Compositionality and sandbag semantics

Elmar Unnsteinsson

Postprint. Please cite published version:
Synthese (2014) 191(14): 3329–3350
doi:10.1007/s11229-014-0449-7

Abstract

It is a common view that radical contextualism about linguistic meaning is incompatible with a compositional explanation of linguistic comprehension. Recently, some philosophers of language have proposed theories of 'pragmatic' compositionality challenging this assumption. This paper takes a close look at a prominent proposal of this kind due to François Recanati. The objective is to give a plausible formulation of the view. The major results are threefold. First, a basic distinction that contextualists make between mandatory and optional pragmatic processes needs to be revised. Second, the pragmatic theory can withstand a Davidsonian objection only by rejecting the importance of a distinction between primitive and non-primitive semantic items. Thirdly, however, the theory is now open to a worry about how it should be understood: either the theory consists in a very broad functionalist generalization about communication, which makes it explanatorily inert, or it boils down to a highly particularist view about linguistic meaning. Finally, I argue that Recanati's notion of 'occasion meaning' is problematic and suggest replacing it with the notion of speaker meaning, which is explanatorily more basic.

> *[Funes] invented a numbering system original with himself [...]. Instead of seven thousand thirteen (7013), he would say, for instance, 'Máximo Pérez'; instead of seven thousand fourteen (7014), 'the railroad'; other numbers were 'Luis Melián Lafinur,' 'Olimar,' 'sulfur,' 'clubs,' 'the whale,' 'gas,' 'a stewpot,' 'Napoleon,' 'Agustín de Vedia.' [...] I tried to explain to Funes that his rhapsody of unconnected words was exactly the opposite of a number system. [...] Funes either could not or would not understand me.*
>
> (Jorge Luis Borges – Funes, His Memory)

*(✉) elmar.geir@gmail.com

*I want to thank Daniel Harris, Stephen Neale, Ben Phillips, Guillermo Del Pinal, participants in the 9th BPPA Masterclass with Recanati in London in 2011, and many anonymous referees, for all their comments and discussions. Special thanks to François Recanati for helpful conversations and for his comments on an early version of the paper.

1 Introduction

Recently a number of theorists have argued for a novel conception of the compositionality of meanings in natural languages (cf. Recanati 2010; Westerståhl 2012; Pagin 2005; Pagin and Pelletier 2007). We can call this 'pragmatic' compositionality. On this contextualist view, natural languages exhibit semantic flexibility in that the same expression may assume different meanings on different occasions without this being simply a matter of polysemy. This has been thought by many to involve a necessary violation of compositionality (cf. Fodor 2003). François Recanati is one prominent defender of pragmatic compositionality and argues, in his recent book *Truth-Conditional Pragmatics*, for revisions in the way we think about compositional rules so that they can accommodate not only so-called *standing* meanings of expressions but also their *occasion* meanings, i.e. the meaning the expression token has in some specific context.

In this paper I will explain what pragmatic compositionality amounts to. Many interesting issues are raised by this idea, and in this paper I focus on its proper formulation. First (§2), I explain why pragmatists want to abandon the classical notion of composition. Secondly (§3), I show that Recanati's distinction between mandatory and optional pragmatic processes needs to be reformulated. Then, in §4, I develop a Davidsonian objection to pragmatic compositionality, which depends on making a distinction between the finitely many semantic primitives and the indefinitely many compositional non-primitives of the theory. The objection falls short and this is taken to show that the contextualist should reject the distinction itself.

After these clarifications there remain, however, a few problems with the pragmatic theory. Taken one way, the theory postulates a single function which is supposed to form part of an explanation of successful linguistic communication. One ought to worry about whether this really provides a substantive explanation (§5.1). Taken another way, the theory only makes claims about *particular* contexts of utterance and gives up on the possibility of useful generalizations (§5.2). These are the most plausible interpretations and I don't see any middle way opening up. These are not devastating objections and perhaps they just point in the direction for more research on pragmatic compositionality. However, as I argue in §6, it seems more plausible to reject the problematic notion of 'occasion meanings' and adopt a more traditional view of semantics where composition is driven by the meanings of expression types. Adding speaker meaning into the mix provides all the alleged explanatory benefits of occasion meanings without adding idle wheels to the theory. Thus I give a brief sketch of a view that makes do with the meanings of linguistic expression types and speaker meanings.

2 The case for pragmatic compositionality

A language L is compositional if the meaning of a complex expression Σ in L is determined by the meanings of its immediate constituent expressions $\varphi_1, \ldots \varphi_n$ and principles determined by their immediate syntactic mode of combination. If the immediate constituents are simple they are directly related to their own lexical meanings and there is a finite number of such meanings (Recanati 2010: 28). Syntactic rules of combination are also finitely many and thus the meaning of every well-formed compound expression is ultimately determined by a finite number of lexical meanings and compositional principles. Since a syntactic rule can be applied recursively this will generate a potential infinity of complex expressions in L. Many theorists consider this idea to be an essential part of any plausible account of the so-called productivity of language. Seeing that natural language contains a potential infinity of compound expressions (because of recursion), it is impossible that people learn their meanings directly one by one—there must be some computational process involved. Recanati (2010) proposes two rules to mark this distinction.

Lexical Rule: $I(\alpha) = m$

Compositional Rule: $I(\alpha^{\star}\beta) = f(I(\alpha), I(\beta))$

According to the lexical rule, the interpretation (i.e. content) of the expression α is directly associated with a lexical meaning m. According to the compositional rule, the interpretation of the compound expression $\alpha^{\star}\beta$ ('$\star$' representing their mode of combination) is the value of the function f for which the interpretation of α and the interpretation of β are arguments. Let's call this picture *simple compositionality*: if these are the right rules then the meanings of compound expressions cannot depend on anything other than the meanings of their immediate constituents and their mode of combination. Thus we derive a plausible explanation of how language users can understand a potential infinity of sentences from only being acquainted with a finite number of simple expressions and a finite number of compositional principles.

Recanati argues, however, that the picture is too simplistic. We must somehow accommodate the phenomenon of 'semantic flexibility.' This is the apparent fact that the meaning of a simple expression, and thereby the meaning of the immediate constituent expressions of the sentence, can vary considerably from one occasion to another. A nice example is the verb 'cut.'

(1) Mia cut the grass.

(2) Mia cut the cake.

Following Searle (1980), Recanati notes that what constitutes cutting grass is quite different from what constitutes cutting cake. Thus when (1) is uttered, 'cut' takes

on the meaning of, say, 'cut *with a lawnmower*' rather than 'cut *with a cake knife*.' But in the case of (2) 'cut' seems to take on something like the latter meaning rather than the former. Thus 'cut' can take on different meanings on different occasions of utterance. In this particular example, the variation seems partly due to the fact that 'cut' takes different complements in (1) and (2). However, this cannot be the whole story since, for one thing, (1) *can* certainly be uttered in a context in which Mia, right after cutting the cake, picks up a handful of grass and cuts it to pieces with the cake knife. Furthermore, there is an abundance of similar counterexamples to simple compositionality in the literature—and no need to rehearse them here.[1]

If this is on the right track two different conceptions of semantics seem to be possible. First, there is simple compositionality or *brick-wall semantics*: the constituent expression φ makes the same semantic contribution to any sentence Σ in which φ occurs. The relation between bricks and brick walls is analogous. Second, there is what I will call *sandbag semantics*: the semantic contribution of φ to any Σ varies from occasion to occasion. The variation can be a function of many different factors (more anon). Similarly, a particular sandbag takes on different *shapes* as a function of the whole of which it is a part (cf. Recanati 2010: 31; Cohen 1986: 223).

If a sandbag semanticist wants to retain the apparent advantages of a compositional language, she must, according to Recanati, make a distinction between standing meaning and occasion meaning. The former is the abstract meaning an expression has in virtue of linguistic conventions. The latter is the specific meaning that an expression token assumes on particular occasions of utterance. The standing meaning of 'cut' as it occurs in (1) and (2) is one and the same while the verb seems to take on different *occasion* meanings. The occasion meaning varies as a function of the context: here, the sentences are uttered either in *grass contexts* or *cake contexts*. The occasion meaning of 'cut,' Recanati argues, is thus similar to the occasion meaning of indexicals, demonstratives and other context-sensitive expressions (2010: 37). The standing meaning of 'I' is something like the rule that it should refer to the speaker, whereas the occasion meaning of a token of 'I' is the specific speaker who uttered 'I' on the occasion in question. The rule is the expression's Kaplanian character, while a particular speaker will be its content. According to Recanati, this calls for a revision of the lexical rule.[2]

Revised Lexical Rule: $I_c(\alpha) = f(c)$

On this rule, the interpretation of the constituent expression α in context c is the content given by the function f, f being determined by the standing meaning of α, when f takes context c as its argument. Whenever α is context-invariant f will

[1] On problems related to compound nominals, see (e.g.) Downing (1977) and Weiskopf (2007). On adjectival modification, see (e.g.) Travis (1997). Genitive constructions are discussed in §6 below.

[2] I have changed the formulation of these rules so that the context-subscripts follow the function letters rather than the parentheses. I thank an anonymous reviewer for this suggestion.

just be a constant function and deliver m as before, regardless of the value of c. For example, 'the grass' as it occurs in (1) simply receives its literal meaning *the grass* as an interpretation.

Although 'cut' and 'I' (for example) are *similar* they are also importantly different. Indexicals, personal pronouns and demonstratives are like free variables in that they must be assigned values by the context, otherwise nothing truth conditional can be retrieved by the hearer. One doesn't know the truth condition of 'She's ready,' on an occasion, without knowing *who's ready* (and for what). Assigning a person as semantic value to 'she' is *mandatory* in this sense. Recanati calls this 'saturation.' And then there is also the *optional* process of 'modulation.' Varying the contextual meaning of 'cut' from occasion to occasion is an instance of modulation and should not be explained in terms the distinction between character and content. It should, rather, be explained in terms of the distinction between literal and non-literal meanings of expressions: such a process is 'optional' in the sense that no modulation need occur when one speaks literally.[3] But sometimes the context will force a non-literal interpretation for an expression like 'cut' (clearly something else is going on with indexicals) (2010: 42). In the following section we will see why this account of the mandatory/optional distinction doesn't work and calls for revision.

Recanati points out that modulation is quite widespread in natural language. He gives another example (2010: 41–46; 2004: 34–36),

(3) The city is asleep.

Since (3) cannot be taken literally, we are to assume that 'the city' is meant in the metonymic sense of *the inhabitants of the city*.[4] Recanati proposes that the modulation function *mod* takes expressions e and contexts c as arguments and delivers a salient pragmatic function g. In this case the standing meaning m of 'the city' is mapped on to a distinct meaning $g_{513}(m) =$ *the inhabitants of the city*. Thus g_{513} is a contextually salient metonymic function from city-meanings to inhabitant-of-city-meanings. Obviously, then, we need to revise the compositional rule (but the lexical rule is unaffected).

Revised Compositional Rule:
$I_c(\alpha^\star\beta) = f(c) = f(mod(\alpha,\, c_1)(I_{c^1}(\alpha)),\, mod(\beta,\, c_2)(I_{c^2}(\beta)) = f(g_1(I_{c^1}(\alpha)),\, g_2(I_{c^2}(\beta)))$

Thus the interpretation of the compound $\alpha^\star\beta$ is a function of the contextually modulated (context c_1 and expression α supply modulation function g_1) part-meanings and their mode of combination. It will be important for our argument later that

[3] Recanati acknowledges that it is hard to pinpoint the difference between optional and mandatory processes (2010: 20–22).

[4] As Recanati points out (3) can also be interpreted such that 'asleep' has the metaphorical meaning of 'quiet' or 'silent.'

this framework makes ineliminable use of *sub-contexts*. That is, when sentence Σ is uttered in context C the immediate constituent expressions $\varphi_1, \ldots \varphi_n$ in Σ are in one-to-one correspondence with sub-contexts $c_1, \ldots c_n \in C$. Furthermore, the meaning of Σ itself may be modulated more globally: after the part-modulation has occurred some sort of whole-modulation may also occur. Thus modulation applies to all expressions, also whole sentences (2010: 43–46).

Does this account show that there is such a thing as compositional pragmatics? Recanati acknowledges that it is not exactly the same thing as given by brick-wall semantics. But he thinks there is no essential distinction here. Sandbag compositionality provides all the benefits (and more) of compositional semantics by applying these revised rules at the level of occasion meaning. Indeed, it may seem, we have a more plausible picture of communication—one in which the gap between mere standing meaning and occasion meaning, i.e. the *correct* interpretation of an utterance, is finally bridged.

3 Saturation and modulation

As explained above, sandbag semantics rests on a distinction between two types of interpretive processes. First there is the mandatory process of saturation. Saturation applies to tokens of expressions like indexicals and demonstratives. Such token expressions must be assigned objects as contents in the context for the utterance in question to be meaningful. For this reason, it is thought, saturation is mandatory. Modulation, however, is optional because it is a process that maps literal meanings of expressions onto the non-literal meanings that they acquire in particular contexts. Modulation applies globally since any expression can be intended ironically or metaphorically, but it's not mandatory. No particular expression, or set of expressions, is such that their literal meaning in the language must be modulated so that the utterance in question is meaningful.

Now, it's unfortunate that this distinction is usually explained in terms of optionality. Recanati (2010: 20–22) has also worried about this but for different reasons. In fact, there is only a superficial sense in which saturation must occur but modulation is possibly avoided. On the contextualist view, modulation is really a global feature of every compositional process. This is because the *mod* function must always be playing its crucial role, both in literal and non-literal contexts. In this sense, modulation is a *mandatory* process and cannot be distinguished from saturation in terms of optionality.

The argument for this claim is based on three points about linguistic interpretation. As far as I can tell, Recanati has endorsed all three of them at one point or another. First, so-called 'literal' interpretation is not privileged and doesn't automatically come first in the order of interpretation (cf. Recanati 1995, 2004: 28–32). Secondly, and relatedly, the process by which modulation maps literal meanings

onto non-literal meanings is not different in kind from the process by which modulation maps literal meanings onto themselves. The latter is what Recanati calls zero-modulation and the value of *mod*, in that case, will be the identity function (2010: 45). Thus *mod* is always operative although sometimes, in literal contexts, the pragmatic function g will just be the identity function. There are only two differences between an occasion on which the meaning of sentence type Σ is modulated and another occasion on which Σ is zero-modulated: (i) in one case the input to g is not identical to the output but in the other the two are identical, and (ii) the context c must be non-literal in one case and literal in the other. But zero-modulation, since it models 'literal' interpretation and 'literal' contexts, is in no way privileged as compared to non-literal modulation or non-literal contexts in general. Thirdly, and lastly, let's assume that all linguistic interpretation is basically intention-recognition where the hearer's job is to take the speaker's utterance as evidence on the basis of which she ascribes a communicative intention to the speaker (cf. Recanati 2001: 83, 2004). If positing a process of modulation is to serve any purpose in the theory of interpretation it must model a process of this kind, i.e. where the utterance is explained by the fact that the speaker has a communicative intention.[5] And again, from this standpoint, there is no fundamental distinction between zero-modulation and non-literal modulation. The point being that even when the final interpretation happens to be a 'literal' one, where the value of *mod* is the identity function, the hearer must engage in the same type of process of interpretation. The hearer, for example, must assume that the speaker is *not* speaking ironically or metaphorically to reach the zero-modulated interpretation. And this assumption is not privileged over the assumption that the speaker is being ironical (cf. Searle 1978; Recanati 1995, 2004).

I take this to show that the *mod* function is, on the contextualist view, a crucial part of every linguistic interpretation that needs to be explained compositionally. It doesn't merely come into play when the context is not a literal one: it's a mandatory process. Modulation via identity is still modulation, just like multiplying by one is still multiplication.

4 Davidson and the finite base of semantic primitives

Let's suppose, as Recanati seems to do, that compositionality is an important part of any explanation of linguistic comprehension.[6] In its absence, explaining how humans can understand an indefinite number of novel sentences is not easy. On one influential version of this view (Davidson 1965) there must be a *finite base of*

[5] I would like to call this process 'inference to the best explanation' but Recanati would disagree (cf. Recanati 2002). See §5.1 below.

[6] Schiffer (1987: 179–209) thinks this is doubtful. See also Schiffer (2003: 160–168).

semantic primitives and a finite number of compositional principles to generate this potential infinity of well-formed sentences.[7] Davidson claimed that if a language doesn't have such a finite base of primitives it will simply be unlearnable.[8] Recanati does seem to presume this type of Davidsonian picture. He writes that it is in principle possible for a finite mind to know the meanings of all the simple expressions (and the number of readings which any ambiguous expression may have) since these are finitely many. He also states that any such *direct* knowledge of the meanings of all the *compound* expressions of a recursive language would be impossible (Recanati 2010: 27–28). Below I will show, however, that pragmatic compositionality is not compatible with the traditional Davidsonian framework. More specifically, the idea that there is a *finite* base of primitives needs to be dropped.

If there is a distinction, as outlined above, between the standing meaning and the occasion meaning of any expression then every expression can take on a very broad range of meanings. The occasion meaning of a token expression is simply the correct interpretation of that token in that context. These meanings are, then, potentially infinite in number. Further, as Recanati points out, we should be interested in the composition of occasion meanings, because these are the meanings expressions will assume in particular contexts—standing meaning does *not* vary with context.[9] Occasion meaning is determined by the context of utterance, represented by *c* in the revised compositional principle. This both applies to sentences and their constituent expressions, such that every constituent has a corresponding sub-context. Let's focus, for now, only on the subset of occasion meanings that has undergone non-literal modulation of some sort. If types of contexts are to determine modulated occasion meanings for token expression they must be at least as many as those meanings themselves. Otherwise the context type couldn't fully determine the meaning in question. Nothing less fine grained than a *cake sub-context* will determine the desired occasion meaning for 'cut' as it occurs in sentence (2) above. This suggests that the explanatory ambitions of the sandbag semanticist requires a commitment to an indefinite number of possible context types: there is no upper bound on the potential ways in which objects may be cut, to use only one example.

But how exactly are these occasion meanings known by speakers? On the Davidsonian picture, we are only allowed to countenance two categorical options: either occasion meanings are (i) part of the finite base of semantic primitives or they are (ii) derivative in some sense, i.e. arrived at by generative principles on the basis of the primitives. This presents us with a dilemma, because neither (i) nor (ii)

[7] To see that this way of thinking about compositionality is still alive and kicking, see Lepore and Ludwig (2011).

[8] Davidson (1965: 8–9).

[9] In terms of Pagin's (2005: 303–305) helpful distinction: standing meaning is context-*sensitive* because its *output* is different in different contexts, but it's not context-*dependent* since it itself remains the same in all contexts. Only the output or semantic value of the function determined by standing meaning is context-dependent.

can be correct. Let's start with the latter option. The revised compositional rule says that for any context *c* and expression α there will be a function *g* that delivers the correct occasion meaning of α in *c*. This can only be compositional if the inputs, i.e. α and *c*, are in fact primitive semantic items the combination of which delivers a complex meaning. For example, the expression 'cut' in (2) and its *cake sub-context* should be primitives the composition of which delivers the occasion meaning, say, *to cut in the manner of a cake*. The standing meaning of 'cut' poses no problem, Davidsonian compositionality was already committed to it being in the semantic base. But the fact that the context has the property of being about cakes rather than something else cannot belong in the base. First, since there is an infinite number of potential modulation-determining contextual properties of this sort, it would make the base infinite. This objection doesn't apply to traditionally Kaplanian accounts of indexicals, as they define contexts as *n*-tuples of very specific features that go into determining contents for a closed set of linguistic expressions. Such *n*-tuples will at least include a speaker, addressee, time, and place. Sandbag compositionality requires complete liberality in contextual properties and, since *mod* applies globally, no restriction on expression types. Second, such properties of contexts are not linguistic expressions with meanings. Contexts do not say anything or have any associated conventional meanings and thus it's entirely unclear what it means to *compose* a meaning and a context, in this sense. Which syntactic mode of combination could be involved here? Thirdly, this would imply that every occurrence of a simple linguistic expression φ_1 is in fact an occurrence of a compound expression Σ, namely φ_1 concatenated with an associated 'contextual expression' φ_2.[10]

This seems to leave us only with the first option: (i) occasion meanings are part of the lexicon just like standing meanings and belong themselves in the base of semantic primitives. But here the same problem just becomes more pressing. Exactly how many primitive occasion meanings are there? Well, at least for indexicals and demonstratives there is no end to the number of possible content-determining context types. The standing meaning of 'that,' for instance, determines a function which delivers a demonstrated/intended content in any context *c*. But there doesn't seem to be any rule or convention that can guide interpreters from that standing meaning to the specific content intended on some occasion. It's IBE, so any method that brings the right result is good enough. Worse still, the number of occasion meanings is potentially infinite and cannot belong in a finite base of primitives.

Quite generally, the word 'context' doesn't appear to be a sortal akin to *house*, *word* or *duck*. Rather, it expresses something without any clear criteria of countability, just like *yellow*, *striped* or *heavy*. The question 'How many *yellow* things

[10]Note that the last two objections, need to be addressed even by Kaplanian theories of indexicals. By which syntactic principle exactly, for instance, does one compose an addressee and a second-person pronoun? But since such theories are more restricted in their scope, they stand a better chance of giving credible answers.

are there in the world?' doesn't have any clear answer. The average adult knows around 40,000 simple expressions but it's unclear how many expressions-in-context, meanings-in-context, or context types he or she knows. Thus contexts types must be indefinitely many and speakers, on the Recanati–Davidson model we are examining, must learn an indefinite number of semantic primitives directly. That is, speakers must learn all the contextually possible modulations of 'cut' or 'the city' one by one. This seems to be implausible and conflicts with our original assumption about the explanatory importance of compositionality. Part of that assumption was the claim that there is no way for a speaker to learn the meanings of all the semantic primitives if they are indefinitely many.

5 Pragmatic compositionality and pragmatic functions

Fortunately, there is a way out for the sandbagger. We need only drop a part of the Davidsonian picture, specifically the requirement that there be a finite base of primitive semantic *items*, which need to be learned one by one. Learnability does not necessarily require a finite domain of primitives. It only requires that the compositional theory itself can be finitely stated. Thus the statement of the theory can easily contain variables ranging over an infinite domain of values, e.g. sequences of objects. The fact that an expression can have have an infinite number of semantic values is no bar to a compositional theory with respect to that expression. And we want our theory to cover natural language expressions that behave like variables, e.g. demonstratives and pronouns. Most importantly for our purposes, the number of primitive or non-primitive linguistic items can fall were it may and the distinction itself loses its importance. The only constraint that remains is that learning a language cannot consist in learning a *new* primitive meaning for every new compound expression.

As I understand Recanati's view, it fits this picture very well. One theoretical primitive in his account is a uniform modulation function *mod* which takes expressions e in some language and contexts in domain C as arguments and delivers a function in domain G of *pragmatic* functions. That is mod: $E \times C \rightarrow G$. The interpretation function maps expressions in E and contexts onto meanings in domain M, i.e. I: $E \times C \rightarrow M$. The pragmatic functions $g \in G$ are, then, just mappings within the domain of meanings, g: $M \rightarrow M$. If this is right, the sandbagger's theory can be stated by (M).

(M) $mod(e, c)(I_c(e)) = g(m) = m'$

As noted in §2, the expression e can be either a simple or a complex expression. If e is complex g will take as argument the content m which results from composing the modulated meanings of the immediate constituents of e. This composed modulated

content m is mapped by g to the modulated meaning m'. Or, in cases of zero-modulation, g maps the composed modulated content m onto itself. Thus there is a sense in which modulation is not optional: *mod* always applies to all expressions in the process of interpretation (cf. §3). The only thing that's really optional is whether or not the value of *mod* happens to be the identity function in a particular case.[11] This is important for the arguments below.

There are two types of worries that I want to raise against principles like (M) in this section. Taken one way, I worry that the *mod* function is too general to be part of an *explanation* of linguistic communication. It seems rather to restate what theorists are aiming to explain in the first place. Taken another way, I worry that the function is too *particular* to specific contexts to provide a generalizable result. Although these worries are not conclusive objections to sandbag semantics, I take them to suggest that an alternative approach should be explored.

5.1 The argument from generality

The two functions of (M), *mod* and *g*, are powerful creatures. So powerful that it seems unlikely that a single function like *mod*, that ultimately delivers *the right interpretation* of any linguistic utterance, explains the capacity of humans to interpret linguistic utterances. Indeed, *mod* simply takes every expression-context pair onto the meaning intended by the speaker. Although it is trivially true that there *exists* a mapping from utterances U to what the speaker intended by U it is unclear how pointing this out can help explain the hearer's ability to interpret U correctly. But Recanati claims that modulation is required for composition to "deliver the correct interpretation for complex expressions" (2010: 44) and the fact that interpreters use or think in accord with *mod* is supposed to partly explain this. I think, however, that the true explanation must posit some mechanism of inference to the best explanation (IBE). Surely, such a mechanism can be specified in functionalist terms. But of course it is no explanation of the human capacity to perform IBE that humans can compute a function from every x to the explanation of x.[12] The situation with *mod* is similar. If linguistic comprehension is basically intention-recognition, which is IBE-based, it isn't explained, in part or in whole, by positing a function like *mod* (more on this in §6).

[11] Arguably, it is more accurate to assign meaning-context pairs as inputs to *mod*, rather than expression-context-pairs. First, as Recanati (2010: 41) stresses, modulation is determined by an incompatibility between *meanings* and contexts. Second, since higher-level modulations such as *irony* must take outputs of lower-level modulation as inputs, and those outputs are definitely *meanings*, it seems like *mod* must take meanings rather than expressions as inputs. It is, for example, the metonymic occasion meaning of (3) that would be intended ironically, not its invariant standing meaning. But, of course, Recanati's own formulation should be followed here.

[12] Fodor (1981: 11–14) makes a similar point about functionalist explanations in general. But of course, as Fodor points out, this doesn't show that *all* functionalist explanations are vacuous. There must be constraints and these can be given, e.g., by the notion of Turing-computability.

The point is not that 'functionalist explanations' are never any good. For example, according to Bradley (2014), one thing that makes such explanations informative is that they rule out other explanations. A true functional explanation in psychology posits some intrinsic state of the agent, which excludes extrinsic or merely accidental explanations. And surely this ought to be an upshot of positing *mod* to explain linguistic interpretation. It would exclude, say, an explanation according to which an agent understood the metonymy in an utterance of (3) only because some brain disorder makes her think that everyone is always referring to the city's inhabitants when 'the city' is uttered. Even this much, however, isn't quite clear. The problem is that pragmatic compositionality casts the explanatory net too wide. The agent who thus accidentally understood (3) is still one who, we can assume, took account of the fact that a sentence with a specific standing meaning was uttered in a specific context. And she got it right. How exactly does she not count as having applied *mod* correctly in this case? On the contextualist view, there is no reason to suppose that she didn't. Since an accidental understander can also act in accord with some function from expression-context pairs to some mapping between two meanings, it's unclear how the contextualist theory helps us distinguish cases of actual interpretive success from cases that are not part of what theorists are trying to explain.

The main point, however, is that the explanatory ambitions of sandbag semantics are bloated. We don't have a handle on *mod* except insofar as it is the function that delivers what the hearer came up with as an interpretation. It is a statement of the explanandum, not explanans.

The sandbagger has a reasonable answer to this charge. She can concede that *mod* is too intelligent a creature on its own, but insist that the point of the theory is to break the function down into its subpersonal proper parts. And positing the pragmatic function g is just the first step in this direction. In the final analysis, the sandbagger will have discovered a multitude of simple and small functions designed to perform specific tasks in the interpretive process: metonymy, metaphor, irony, literalness, etc. This methodology is reminiscent of Dennett's (1991) well-known 'homuncular' functionalism designed to explain consciousness, where each functional posit is relatively stupid.

Unfortunately, however, the sandbagger's theory is a far cry from reaching this goal. Each function, *mod* or g_n (for all n), is too unconstrained for this answer to carry conviction. We need utterly dumb, less-than-whole-mind operations to be modeled by these functions. First, the way in which *mod* gives rise to its particular g_n in context is an IBE. Second, any pragmatic function itself, such as the metonymy function, seems to require the resources of a full-blown mind to compute. Pragmatic functions are defined as mappings from meanings to meanings. And as we have seen, just about anything in the universe is potentially a meaning of some expression—either a standing meaning or an occasion meaning. Pretty much everything is thus in its domain. This fact isn't problematic in itself. The

identity function, for example, ranges over absolutely everything but it can still easily be understood and represented by humans. But we are in the business of uncovering a finitely statable theory that models some actual interpretive mechanism and plays its part in explaining linguistic comprehension.

Not only are both of these functions unconstrained in their applicability, they are also supposed to model a genuinely productive and creative process. Knowing when and how to apply the metonymic function in interpretation is very different from knowing when and how to apply, say, the addition function for numbers. This is exactly why IBE is taken by many to be a plausible first step in understanding processes of interpretation (cf. Wilson and Sperber 2012). Interpreters consider all the possible evidence, typically taking the actual utterance as primary, and creatively come up with a hypothesis as to what communicative intention best explains the speaker's behavior. But who would take seriously the following explanation of IBE: it is a function from pairs of contexts and events (or whatever) onto explanations of those events? If interpretation is IBE-based such an explanation won't do for it either.

In §6 I will indicate the kind of theory that I favor. But there are no surprises. I am agnostic about whether a principle of compositionality is a non-negotiable condition on any theory of linguistic comprehension. The kind of principle one could adopt, however, is one that takes linguistic meanings of constituent expression types as inputs and yields meanings of compound expression types. Modulation functions are eliminated and the explanatory slack is picked up by a theory of communicative intentions and the mindreading capacities of human beings.

5.2 The argument from particularity

The standard contextualist recoil from the worry about generality is some type of particularism. This is the idea that it is only as 'particular historical events' and in 'particular circumstances' that utterances are true or false (cf. Travis 2008: 254). Recanati equates this with 'meaning eliminativism' which "gets rid of abstract meanings for types, in favor of particular uses" (Recanati 2004: 151).[13] Particularism thus construes (M) not as general rule that subsumes each and every successful linguistic interaction, rather, (M) only applies to particular cases individually.

If this means that the theory cannot generalize over types of meanings and types of contexts it doesn't get us very far. Then there is no way for the theory to state rules that apply to distinct occasions of utterance. In every case, we will have a function or rule that says that when applied to *this* meaning and *this* context, *mod* yields such and such a pragmatic function. The rule will be *about* this particular meaning and this context, i.e. these will be part of the *content* of the rule. This is arguably what happens in (3), with respect to which the particularist asserts that

[13] Travis (1989, 1997) is one example of a particularist in this sense. Recanati (2004: 151) argues that the view has a "surprising viability."

in this particular context, this particular meaning (CITY), being incompatible with something in the particular context, gives rise to metonymy. What does this tell us, as interpreters, about how to recognize metonymy elsewhere? Nothing at all; there is no such thing as metonymy *in general*. The particularist might, at this point, say that interpreters can still see similarities and patterns in different linguistic contexts. And this is certainly true. But I can't understand this proposal unless it simply reintroduces types of contexts to generalize over (some of which are of the metonymic type). If so, I am all ears. But for reasons that should be clear by now I doubt that this can be done while respecting the requirements of a compositional language.

There is a distinction between a rule that can be applied to an indefinite number of cases and a rule that is stated again for each new case. Take chess. The rule for the bishop, say, states that it is limited to diagonal movement. This applies to all bishops in all normal chess contexts. A difference in context will make for a difference in the token objects to which the rule applies but not a difference in *what the rule states*. The content of the rule is not that *this* bishop is limited to diagonal movement. If it were, it wouldn't be a rule of chess—it would be a 'rule' about this particular piece of wood, stating the 'function' of this piece of wood, with no implications whatsoever for any other bishop-shaped pieces of wood. Thus one learns nothing about the movement of the bishop by being told a rule which only applies to some particular bishop.

Projecting (M) from one context to another is similarly problematic. Knowing that the pragmatic function g_{513} maps the meaning of 'the city' onto the meaning of 'the inhabitants of the city' doesn't help in other contexts because g513 is itself particular to its own context. When the meaning of 'Homer' is mapped on to the meaning of 'the writings of Homer' by some function g there is nothing about the context that tells you that this is the very same metonymic function. And it isn't, since the function is a distinct mapping between members of the domain M. The functions are the same only in the sense that they are both mappings within domain M. But it should be clear by now that this doesn't unify the functions in any substantive way. The problem with particularism does not lie in the potential infinity of primitive meanings or contexts, but in the potentially infinite number of pragmatic *functions*, each particular to a given context, which one must master as a speaker of the language. Knowledge of these functions isn't explained compositionally, as the identity of the function itself is context-dependent.[14]

[14]See footnote 9. Pagin (2005: 318) discusses this problem for a specific compositional analysis of 'It rains': "…to introduce context sensitivity in the composition rule, may be claimed to violate compositionality, the reason being that in compositional semantics, the rules must not be context sensitive." But he thinks this is both right and wrong. Stanley (2007: 34) makes a similar point.

6 Occasion meaning versus speaker meaning

I believe that the source of the sandbagger's problems is the very distinction between standing meaning and occasion meaning. The category of occasion meanings doesn't seem to be a basic category for a theory of interpretation. As Bach (1987: 85–88) and Neale (2005) have argued it's clear that such a theory will need at least two notions of meaning. First a notion of what a speaker means on a given occasion of utterance and, second, a notion of the meaning encoded by linguistic expression types relative to the language to which they belong. If occasion meanings don't reduce to either of these two categories it's not clear why they're needed. Recanati (2010: 32–33) writes that expression tokens or expressions relative-to-contexts can 'assume' or 'take on' so-called occasion meanings. And it's important to remember that the set of occasion meanings contains *both* modulated meanings *and* the contextually determined contents of indexicals and demonstratives, i.e. expressions with Kaplanian characters (2010: 42). But no good reason has been given for this terminology. At any rate, it's more plausible to think that occasion meanings are simply meanings intended by the speaker on a specific occasion of utterance. And it's not necessary to think that these meanings become properties of the expression tokens uttered. This goes against a common way of talking in philosophy of language. Many theorists claim that the expression *token*, the *tokening* of the expression, the expression *relative to the context*, or some such entity acquires a contextually determined meaning that is not identical to the linguistic meaning of the expression type.[15] Occasion meanings as properties of expression tokens is what makes the arguments from generality and particularity possible. And as Neale (2005: 183) points out, positing token meanings as distinct from type meanings makes unwanted equivocation all too easy for theorists. As a working hypothesis, our theory should use only two kinds of entities to replace *X* in sentences of the form '*X* means *such and such*,' namely speakers and expression types. In this section I present two arguments against occasion meanings as such.

6.1 The argument from token invariance

When the sandbagger posits occasion meanings as properties acquired by expression tokens she incurs the debt of explaining exactly what causes this to happen. It is not sufficient simply to posit a pragmatic function specifically designed for this purpose. The explanatory power gained by positing the function can just as well be achieved without assuming that the function's output becomes a property of token expression (see §6.2). The output can, for instance, be a property of the speaker's communicative intention. At any rate, there are two plausible candidates to explain this kind of token variability. I am not sure which fits sandbag semantics better, but it doesn't matter since neither of them is very convincing.

[15] Cf. Devitt (1981, 1996), Stanley (2007). Travis (1989, 2008) seems to use 'speakings' similarly.

First, one might argue that there is a convention or rule that holds within the linguistic community in question to the effect that, under specific circumstances, the meaning m' of a token of the expression e will be distinct from the meaning m that e has in the language. For example, it's a convention that sometimes a token of the expression 'the city' will have the meaning of the expression type 'the inhabitants of the city.' Similarly, it's a convention that sometimes a token of

(4) He's a fine friend,

is uttered ironically so that it acquires something like the meaning of the expression type 'He's a terrible friend.'[16] But it seems like all expressions in the language can be used non-literally on some occasions. This includes irony, sarcasm, pretense, metaphor and many other tropes. Furthermore, any such convention can be violated. These different tropes are compatible with there being no distinguishing factor at the level of the meaning of the expression type.

Now, let's assume for sake of argument that the convention in question is explicit in the language itself. In this language L speakers will always prefix a trope marker to the original sentence. Thus instead of 'He's a fine friend' one would utter (5) in order to distinguish the non-literal meaning from the intended literal meaning.

(5) I mean the following literally: He's a fine friend.

One might utter (5) to make sure no irony at all is suggested. But the notion of literal meaning simply doesn't carry such weight by itself. If L is a natural language it is likely that (5) could be uttered non-literally.[17] Is the trope marker in (5) also meant literally? And if so, will the speaker not need a metatropical metamarker to make sure this is communicated? If so we appear to have an infinite regress. An 'irony marker' would obviously have the same consequences. Since such an explicit convention cannot fully explain how token expressions acquire their non-literal occasion meanings I take it that an implicit convention, depending on emphasis or tone of voice for example, cannot do so either. Such a convention can always be successfully violated. If Gibbs and Colston (2012) are to be trusted, empirical studies definitely support this conclusion. As they put it: "To date [...] no reliable procedure has been established for irony identification" (p. 53). Thus we need some other type of explanation for occasion meanings to be accepted.

The other candidate explanation is to say that occasion meanings become properties of expression tokens because of the speaker's communicative intention.

[16] The example is from Grice (1989: 53–54). See Camp (2012) for discussion on this point.

[17] Anglophones correct each other in some cases where 'literally' is used, because the expression 'figuratively' would have been more appropriate: "My mother *literally* hangs around the house all day."—"Is your mother a bat by any chance?" I think this is simply wrong; speakers are not guilty of any error and 'figuratively' would just not do the work which it would be intended to do.

Thus, in a typical case of modulation, the speaker intends on the occasion of uttering a token of the expression type Σ that the token itself shall have the meaning m where m is the output of the modulation function. The occasion meaning of personal pronouns is explained in a similar manner. The occasion meaning of a token of the type 'she' is a certain salient female person since the speaker intended this token to take on that occasion meaning. It seems implausible to suppose that ordinary speakers have such intentions with respect to expressions—token or type. Speakers clearly don't intend to modulate the meaning of the expression *type* by uttering a token of the type. This would be like intending to change what the dictionary says by uttering an expression non-literally. What about the actual token uttered? This issue is more complicated but I believe it can be resolved by appealing to a distinction inspired by Grice's (1989: 100–102) and Schiffer's (1972: 27–30) responses to Searle's (1965, 1969: 44–45) well-worn example of the American soldier. Searle imagines an American soldier that is captured by the Italians in the Second World War. The soldier would like to convince the Italians that he is a German soldier so he won't be imprisoned. Hoping that the Italians don't understand German very well, he utters the only German sentence he can remember, 'Kennst du das Land who die Zitronen blühen?' Now, it seems to follow from the Gricean account of speaker meaning that what the speaker meant in uttering the German sentence was *that he is a German soldier*. This was the the proposition he intended the Italians to believe as a result of recognizing his communicative intention. Searle thinks this is implausible. As he puts it, the soldier cannot have meant this because he knows that the sentence *itself* means something quite different.

Without going into the details of the Gricean view there is a simple distinction that can be made here. It was no part of the soldier's intention that the token sentence he uttered should mean *that he is a German soldier*. The soldier need not have any such intention. All he needs is the intention to communicate this proposition *by the act of* uttering the German sentence. And this is what the basic category of speaker meaning requires. Speakers mean such and such by producing some evidence that they mean such and such. In the context of the soldiers, uttering the German sentence can very well constitute evidence of meaning *that one is a German soldier*. Of course, the connection between the meaning of the sentence in the language and the meaning intended by the speaker is usually closer than this, but it need not be very close.[18] Moreover, if it were a necessary condition on speaker meaning that the expression Σ (type or token) should mean in L exactly what the speaker means by uttering Σ on some occasion, then a non-circular account of linguistic meaning in terms of communicative intentions would be impossible (cf. Szabó 2005: 4n7). Thus we should assume that the meaning of the token expression remains invariant insofar as the meaning of the expression type doesn't change.[19]

[18] Cf. Carston (2002), Neale (1992), Sperber and Wilson (1986, 1995), Wilson and Sperber (2012).

[19] I'm ignoring at least one alternative explanation, but it would require a much longer discussion. I have in mind Ruth Millikan's claim that expression types (or 'linguistic devices' in general) have

6.2 The idle wheel argument

I agree with Bach (1987: 87–88) that the semantic properties of expression tokens play no explanatory role in the theory of communication. If tokens have any 'semantic' properties those properties can always be reduced to the semantic properties of speaker's communicative intentions. In terms of explanatory power, positing an extra layer of occasion meaning that attaches to tokens is just an idle wheel.

Let me explain this in more detail. As I have already indicated, there are good reasons to think that the proposition expressed, or the content of what is said, by a speaker in uttering an expression Σ in language L on an occasion is always underspecified by any linguistic meaning Σ itself has in L. This follows straightforwardly from the fact that no natural language expression need, on any occasion, be uttered such that the speaker only means what Σ means in L. One can know the abstract meaning of Σ by knowing (some part of) the linguistic conventions or regularities in virtue of which it has that meaning in L. But this knowledge is insufficient for the hearer to determine completely what the speaker said on some occasion. Indeed, as Kent Bach notes, even if the content of what the speaker said happens to correspond *exactly* to the linguistic meaning of Σ, this fact will not hold in virtue of the linguistic meaning of Σ. This fact is constituted by something else, most likely the speaker's intentions.[20] If this is right, the epistemological problem of the hearer, i.e. that of discovering the speaker's communicative intention, is not different in any essential way when the interpretation happens to be 'literal.'

In the framework of Recanati's truth-conditional pragmatics, the gap between speaker's meaning and linguistic meaning is to be closed, at least partly, by invoking the occasion meaning of Σ in addition to the standing meaning of Σ. Ultimately, I don't think the gap gets any smaller by postulating occasion meanings. Compositional rules range over standing meanings of expression types, not their occasion meanings, and their output is the standing meaning of compound expression types.[21] Consider the example of genitive case constructions in English.

(6) I read John's book.

Suppose that the occasion meaning of (6) involves a relation R_{54} between authors and their works as opposed to, say, another relation between owners and their

direct proper functions while token expressions may have specific *derived* proper functions. I hope to address this theory elsewhere. See Millikan (1984, 1989a, 1989b), but also Origgi and Sperber (2000).

[20] Bach (2005: 26–27), cf. Carston (2002: ch. 1).

[21] Thus I disagree with King and Stanley (2005: 123–124) when they claim that only contents compose, never characters (or 'unrelativized semantic values'). They think a theory need not assign characters to sentences since it is sufficient to assign them only to simple expressions. But it seems trivial that once one assigns character to the constituents of a sentence one has already assigned a character to the sentence as a whole (cf. Pagin and Pelletier 2007).

books. It seems to be Recanati's view (2004: 63, 2010: 182) that R_{54} is part of the occasion meaning of (6) in virtue of the speaker's communicative intention. Since he also thinks that intentions are part of the context (2010: 184) we can say that R_{54} is part of the occasion meaning because (6) is uttered in the *author-work context*, i.e. c_{54}. But it doesn't add to the explanatory power of the theory to say that the relation is part of the occasion meaning which is a property of the expression token. What work does this epicycle do that isn't done by the speaker's intention in context c_{54} according to this theory? None it seems, for if the hearer can infer that (6) was uttered in c_{54} she doesn't need the extra step of supposing that the expression token (6) assumes R_{54} as part of its occasion meaning. According to an influential alternative view, the only compositional information provided by a simple sentence like

(7) New York snores,

is that the name 'New York' is an expression of (Montagovian) type *e*, standing for an object, that (7) is an expression of type *t* and stands for a truth-value, and that 'snores' is of type $\langle e, t\rangle$, denoting a function from individuals to truth-values.[22] As such, any snoring object called 'New York' can make (7) true although only one such can make what a speaker said in uttering (7) on some occasion true. Thus the sentence merely puts constraints on which relations the speaker can (reasonably) intend to say and mean. Supposing that the composition of 'New York' and 'snores' takes their occasion meanings (where the former meaning might be metonymic) as input, rather than their linguistic type meanings, will not change this. The hearer's job is always to work out the intention with which (7) is being uttered. The type-token and standing-occasion distinctions won't give her any advantage in interpreting an utterance of (7). She will always rely on some kind of IBE based, among other things, on the linguistic meaning of (7). If the token has any meaning as such, this meaning is derived from the type.[23] Thus I conclude that occasion meanings are superfluous in a theory of communication.[24]

[22] Cf. Heim and Kratzer (1998). Yet it's important that the argument in this paper doesn't depend on the idea that composition is (largely) function application. I think more recent ideas about propositions as act types or event types may be more promising—even from the standpoint of compositionality (cf. Hanks 2011; Soames 2010).

[23] Cf. Searle (1978: 209): "Barring diachronic changes, special codes, and the like, the meaning of the token is always the same as the meaning of the type."

[24] A reviewer suggests that the meaning arrived at by the hearer in interpreting an utterance might be construed as the occasion meaning of the utterance. Of course a theory of interpretation needs such an entity, but, in my opinion, it is ill-suited as occasion meaning. Occasion meaning is the *correct* interpretation of the utterance and it's supposed to be a property of the actual token uttered. And our theory must allow for the possibility of misunderstanding: the hearer may come up with the wrong interpretation. I don't see much use in calling wrong interpretations occasion meanings. Neale's (2005: 179–180) remarks on epistemic asymmetry in communication are relevant here.

The basic reason why Recanati thinks we need an intermediate category of occasion meanings, between expression (type) meaning and speaker meaning, is to avoid defeating what Perry calls the principle of 'homomorphic representation' and Crimmins the principle of 'full articulation.'[25] A sentence will violate the principle if the proposition expressed by uttering it has constituents not corresponding to any component of the sentence itself. Thus uttering 'It is raining' to express the proposition *that it is raining in Paris* violates the principle. Recanati thinks that, rather than violate such a principle, it is more satisfactory to say that some word acquires a specific sense in the context. Thus 'raining' is made equivalent to 'raining in Paris' and, similarly, 'eat' becomes equivalent to 'eat dinner' in some contexts (2010: 124).

I don't think this terminology is helpful. There is a distinction between mere abbreviation and underspecification. The former is a relation between two things of the same type, e.g. two expressions. Thus the expression 'CEO' abbreviates and in *some* sense underspecifies the expression 'chief executive officer.' Such relations are symmetric, so that the latter expression will actually 'lengthen' the former. The only way, I think, to understand Recanati's idea is to see it as, exactly, a relation that is supposed to hold between two (type or token) expressions. But that cannot be the case. The expressions (type or token) 'raining' and 'raining in Paris' are not related in such a way that one is a shortening or a lengthening of the other. There is no conventional relation of abbreviation between the two expressions. Specifically, speakers do not utter the shorter expression in order to convey or indicate the longer *expression*. Suggesting that there is abbreviation here immediately falls foul of the meaning-intention problem. Speakers do not have any intention, with respect to the expression 'rain,' that *it* (type or token) should abbreviate 'rain in Paris' (which is an expression, not a meaning) on some occasions. They do, however, regularly utter the shorter expression to say something that could just as well have been said by uttering the longer expression. I think the same can be said about the pairs ⟨'eat', 'eat dinner'⟩ and ⟨'drink', 'drink alcohol'⟩ mentioned by Recanati, but I won't pursue the point here.

Underspecification, however, seems to be a relation between things of different types. Abbreviation is mere variation in notational devices. But the kind of underspecification that we are interested in holds between the linguistic meaning of the expression and the proposition expressed by a speaker in uttering that expression: one being (part of the) evidence provided and the other being the proposition intended. This relation is asymmetric and rules out the possibility of overspecification (i.e. 'lengthening'): it doesn't matter how verbose or repetitive the speaker is, the meaning of the expressions she utters will always underdetermine the proposition intended. Consider sentence (8).

[25] Perry (2000: 174), Crimmins (1992: 10). Quoted in Recanati (2010: 121–122).

(8) I have already had breakfast earlier today.[26]

In uttering (8) the speaker used at least three expressions to indicate the time at which his breakfast-eating occurred, i.e. 'already,' 'earlier' and 'today' (I say 'at least' because the verbs have tenses). That does not change the fact that the temporal specification actually intended by the speaker is underspecified by the encoded meaning of (8). This is simply because the hearer needs, at least, also to know the time of utterance and what day it is in order to understand what was said. Adding more temporal components to the sentence will not change this, it will merely invite different ways in which the sentence may underspecify the proposition expressed.[27] Thus, even in a case where someone utters,

(9) It's raining in Paris

the expression itself will underspecify the proposition expressed. This really should not be surprising, since the speaker might be talking about a town in Illinois or a town in his favorite novel. And even if the intended interpretation of 'Paris' is successfully conveyed in an utterance of (9) it doesn't follow that the *expression* in (9) is sufficient to constitutively determine this interpretation. Only the speaker's intention will constitute such a determination.[28]

To sum up, the occasion meaning of Σ cannot be conventionally related, via some type of abbreviation, to Σ's standing meaning for the former goes beyond what is linguistically encoded. But neither can the occasion meaning of Σ, as it turns out, be identified as what the speaker intended by uttering Σ on that occasion. This is because the speaker only intends, say, a specific place *by* uttering 'It rains.' The speaker doesn't intend for the expression 'rains' to *take on* or *assume* the meaning of the expression 'rains in Paris, Illinois.' Same goes for 'the city' and 'the inhabitants of the city.' This would be absurd and the absurdity isn't avoided by adding that the intention is directed at the tokening of the expression rather than the expression type.

7 Conclusion

Suppose for a moment that linguistic expressions are tools. Tools usually have some abstract function. Not without simplification, one could say that the function of a hammer is to drive nails. Knowing this makes it possible for one to say of some particular hammer that it has the function of driving nails. But what should

[26] Cf. Sperber and Wilson (1986, 1995: 189–190)

[27] Even if we consider an 'eternalized' version of (8), such as 'I had breakfast at 9am, May 24, 2011,' this will still be the case. But the point is not important for my present argument, since I can concede that full determination is possible *in principle*. But I'm not sure how interesting this is in light of the fact that people almost never communicate by using such 'eternalized' sentences.

[28] Here I agree with Bach (2005) and Neale (2005).

one say about some 'token use' or 'occasion function' of a hammer? Will some particular tokening of a hammer-in-use acquire some specific function? I don't think so. If one hangs one's favorite hammer up on the wall for decoration (not intending to drive nails any more) that token use of the hammer does not take on the decoration function. Rather, the agent uses the hammer for decoration. The decoration function does not *thereby* become a new possible function for hammers or the function of this token use of a hammer. Hammers still drive nails.

Similarly, speakers utter expressions with the intention to say something. They do not utter expressions with the intention that the expression take on some specific contextual meaning they have in mind. Austin's lectures were about the things one can do with words, not the things one can make words do. In this paper, I have given reasons to doubt the notion of pragmatic compositionality. First, if we assume a Davidsonian notion of compositionality, it becomes doubtful whether the requirement that the base of semantic primitives be finite is satisfied by pragmatic composition. This turns out not to be a problem, because we can have the fruits of compositionality (learnability at least) without a finite base. Secondly, however, it turns out that the pragmatic theory, after taking on the requirement of compositionality, doesn't identify a function that's plausibly thought of as part of an explanation of linguistic comprehension. Thirdly, I showed how the contextualist provides a false picture of linguistic underspecification, where it amounts to a relation between the standing meaning of expression types and the occasion meaning of expression tokens. The relation we are interested in is, rather, the relation between the meaning of the expression type and the proposition the speaker intended in making the utterance. Furthermore, these arguments provided reasons to doubt the propriety of the very notion of occasion meanings or the meaning an expression has in context. There seems to be no theoretical reason to postulate such meanings. Recognizing the intentions of speakers is based on the linguistic evidence they provide and any contextual clues one can use in the process. Our ability to do this is guided by assumptions of rationality and cooperation. Adding an indefinite number of occasion meanings and pragmatic functions into the computational mix is only going to make the process itself unduly complex and psychologically implausible.

Bibliography

Bach, K., 1987. *Thought and reference.* Clarendon.

———, 2005. "Context *ex machina*." Z. G. Szabó (ed.), *Semantics versus pragmatics*, Clarendon, pp. 15–44.

Bradley, D., 2014. "Functionalism and the independence problems." *Noûs*, 48(3):545–557.

Camp, E., 2012. "Sarcasm, pretense, and the semantics/pragmatics distinction." *Noûs*, 46(4):587–634.

Carston, R., 2002. *Thoughts and utterances.* Blackwell.

Cohen, L. J., 1986. "How is conceptual innovation possible?" *Erkenntnis*, 25(2):221–238.

Crimmins, M., 1992. *Talk about belief.* MIT Press.

Davidson, D., 1965. "Theories of meaning and learnable languages." Y. Bar-Hillel (ed.), *Proceedings of the 1964 International Congress for Logic, Methodology and Philosophy of Science*, North Holland Publishing Co, pp. 383–394. Repr. in Davidson (1984), pp. 3–15.

———, 1984. *Inquiries into truth and interpretation.* OUP.

Dennett, D., 1991. *Consciousness explained.* Back Bay Books.

Devitt, M., 1981. *Designation.* CUP.

———, 1996. *Coming to our senses: A naturalistic program for semantic localism.* CUP.

Downing, P., 1977. "On the creation and use of English compound nouns." *Language*, 53(4):810–842.

Fodor, J., 1981. *Representations: Philosophical essays on the foundations of cognitive science.* The Harvester Press.

———, 2003. *Hume variations.* OUP.

Gibbs Jr, R. W. & Colston, H. L., 2012. *Interpreting figurative meaning.* CUP.

Grice, P., 1989. *Studies in the way of words.* HUP.

Hanks, P. W., 2011. "Structured propositions as types." *Mind*, 120(477):11–52.

Heim, I. & Kratzer, A., 1998. *Semantics in generative grammar.* Blackwell.

King, J. & Stanley, J., 2005. "Semantics, pragmatics and the role of semantic content." Z. G. Szabó (ed.), *Semantics versus pragmatics*, Clarendon, pp. 111–164.

Lepore, E. & Ludwig, K., 2011. "Truth and meaning redux." *Philosophical Studies*, 154(2):251–277.

Millikan, R. G., 1984. *Language, thought, and other biological categories: New foundations for realism.* MIT Press.

———, 1989a. "Biosemantics." *The Journal of Philosophy*, 86(6):281–297. Repr. in Millikan (1993), pp. 83–102.

———, 1989b. "In defense of proper functions." *Philosophy of Science*, 56(2):288–302. Repr. in Millikan (1993), pp. 13–30.

———, 1993. *White queen psychology and other essays for Alice.* MIT Press.

Neale, S., 1992. "Paul Grice and the philosophy of language." *Linguistics and Philosophy*, 15(5):509–559.

———, 2005. "Pragmatism and binding." Z. G. Szabó (ed.), *Semantics versus pragmatics*, Clarendon, pp. 165–285.

Origgi, G. & Sperber, D., 2000. "Evolution, communication and the proper function of language." P. Carruthers & A. Chamberlain (eds.), *Evolution and the human mind*, CUP, pp. 140–169.

Pagin, P., 2005. "Compositionality and context." G. P. G. Preyer (ed.), *Contextualism in philosophy: Knowledge, meaning and truth*, OUP, pp. 303–348.

Pagin, P. & Pelletier, J., 2007. "Context, content and communication." G. P. G. Preyer (ed.), *Context-sensitivity and semantic minimalism*, OUP, pp. 25–62.

Perry, J., 2000. *The problem of the essential indexical and other essays: Expanded edition.* CSLI Publications.

Recanati, F., 1995. "The alleged priority of literal interpretation." *Cognitive science*, 19(2):207–232.

———, 2001. "What is said." *Synthese*, 128(1):75–91.

———, 2002. "Does linguistic communication rest on inference?" *Mind & Language*, 17(1-2):105–126.

———, 2004. *Literal meaning.* CUP.

———, 2010. *Truth-conditional pragmatics.* OUP.

Schiffer, S., 1972. *Meaning.* OUP.

———, 1987. *Remnants of meaning.* CUP.

———, 2003. *The things we mean.* Clarendon.

Searle, J., 1965. "What is a speech act?" M. Black (ed.), *Philosophy in America*, Allen and Unwin, pp. 221–239.

———, 1969. *Speech acts: An essay in the philosophy of language.* CUP.

———, 1978. "Literal meaning." *Erkenntnis*, 13(1):207–224.

———, 1980. "The background of meaning." J. Searle, F. Kiefer, & M. Bierwisch (eds.), *Speech act theory and pragmatics*, Reidel, pp. 221–232.

Soames, S., 2010. *What is meaning?* Princeton University Press.

Sperber, D. & Wilson, D., 1986. *Relevance: Communication and cognition.* Blackwell.

———, 1995. *Relevance: Communication and cognition.* Blackwell, 2 edn.

Stanley, J., 2007. *Language in context.* Clarendon.

Szabó, Z. G., 2005. "Introduction." *Semantics versus pragmatics*, Clarendon, pp. 1–14.

Travis, C., 1989. *The uses of sense.* OUP.

———, 1997. "Pragmatics." B. Hale & C. Wright (eds.), *A companion to the philosophy of language*, Blackwell, pp. 87–107. Repr. in Travis (2008), pp. 19–64.

———, 2008. *Occasion-sensitivity: Selected essays.* OUP.

Weiskopf, D. A., 2007. "Compound nominals, context, and compositionality." *Synthese*, 156(1):161–204.

Westerståhl, D., 2012. "Compositionality in Kaplan style semantics." M. Werning, W. Hinzen, & E. Machery (eds.), *The Oxford handbook of compositionality*, OUP, pp. 192–219.

Wilson, D. & Sperber, D., 2012. *Meaning and relevance.* CUP.