accept that diachronic semantic flexibility has a substantial impact on linguistic meaning. Indeed, the empirical study of concepts converges on the following finding: at least some of the components of the prototypes associated with descriptive terms are sensitive to task demands and background information (Barsalou 1987; Murphy 2002; Rosch 2011; Casasanto and Lupyan 2015; Machery 2015; Del Pinal 2016). For example, which perceptual prototype is associated with *lion*, in a given context, depends on the presence of information such as whether the topic is female, male or baby lions. In short, the assumption that the lexicon includes stable entries such as (22) is a simplification that we have to abandon, at least once we decide to model diachronic semantic flexibility.

Still, an advantage of adopting Putnam-style lexical entries is that they encode structural information that is crucial to understand the dynamics of diachronic semantic flexibility. For example, there is substantial evidence in support of the following claims (see, e.g., Murphy 2002; Carey 2009). A central dimension in our conceptions of natural kinds is the way in which the relevant entity comes into being, i.e., the AGENTIVE (A). The PERCEPTUAL (P) dimension is less central. This also holds of our conceptions of artifact kinds, although in this case the function or TELIC (T) also seems to be quite central. To capture these results, we can represent the centrality of each dimension by its relative position in C-structure. So consider again the entries for *lion* and *gun*, where we now represent the relative centrality of the dimensions:

$$(64) \quad [[lion]]_M^c =$$

E-structure: $\lambda x. LION(x)$
C-structure: $\langle A, C, P, T \rangle$

$$(65) \quad \llbracket \text{gun} \rrbracket_M^c =$$

E-structure: $\lambda x. \text{GUN}(x)$
C-structure: $\langle A, T, C, P \rangle$

Why is the relative centrality of dimensions crucial to model diachronic semantic flexibility? The more central a feature, the less likely it is to change as information and task demands change, hence the more diachronic stability it has (Del Pinal and Spaulding 2017). For example, one's image of a typical gun changes substantially depending on whether one is thinking of petty crimes in Chicago or a guerrilla war in Colombia. One's conception of the basic function of guns, however, is more stable across those same contexts. Similarly, one's representation of a typical lion changes depending on whether one is thinking of a female, male or baby lion. However, one's beliefs about their typical way of coming into being (i.e., by being born of lion parents) are stable across those same sub-categories.

The view that centrality is a strong determinant of diachronic flexibility—i.e., the more central is feature f in concept C , the more stable is f in the diachronic life of C —has a distinguished philosophical pedigree, even if, usually, it is not presented in these terms. Indeed, this is one of the reasons why philosophers who emphasize that one function of concepts is to provide thought with substantial cross-contextual stability, have criticized theories of concepts that include features or dimensions that have a low degree of centrality, including the original versions of prototype theory (cf., Rey 1983; Fodor 1998). Furthermore, there is substantial empirical evidence in support of this view: e.g., it has been shown that the more central a feature or dimension in concept C , the more likely it is to survive into compositions and sub-categorisations involving C (Hampton 1987, 2006).³⁴

Summing up, the radical contextualist objection forced us to acknowledge that, when invoking the notion of a default lexical entry in our accounts of vertical semantic flexibility, we abstracted away from diachronic semantic effects.³⁵ Furthermore, diachronic semantic processes are the product of cognitive processes which are not strictly part of our linguistic competence. Indeed, a full account of diachronic effects will have to tackle questions like: How do task demands affect the amount/kind of information retrieved from long term memory? How does information about the social category of our interlocutors help us decide which fine-grained idiolect to adopt? Still, promising psychological accounts of diachronic processes make use of conceptual structures that encode, at least, the information encoded in the lexical entries of our

³⁴ Even researchers who emphasize the relative instability of C-structures, or concepts in general—e.g., Barsalou (1987), and more recently, Casasanto and Lupyan (2015)—do not question, or at least provide any reason to question, the claim being made above, namely, that degree of centrality is a determinant factor of cross-contextual stability. Furthermore, various theorists have recently argued in detail that the supposed diachronic instability of even non-central features, and of C-structure in general, has been greatly exaggerated (see e.g., Prinz 2012; Machery 2015; Del Pinal 2016).

³⁵ To be clear, this does *not* affect the basic form of our account of vertical semantic flexibility. To see this, it is useful to make a distinction between 'working' and 'long-term' lexical entries. Diachronic semantic effects interact with long-term lexical entries to produce a working entry (e.g., to produce the entry used for *lion* in a context in which one is clearly talking about female lions). Strictly speaking, the terminal nodes of linguistic structures are all working lexical entries. In other words, meanings at terminal nodes should not be identified with structures stored in long-term memory.

Putnam-style semantics, such as the relative centrality of different dimensions and features.

5.3 Multidimensional Putnam-style semantics and traditional theories of linguistic meaning

I now briefly compare Putnam-style semantics with some influential theories of linguistic meaning and concepts.

Putnam-style semantics requires some degree of reconciliation between philosophical and psychological theories of lexical meaning and concepts. Now, some philosophers operate under a general distrust of the results, or of the commonly drawn theoretical implications, of the scientific study of concepts. When this position is explicitly defended, what is usually objected to is the claim that ‘prototypes’ (or richer structures such as our C-structures) are, strictly speaking, components of concepts (e.g., [Rey 1983](#); [Burge 1993](#); [Fodor 1998](#)). In those discussions, it is clear that by ‘concepts’ philosophers tend to include only representations which have an extension determining role (roughly, our E-structure). Still, that prototypes and dependency networks are important psychological variables is usually not disputed—indeed, it has been explicitly accepted by important critics of ‘psychological’ theories of concepts, such as [Rey \(1983\)](#) and [Fodor \(1998\)](#).

Importantly, our Putnam-style semantics is based on the idea that the C-structure of a given term usually does *not* determine its E-structure. In this, Putnam-style semantics respects a longstanding philosophical tradition. However, it is often also assumed that the only thing that a term contributes to complex expressions is what we here call the ‘E-structure’. One of the aims of this paper is to show that this second assumption is, at best, an idealization which it would now be useful to drop. According to our Putnam-style semantics, the ‘conceptions’ associated with terms sometimes enter into the compositional processes that determine the (literal) E-meaning of complex expressions of which they are part. If this is correct, conceptions are part of what has to be grasped to possess full linguistic competence, as Putnam originally suggested (see also [Bilgrami 1992, 1998](#); [Chomsky 2000, 2012](#)).

Indeed, theorists such as [Burge \(1993\)](#) and [Soames \(2015\)](#) hold that, for classes of words such as common nouns, grasping something like C-structure might be required to possess ‘full linguistic competence’. Now, the importance of C-structure is often illustrated by invoking the notion of a ‘canonical’ categorization or recognition procedure. For example, there is clearly some key competence missing from a mature speaker whose perceptual systems are functioning, who seems to otherwise know the meaning of *lion*, but who can nevertheless not recognise (in ‘normal’ conditions) a typical lion as a lion. However, accepting that is compatible with also holding that categorization is not, properly speaking, part of linguistic competence. In contrast, the argument presented here, according to which C-structure is involved in the *compositional* linguistic processes which determine E-structure, strongly suggests that C-structure is part of our linguistic competence. In other words, if the C-structures of expressions often determine their compositional contribution to the E-structures of

which they are part, then we have to accept that C-structures are, strictly speaking, part of linguistic meaning.³⁶

To be sure, some philosophers have argued that concepts and linguistic meaning should include something like C-structures. Good examples include the empirically informed theories of concepts defended by Prinz (2002), Weiskopf (2009), Machery (2009) and Gärdenfors (2004, 2014).³⁷ These theories are *broadly* compatible with key components of a Putnam-style semantics. However, the first two authors adopt an implementation which blocks the dynamic interaction between C-structures and the compositional determination of E-structures. For example, Weiskopf (2009) argues that E and C-structures are computed in parallel, so that each process is basically isolated from the other. In this paper, we have examined the advantages of adopting a more dynamic view of compositionality, in which the processes which determine the E-structure of complex expressions can access the C-structure of their constituents. We have also seen that this view can be implemented in a type-driven semantic theory. For these reasons, I think theorists who already accept rich informational views of the lexicon should drop the assumption of strictly parallel composition and adopt instead a dynamic interpretation function along the lines of $\llbracket \cdot \rrbracket_M^C$.

6 Conclusion

According to truth-conditional pragmatics, we can model paradigmatic cases of semantic flexibility within a compositional truth-conditional theory. To achieve that, pragmatists tend to assume that we should adopt an interpretation function that has the expressive power of free modulation. Despite some advantages, the resulting function over-generates unavailable meanings, and it is hard to see how this could be fixed by appealing to external constraints such as ‘performance’ limits. Furthermore, this impasse supports the main critics of truth-conditional pragmatics: minimalists and neo-Wittgenstenians. Minimalists can take it as evidence against the view that we should account, in our formal semantic theories, for the ‘intuitive’, flexible meanings associated with ‘what is said’ (cf. Borg 2004, 2007; Cappelen and Lepore 2005). Neo-Wittgenstenians can take it as evidence that the pursuit of formal, compositional theories of our linguistic competence is deeply misguided, or at least naively optimistic (cf. Travis 1994; Chomsky 2000).

³⁶ In our Putnam-style semantics, the conception associated with e does not determine e 's E-structure, and yet still has an important role in determining the compositional contribution of e to complex expressions. For this reason, this position has some similarities to conceptual role theories (e.g. Block 1986, 1993; Harman 1987; Greenberg and Harman 2008); however, in the implementation adopted here, all content is truth-conditional content.

³⁷ Other philosophical accounts also take something like C-structure as an integral part of meaning. For example, the notion of C-structure is one way of implementing Bilgrami (1992, 1998)'s Fregean requirement that a subject's own concepts and beliefs be used to specify their (externally determined) concepts. This allows the following possibility (which some argue is required to deal with familiar Frege-cases): If two subjects associate different C-structures with term e , then, even if we equate their E-structure for e , their corresponding (full) meanings are, strictly speaking, different. Another account that emphasises the importance of C-structure for modelling competence, but more radically than we have done here (it abandons the notion of E-structure altogether), is Rayo (2013)'s grab-bag semantics.

I have argued that, to best pursue the project of truth-conditional pragmatics, we should drop the assumption that most lexical items are atomic and one-dimensional. I proposed that we adopt a view of the general structure of the lexicon, originally defended by Putnam, according to which full competence with many lexical items requires that we grasp their dual E/C-structure. I showed that this view can be integrated into a type-driven, compositional theory of meaning, and that the resulting multidimensional framework provides us with a promising set of tools to deal with semantic flexibility. By introducing Putnam-style lexical items, and compositional operations which can interact with them in constrained ways, we can keep some of the flexibility of the modulated interpretation function, while reigning-in some of its expressive power.

This paper focused on open class terms such as common nouns and nominal modifiers. In future work, I argue that a Putnam-style multidimensional semantics can be used to deal with flexibility-related puzzles involving quantifiers, generics and propositional attitudes. Ultimately, I hope that, in addition to furthering the project of truth-conditional pragmatics—and showing in particular how it can be usefully implemented in a multidimensional semantics (cf. [McCready 2010](#); [Gutzmann 2015](#))—the system developed here also contributes to recent efforts to model how representations akin to conceptions or C-structures can play a role within the compositional semantics. Particularly important in this regard are accounts such as [McNally \(2017\)](#) and [McNally and Boleda \(2017\)](#), which try to unify formal and distributional semantics to model C-structures and their combinatorics, and [Sassoon \(2011, 2017\)](#), which provides an account of nominal comparisons using a semantics which has access to multidimensional, C-structure-like information associated with nouns.

Acknowledgements Versions of this paper were presented at the University of Beijing, ZAS Berlin, the New York Philosophy of Language Workshop, Universitat Pompeu Fabra, and the Berlin Philosophy of Language reading group. Thanks to audiences at those talks, and especially to Akeel Bilgrami, Emma Borg, Robyn Carston, Stephen Crain, Alexander Dinges, Nathan Hansen, Eleonore Neufeld, Martin Stokhoff, François Recanati, Daniel Rothschild, Uli Sauerland, Emanuel Viebahn and Julia Zakkou. Two anonymous referees for Linguistics and Philosophy provided me with excellent, constructive comments which led to substantial improvements of this paper. I am also grateful to Louise McNally for many illuminating discussions about the interface between lexical and compositional semantics, and to Barbara Partee for an insightful review of another, related paper of mine which provided me with extremely valuable material to develop my project on multidimensional semantics.

References

- Asher, N. (2011). *Lexical meaning in context: A web of words*. Cambridge: Cambridge University Press.
- Barsalou, L. 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* (pp. 101–140). Cambridge: Cambridge University Press.
- Bezuidenhout, A. (2002). Truth-conditional pragmatics. *Philosophical Perspectives Language and Mind*, 16, 105–134.
- Bilgrami, A. (1992). *Belief and Meaning*. Oxford: Basic Blackwell Ltd.
- Bilgrami, A. (1998). Precis of belief and meaning. *Philosophy and Phenomenological Research* LVIII, 1(3), 595–605.
- Block, N. (1986). Advertisement for a semantics for psychology. *Midwest Studies in Philosophy*, 10(1), 615–78.
- Block, N. (1993). Holism, hyper-analyticity and hyper-compositionality. *Philosophical Issues*, 3, 37–72.

- Borg, E. (2004). *Minimal semantics*. Oxford: Oxford University Press.
- Borg, E. (2007). Minimalism versus contextualism. In G. Preyer & G. Peter (Eds.), *Context-sensitivity and semantic minimalism: New essays on semantics and pragmatics* (pp. 546–571). Oxford: Oxford University Press.
- Burge, T. (1993). Concepts, definitions and meaning. *Metaphilosophy*, 24(4), 309–325.
- Cappelen, H., & Lepore, E. (2005). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*. Oxford: Blackwell.
- Carey, S. (2009). *The origin of concepts*. New York, NY: Oxford University Press.
- Carlson, G., & Pelletier, F. J. (2002). The average American has 2.3 children. *Journal of Semantics*, 19, 73–104.
- Carston, R. (2012). Relevance theory. In G. Russell & D. G. Fara (Eds.), *The Routledge companion to the philosophy of language Chap 1.14* (pp. 163–176). New York: Routledge.
- Casasanto, D., & Lupyan, G. (2015). All concepts are ad hoc concepts. In E. Margolis & S. Laurence (Eds.), *The conceptual mind: New directions in the study of concepts, Chap. 19* (pp. 543–566). Cambridge, MA: The MIT Press.
- Chomsky, N. (1986). *Knowledge of language: Its nature, origin, and use*. Connecticut: Praeger Publishers.
- Chomsky, N. (2000). *New horizons in the study of mind language*. Oxford: Oxford University Press.
- Chomsky, N. (2012). *The science of language: Interviews with James McGilvary*. Cambridge: Cambridge University Press.
- Coulson, S., & Fauconnier, G. (1999). Fake guns and stone lions: Conceptual blending and privative adjectives. In B. Fox, D. Jurafsky, & L. Michaels (Eds.), *Cognition and Function in Language* (pp. 143–158). Palo Alto, CA: CSLI.
- Del Pinal, G. (2015a). Dual content semantics, privative adjectives and dynamic compositionality. *Semantics and Pragmatics*, 8(7), 1–53.
- Del Pinal, G. (2015b). The structure of semantic competence: Compositionality as an innate constraint of the Faculty of Language. *Mind & Language*, 30(4), 375–413.
- Del Pinal, G. (2016). Prototypes as compositional components of concepts. *Synthese*, 193(9), 2899–2927.
- Del Pinal, G., & Reuter, K. (2017). Dual character concepts in social cognition: Commitments and the normative dimension of conceptual representation. *Cognitive Science*, 41(S3), 477–501.
- Del Pinal, G., & Spaulding, S. (2017). Conceptual centrality and implicit bias. *Mind & Language* (forthcoming).
- Fodor, J. (1998). *Concepts: Where cognitive science went wrong*. Oxford: Oxford University Press.
- Fodor, J. (2001). Language, thought and compositionality. *Mind & Language*, 16, 1–15.
- Fodor, J. (2003). *Hume variations*. Oxford: Oxford University Press.
- Fodor, J., & Lepore, E. (2002). *The compositionality papers*. Oxford: Oxford University Press.
- Gärdenfors, P. (2004). *Conceptual spaces: The geometry of thought*. Cambridge, MA: The MIT Press.
- Gärdenfors, P. (2014). *The geometry of meaning: Semantics based on conceptual spaces*. Cambridge, MA: The MIT Press.
- Greenberg, M., & Harman, G. (2008). Conceptual role semantics. In E. Lepore & B. C. Smith (Eds.), *The Oxford handbook of the philosophy of language, Chap. 14* (pp. 295–322). Oxford: Oxford University Press.
- Gutzmann, D. (2015). *Use-conditional meaning: Studies in multidimensional semantics*. Oxford: Oxford University Press.
- Hall, A. (2014). ‘free’ enrichment and the nature of pragmatic constraints. *International Review of Pragmatics*, 6(1), 1–28.
- Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, 15(1), 55–71.
- Hampton, J. A. (2006). Concepts as prototypes. *The Psychology of learning and motivation: Advances in research and theory*, 46, 79–113.
- Hampton, J. A., & Jonsson, M. L. (2012). Typicality and compositionality: The logic of combining vague concepts. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality, Chap. 18* (pp. 385–402). New York: Oxford University Press.
- Hansen, N. (2011). Color adjectives and radical contextualism. *Linguistics and Philosophy*, 34, 201–221. <https://doi.org/10.1007/s10988-011-9099-0>.
- Harman, G. (1987). (nonsolipsistic) conceptual role semantics. In E. Lepore (Ed.), *New directions in semantics* (pp. 55–81). London: Academic Press.
- Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. Oxford: Blackwell Publishers Ltd.

- Keil, F. C. (1989). *Concepts, kinds and cognitive development*. Cambridge, MA: The MIT Press.
- Kennedy, C., & McNally, L. (2010). Color, context, and compositionality. *Synthese*, 174, 79–98. <https://doi.org/10.1007/s11229-009-9685-7>.
- Knobe, J., Prasada, S., & Newman, G. E. (2013). Dual character concepts and the normative dimension of conceptual representation. *Cognition*, 127, 242–257. <https://doi.org/10.1016/j.cognition.2013.01.005>.
- Kripke, S. (1980). *Naming and necessity*. Cambridge, MA: Harvard University Press.
- Lahav, R. (1989). Against compositionality: The case of adjectives. *Philosophical Studies*, 57, 261–279.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago, IL: The University of Chicago Press.
- Laserson, P. (1999). Pragmatic halos. *Language*, 75(3), 522–551.
- Laserson, P. (2012). Contextualism and compositionality. *Linguistics and Philosophy*, 35, 171–189.
- Leslie, S. J. (2015). “Hillary Clinton is the only man in the Obama administration”: Dual character concepts, generics and gender. *Analytic Philosophy*, 56(2), 111–141.
- Lewis, K. (2014). Review of truth conditional pragmatics by Francois Recanati. *Mind*, 123(492), 1234–1238.
- Lidz, J., Pietroski, P., Halberda, J., & Hunter, T. (2011). Interface transparency and the psychosemantics of ‘most’. *Natural Language Semantics*, 19, 227–256.
- Machery, E. (2009). *Doing without concepts*. Cambridge, MA: The MIT Press.
- Machery, E. (2011). Concepts: A tutorial. In R. Belohlavek & G. J. Klir (Eds.), *Concepts and fuzzy logic Chap. 2* (pp. 13–44). Cambridge, MA: The MIT Press.
- Machery, E. (2015). By default: Concepts are accessed in a context-independent manner. In E. Margolis & S. Laurence (Eds.), *The conceptual mind: New directions in the study of concepts, Chap. 20* (pp. 567–588). Cambridge, MA: The MIT Press.
- Margolis, E., & Laurence, S. (1999). Concepts and cognitive science. In E. Margolis & S. Laurence (Eds.), *Concepts: Core readings, Chap. 1* (pp. 3–82). Cambridge, MA: The MIT Press.
- Margolis, E., & Laurence, S. (2007). *Creations of the mind: Theories of artifacts and their representation*. Oxford: Oxford University Press.
- McCready, E. (2010). Varieties of conventional implicature. *Semantics and Pragmatics*, 3(8), 1–57.
- McCready, E. (2015). Against lexical self-reference. *Linguistic Inquiry*, 46(4), 742–754.
- McCready, E., & Ogata, N. (2007). Adjectives, stereotypicality and comparison. *Natural Language Semantics*, 15, 35–63.
- McNally, L. (2017). Kinds, descriptions of kinds, concepts, and distributions. In K. Balogh & W. Petersen (Eds.), *Bridging formal and conceptual semantics. Selected Papers of BRIDGE-14* (pp. 39–61). DUP: Düsseldorf.
- McNally, L., & Boleda, G. (2017). Conceptual vs. referential affordance in concept composition. In Y. Winter & J. A. Hampton (Eds.), *Compositionality and concepts in linguistics and psychology* (pp. 245–268). Berlin: Springer.
- McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavioral Research Methods, Instruments, & Computers*, 37(4), 547–559.
- Moravcsik, J. (1998). *Meaning, creativity, and the partial inscrutability of the human mind*. California: CSLI Publications.
- Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: The MIT Press.
- Pagin, P. (2014). Pragmatic enrichment as coherence raising. *Philosophical Studies*, 168, 59–100.
- Pagin, P., & Pelletier, J. (2007). Content, context, and composition. In G. Preyer & G. Peter (Eds.), *Context-sensitivity and semantic minimalism: New essays on semantics and pragmatics* (pp. 25–62). Oxford: Oxford University Press.
- Pagin, P., & Westerstahl, D. (2010a). Compositionality I: Definitions and variants. *Philosophy Compass*, 5, 250–264.
- Pagin, P., & Westerstahl, D. (2010b). Compositionality II: Arguments and problems. *Philosophy Compass*, 5, 265–282.
- Pagin, P., & Westerstahl, D. (2010c). Pure quotation and general compositionality. *Linguistics & Philosophy*, 33, 381–415.
- Partee, B. (2007). Compositionality and coercion in semantics: The dynamics of adjective meaning. In G. Bouma, I. Kraemer, & J. Zwirts (Eds.), *Proceedings of the colloquium “cognitive foundations of interpretation”* (pp. 145–161). Amsterdam: Royal Netherlands Academy of Arts and Sciences.
- Partee, B. (2010). Privative: Subjectives plus coercion? In R. Baule, U. Reyle, & T. Zimmerman (Eds.), *Presupposition and discourse: Essays offered to Hans Kamp* (pp. 273–285). Amsterdam: Elsevier.

- Pietroski, P., Lidz, J., Hunter, T., & Halberda, J. (2009). The meaning of 'most': Semantics, numerosity and psychology. *Mind & Language*, 24(5), 554–585.
- Portner, P. (2007). Instructions for interpretation as separate performatives. In K. Schwabe & S. Winkler (Eds.), *On information structure, meaning and form* (pp. 407–426). Amsterdam: Benjamins.
- Potts, C. (2005). *The logic of conventional implicatures*. Oxford: Oxford University Press.
- Prinz, J. (2002). *Furnishing the mind*. Cambridge, MA: MIT Press.
- Prinz, J. (2012). Regaining composure: A defense of prototype compositionality. In M. Werning, W. Hinzen, & E. Machery (Eds.), *Oxford handbook of compositionality*, Chap. 21 (pp. 437–453). Oxford: Oxford University Press.
- Pupa, F. (2015). Impossible interpretations, impossible demands. *Linguistics and Philosophy*, 38, 269–287.
- Pustejovsky, J. (1995). *The generative Lexicon*. Cambridge, MA: The MIT Press.
- Pustejovsky, J. (2012). Co-compositionality in grammar. In M. Werning, W. Hinzen, & E. Machery (Eds.), *Oxford handbook of compositionality*, Chap. 17 (pp. 371–384). Oxford: Oxford University Press.
- Putnam, H. (1970). Is semantics possible? *Metaphilosophy*, 1(3), 187–201. <https://doi.org/10.1111/j.1467-9973.1970.tb00602.x>.
- Putnam, H. (1975). The meaning of 'meaning'. *Minnesota Studies in the Philosophy of Science*, 7, 131–193.
- Putnam, H. (1988). *Representation and reality*. Cambridge, MA: The MIT Press.
- Rayo, A. (2013). A plea for semantic localism. *Nous*, 47(4), 647–679.
- Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
- Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Oxford University Press.
- Recanati, F. (2012). Pragmatic enrichment. In G. Russell & D. G. Fara (Eds.), *The Routledge companion to the philosophy of language*, Chap. 1.6. New York: Routledge.
- Rey, G. (1983). Concepts and stereotypes. *Cognition*, 15(1), 237–262.
- Rosch, E. H. (2011). Slow lettuce: Categories, concepts, fuzzy sets, and logical deduction. In R. Belohlavek & G. J. Klir (Eds.), *Concepts and fuzzy logic*, Chap. 4. Cambridge, MA: The MIT Press.
- Rothschild, D., & Segal, G. (2009). Indexical predicates. *Mind & Language*, 24(4), 467–493. <https://doi.org/10.1111/j.1468-0017.2009.01371.x>.
- Sassoon, G. W. (2011). Adjectival vs. nominal categorization processes. *Belgian Journal of Linguistics*, 25, 104–147. <https://doi.org/10.1075/bjl.25.06sas>.
- Sassoon, G. W. (2017). Comparisons of nominal degrees. *Language*, 93(1), 153–219. <https://doi.org/10.1135/lan.2017.0005>.
- Sloman, S. A., Love, B. C., & Ahn, W. K. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, 22(2), 189–228.
- Soames, S. (2015). *Rethinking language, mind, and meaning*. Princeton: Princeton University Press.
- Stanley, J. (2002). Making it articulated. *Mind & Language*, 17(1&2), 149–168.
- Stanley, J. (2007). *Language in context*. Oxford: Oxford University Press.
- Szabó, Z. (2010). The determination of content. *Philosophical Studies*, 148(2), 253–272. <https://doi.org/10.1007/s11098-008-9296-z>.
- Travis, C. (1994). On constraints on generality. *Proceedings of the Aristotelian Society*, 44, 165–188.
- Vikner, C., & Jensen, P. A. (2002). A semantic analysis of the English genitive: Interaction of lexical and formal semantics. *Studia Linguistica*, 56(2), 191–226. <https://doi.org/10.1111/1467-9582.00092>.
- Weiskopf, D. (2009). Atomism, pluralism, and conceptual content. *Philosophy and Phenomenological Research*, LXXIX(1), 131–163. <https://doi.org/10.1111/j.1933-1592.2009.00269.x>. <http://www3.interscience.wiley.com/journal/122464216/abstract>.