

1

Content, Context, and Composition

Peter Pagin and Francis Jeffry Pelletier

1. INTRODUCTION

It is traditional, at least since Grice, to make a distinction between what is called *the literal meaning* of an utterance and *what is meant* by that utterance. The former notion is sometimes thought of as “the dictionary meanings of words plus standard semantic effects of the syntactic rules” that were employed in the utterance. The latter notion is often thought of as the “all things considered” information that is conveyed by the utterance in the context it is used. The former is often said to be the “context independent meaning” or “the meaning that is constant across all contexts of use”. The latter is often called “what is conveyed in a context”. The former is sometimes called “the timeless meaning” or “the meaning-in-the-language” or “the meaning of the linguistic *type* of the utterance”. The latter is sometimes called “the meaning of the speech act being performed” or “the meaning of the *token* being uttered”.

Intuitively, the former notion allows *no* features of the context in which an utterance is made to enter into its semantic evaluation. The latter notion seems to claim that *every* aspect of the context might be relevant to the evaluation of the utterance.

The notion of *what is said*—as opposed to “the literal meaning” and as opposed to “what is meant”—is a theory-laden notion that is intended to locate an important semantic feature of linguistic communication. The idea is that there is some feature that identifies the *semantic meaning* of an utterance and separates it from its *pragmatic meaning*. Now, these two notions are also theory-laden terms, but there is at least agreement that “literal meaning” is part of semantics while “what is conveyed” is a part of pragmatics. The question has always been: where should the line be drawn? The point where the line is drawn identifies *what is said*, and any further information that

might be gleaned in some communicative act will be classified as *what is meant*.¹

In two recent publications (Recanati 2004, 2005), François Recanati presents a way to organize different theories of language—that is, theories of *what is said*—in accordance with “how much context” the theories will allow as a part of their semantic component. These theories range from “pure literalism”—theories of linguistic semantics that have no use for any kind of contextual information, not even indexicality, and make “what is said” be the same as “literal meaning”—and continue through a series of ever-more-context until we reach “pure contextualism”, a theory of the role of semantics that Recanati calls “meaning eliminativism” because “what is said” is entirely a matter of context with no contribution of the “literal meaning”. Between these two extremes lie a number of theories that limit the amount of context that is allowed into semantics in one way or another, and are called such things as eternalism, indexicalism, syncretism, quasi-contextualism, and full-blooded contextualism. Recanati notes, for example (2004: 92 n. 20), that these intermediate views come in degrees, so that there are actually a number of different syncretist views, for example.

1.1. Cappelen and Lepore Background

In their recent book, Cappelen and Lepore 2005, Herman Cappelen and Ernie Lepore (CL), propose and defend two views, labeled ‘Semantic Minimalism’ (henceforth Minimalism) and ‘Speech Act Pluralism’ (henceforth Pluralism). Generally speaking, Minimalism is the view that Recanati called ‘indexicalism’. (We say this despite Recanati’s claim that CL’s view is syncretic: Recanati 2004: 92.) CL’s chief arguments for these views consist in discrediting what they portray as the main alternatives, *radical contextualism* (theories such as full-blooded contextualism that approach Recanati’s meaning eliminativism) and *moderate contextualism* (any of Recanati’s versions of contextualism “between” indexicalism and full-blooded contextualism). These arguments fall into two categories: on the one hand direct arguments against Radical Contextualism (part 2 of their book), and on the other hand indirect arguments against moderate contextualism (part 1). The indirect ones are arguments that Moderate Contextualism *leads to* Radical Contextualism, and because of that, the arguments against Radical Contextualism will apply to Moderate Contextualism as well. We shall here be concerned with the indirect arguments, those called the Instability Arguments, and will simply accept their arguments against Radical Contextualism. For it is

¹ If one decides that the line should be drawn all the way at the end—at the level of “what is conveyed”—then there are two choices available. One could say that semantics has all of “what is conveyed” in its scope, or that there is no independent semantic theory and “it’s all pragmatics”. Most of those who hold that this is the only place to draw the line also hold to the position that “it’s all pragmatics”.

our goal to defend a kind of Moderate Contextualism that does not lapse into Radical Contextualism. But first a couple of terminological explanations.

In order to explain what they mean by ‘Semantic Minimalism’ CL give a list of expressions that they characterize as ‘The Basic Set of Context Sensitive Expressions’ (CL 2005: 1–2), comprising the personal pronouns (in their various grammatical forms), demonstratives, the adverbs ‘here’, ‘there’, ‘now’, ‘today’, ‘yesterday’, ‘tomorrow’, ‘. . . ago’, ‘hence(forth)’, the adjectives ‘actual’ and ‘present’, tense and aspect indicators generally, common nouns like ‘enemy’ and ‘foreigner’, and adjectives like ‘foreign’ and ‘imported’.

CL immediately point out that the basic set does not contain a number of terms that in recent decades have been given a contextualist analysis, such as ‘every’ (contextual quantifier domain restriction), ‘know’ (contextual standards of knowledge), or ‘happy’ (contextual comparison class).

With reference to the basic set, CL (2005: 2) characterize Minimalism effectively as follows

- (i) Only expressions in the basic set (‘plus or minus a bit’) are accepted as context-sensitive.
- (ii) All context sensitivity is grammatically triggered.
- (iii) Context has no effect on what proposition an utterance expresses other than to fix the semantic value of these context-sensitive expressions.

Since mainstream natural language semantics provides contextualist analyses of many more expressions (such as those explicitly excluded above) and constructions than CL recognize, Minimalism is a highly controversial position. It also has clearly counterintuitive consequences, in that some sentences that appear context-sensitive, such as

- (1) Tipper is ready

(CL 2005: 60; Bach 94*a, b*), will not be counted as such. Disregarding the tense ingredient of (1), we would still think that nothing one can say with (1) is fully determined by the sentence itself: Tipper will be ready for some things and not for others, and it has to be understood from the context what Tipper is claimed to be ready for. This is denied by CL (2005: 116), who give the general semantic description

- (2) Every utterance *u* of ‘A is ready’ expresses the proposition *that A is ready*.

Pluralism is introduced as follows:

No one thing is said (or asserted, or claimed, or . . .) by any utterance: rather, indefinitely many propositions are said asserted, claimed, stated. What is said (asserted, claimed, etc.)

depends on a wide range of facts other than the proposition semantically expressed. It depends on a potentially indefinite number of features of the context of utterance and of the context of those who report (or think about) what was said by the utterance. (CL 2005: 4)

Pluralism, too, is both controversial and counterintuitive. A counterintuitive consequence explicitly endorsed by CL is that a speaker can *say* something by a sincere utterance without believing what she says (indeed, while explicitly *disbelieving* what she says). For instance by means of an utterance of

(3) That man is shady

the speaker may have *said that* the moronic clown is shady (CL 2005: 202–3), and yet not believe it, since not believing that the man referred to is a moronic clown (in fact, even while *disbelieving* that the man referred to is a moronic clown). The indirect discourse attribution is nonetheless correct if the man in fact *is* a moronic clown.

This view is a consequence of CL's very liberal theory about the correctness of indirect discourse attributions, going back to CL 1997. A consequence of this view, and explicitly part of Pluralism, is the nihilistic claim

(N) *There can be no systematic theory of speech act content*

(CL 2005: 190). This claim motivates separating semantic theory from the theory of speech act content (nearly enough²), for on their view this is the only way of keeping semantic theory itself systematic. We will return to this general question in §6.

Radical contextualism is in one respect close to being the opposite of Minimalism: the former maximizes the ratio between the pragmatic-contextual and the semantic contribution to any speech act content, while Minimalism postulates (for every utterance) one speech act content for which the ratio is minimal. Radical contextualism is characterized as follows by CL:

(RC1) No English sentence *S* ever semantically expresses a proposition.³
Any semantic value that Semantic Minimalism assigns to *S* can

² Although CL argue that there can be no systematic account of whatever else can be said, they do think there can be a systematic account of "semantic content". And since semantic content is always correctly reported as having been said, they say, it follows that there is a systematic account of *part* of what is said. So the separation of semantic theory and speech act content is not total, on their account.

³ Abstractly, there can be an ambiguity here, between the view that "only people, not inanimate words" can express anything and the view that no (finite) amount of verbiage can completely specify a situation in enough detail so as to be determinately true or false. We think the former of the two views is a mere terminological quibble if it has no support from the second view. For, if it were granted that speakers always express the same proposition by means of uttering a particular sentence *s*, then we could define another relation, *x* schempresses *y*, between expressions and content, and by which *s* schempresses that proposition. So this is not a very interesting proposal, unless it is backed

be no more than a *propositional fragment* (or *radical*), where the hallmark of a propositional fragment (or radical) is that it does not determine a set of truth conditions, and hence, cannot take a truth value.

- (RC2) Context sensitivity is ubiquitous in this sense: No expansion of what we are calling the Basic Set of context sensitive expressions can salvage Semantic Minimalism, i.e. however the Basic Set is expanded, the output will never be more than a propositional fragment; something, therefore, not even truth evaluable.
- (RC3) Only an utterance can semantically express a complete proposition, have a truth condition, and so, take a truth value. (CL 2005: 6)

As chief representatives of Radical Contextualism, CL designate John Searle (1978, 1980, 1983) and Charles Travis (1985, 1989, 1996).

Contextualism in general, and perhaps the radical variety in particular, is supported by the difficulties of tricky examples, or perhaps better, by intuitions fueled by recognition of these difficulties. The following example (CL 2005: 43, 64) is taken from Travis (1985: 197). The sentence

- (4) Smith weighs 80 kg

sounds determinate enough at first blush, but it could be taken as true or as false in various contexts, depending on what counts as important in those contexts. For example, it can be further interpreted as being true if Smith weighs

- (4a) 80 kg when stripped in the morning
 (4b) 80 kg when dressed normally after lunch
 (4c) 80 kg after being force fed 4 liters of water
 (4d) 80 kg four hours after having ingested a powerful diuretic
 (4e) 80 kg after lunch adorned in heavy outer clothing

The idea is that the literal meaning of (4) does not settle which of these further interpretations (4a–e), or yet others, is relevant in a context, and so fails to determine a propositional content on its own. Looking at a number of seemingly innocent examples in this light gives the impression that no sentence at all expresses a proposition on its own.

Moderate Contextualism does not go that far. Moderate Contextualism is characterized as follows by CL:

up by the claim in the second view. With this backing we would *not* be able to define something like ‘schempresses’.

- (MC1) The expressions in the Basic Set do not not exhaust all the sources of semantic context sensitivity
- (MC2) Many sentences that Semantic Minimalism assigns truth conditions to, and treats as semantically expressing a proposition, fail to have truth conditions or to semantically express a proposition; they express only fragmentary propositions. Such linguistic expressions are described as providing ‘incomplete logical forms’, ‘semantic skeletons’, ‘semantic scaffolding’, ‘semantic templates’, ‘propositional schemas’ . . . All of these locutions entail that the expression is not fully propositional; it is incomplete *qua* semantic entity; it is not truth evaluable.
- (MC3) For the cases in question, only their utterances semantically express a proposition, and have (interpretive) truth conditions, and so, take a truth value. (CL 2005: 7)

As a typical representative of (one kind of) Moderate Contextualists, CL select John Perry (1986), who made the sentence

- (5) It is raining

the paradigm example of so-called *unarticulated constituents* (Perry 1986: 206). The idea is that a speaker of (5) speaks of the weather at some location, even though there is no constituent of (5) that takes location as semantic value. The truth value of the sentence as a whole depends on assignment of a location value, but no part of the sentence *articulates* that dependence. As opposed to location, the tense of the verb articulates the dependence on a *time* value.

On CL’s view, there is no unarticulated location constituent in (5). Rather, the content of (5) is given by

- (6) ‘It is raining’ express the proposition *that it is raining* and is true iff it’s raining

(CL 2005: 61–3). Although it may seem that adding feature-placing sentences like (5) to the list of context-sensitive expressions is not a drastic addition, it is important for CL to reject it, since according to CL if this is accepted, we would have to accept a lot more. The underlying issue here is that there is no overt syntactic item in (5), or similar sentences, that legitimizes the addition of such features. But if we were to accept Moderate Contextualism of this sort, and give permission to posit additions not tied to anything in the sentence, then we would rationally be led to accepting Radical Contextualism, they say.

On our view, this is not right. In §6 we will endeavor to distinguish various types of contextualism, with the motivation that some of them are not susceptible to the types of considerations that CL employ in their rejection of Moderate Contextualism. The general picture is that we can expand CL's "Basic Set" with certain other words and constructions, such as comparative adjectives and domain restrictions, without having to be a radical contextualist. And as a part of this discussion, we will set out various considerations about compositionality for contextually sensitive terms and sentences. Here we give a framework for characterizing the differences between Recanati's "saturation" and "modulation", and we show by means of examples that it can be coherently applied to support moderate contextualism.

2. CAPPELEN AND LEPORE'S INSTABILITY ARGUMENT

The reason the Moderate position leads to the Radical, according to CL, is that if the *justification* offered for Moderate Contextualism is good enough, it justifies Radical Contextualism as well. This is said to hold of both types of arguments that have been adduced in favor of contextualism: *context shifting* arguments and *incompleteness* arguments.

2.1. Context Shift

The notion of a Context Shifting Argument (CSA) is not given a precise definition in CL, but is characterized as follows.

One way that philosophers of language [go about in establishing that an expression *e* not in the Basic Set is context sensitive] is to think about (or imagine) various utterances of sentences containing *e*. If they have intuitions that a *semantically relevant feature* of those utterances varies from context to context, then that, it is assumed, is evidence [that] *e* is context sensitive. . . . The kinds of features contextualists claim to have intuitions about include

- What is *said* or *asserted* or *claimed* by utterances of sentences containing *e*.
- The truth condition of utterances containing *e*.
- The proposition expressed by utterances of sentence containing *e*. (CL 2005: 17–18)

Summing this up, the form of a Context Shifting Argument would be something like

- (CS) If a *semantic feature* of utterances of sentences containing *e*, and associated with *e*, changes from context to context, then *e* is context-sensitive.

Now we agree with CL that this is not a good argument form. But (apparently) unlike CL, we think that (CS) can be modified to provide a good argument form, and we will come back to that later, in §6.⁴

As examples of Context Shifting Arguments, CL list alleged evidence for

- (a) shifts of quantifier domain
- (b) shifts of comparison class for comparative adjectives
- (c) shifts of content characterizing notions for belief reports
- (d) shifts of possible world comparative similarity for counterfactual conditionals
- (e) shifts of standards for knowledge attributions
- (f) shifts of content-determining factors for moral evaluations
- (g) shifts of standards of precision generally
- (h) shifts of location values for weather reports

To exemplify cases (a) and (b), consider the sentences

- (7) Every bottle is empty
- (8) That basketball player is short

An utterance of (7) (from Stanley and Szabó 2000: 219–20) is likely to be concerned with a restricted range of bottles, not with every bottle in the universe. Moreover, which restricted range of bottles it is concerned with depends on context, for instance every bottle on a particular table at some party, or every bottle on a particular shelf in some wine cellar. The intuitive content is then different, and the truth value can well (intuitively) be different, even if the utterances are made at the same time. Hence, by a (CS) argument, there is context sensitivity induced by ‘every’, or perhaps by ‘bottle’.

An utterance of (8) (Stanley 2002: 377) is likely to be concerned with height evaluation in relation to some salient group of objects of comparison. It can be truly said of a basketball player in relation to a group of other, taller, basketball players, but only falsely said of the same basketball player in relation a group of people of average height. Which group is relevant depends on context, and so the truth value of two utterances of the sentence, at the same time, with reference to the same player, can differ. The (CS) conclusion is that there is context sensitivity induced by ‘short’.

The main objection by CL against arguments of this kind is spelled out in their principle

⁴ As also stressed by our referees, it doesn’t follow from the fact that sentences containing *e* are context-sensitive that their sensitivity depends on the meaning of *e*. What we need is a globally well-supported semantic theory that implies this. This remains a task for CL’s description of Context Shifting Arguments.

- (GEN) With sufficient ingenuity, a CSA can be provided for any sentence whatsoever, and consequently, for any expression.

(CL 2005: 40). In order to establish (GEN), CL provide a series of examples. We shall here briefly recapitulate three of them. First, we have the Travis sentence

- (4) Smith weighs 80 kg

and CL present two scenarios (2005: 43). In the first case, a report is made when Smith has just eaten lunch and is fully dressed, but the conversation concerns his recent dieting, and what is relevant is what the scale registered in the morning of the utterance, before breakfast, with naked Smith. In this case, the utterance is true. In the second scenario, on the other hand, it would be false because Smith is about to enter an elevator which can take no more than an extra 80 kg, and the report is intended to be relevant to that. If, at that time, Smith, with clothes, after lunch, weighs more than 80 kg, the consequences are fatal. So two different things are said by (4) in the two contexts, and hence the sentence by a (CS) argument is deemed context-sensitive.

As a second example, CL give the sentence

- (9) John went to the gym

(CL 2005: 44–5). In the first context of utterance, the topic of conversation is John's walking habits. In this case the utterance of (9) is true if John walked to the vicinity of the gym. In the second context, the conversation is concerned with John's exercising habits, and the utterance would be true only if John did some workout at the gym. In the third context, John is involved in construction work at the gym, and in this case the utterance is true only if John ended up at the gym and also performed the relevant activity (overseeing the construction of a bathroom).

The third example is provided with the sentence

- (10) That's a dangerous dog

(CL 2005: 46–7). In the first context, the conversation is concerned with the disposition of the dog, and is true if the dog is aggressive, false if it is gently disposed. In the second context, the dog is gentle but carries a viral disease that can spread to humans, and in this case the utterance is true.

After giving these and a few more examples, CL think that the reader is able to go on by herself to produce scenarios for context shift for arbitrary sentences. Reflecting on this, one can conclude that any sentence, and thus any expression, is context-sensitive. One must then be a Radical Contextualist, if you give any credence at all to CSAs.

2.2. Incompleteness

The Incompleteness Arguments draw on intuitions that context must contribute something to what is said, since the sentence itself does not have a content that is truth evaluable. CL sum up the nature of Incompleteness Arguments as follows:

A typical Incompleteness Argument, as we think of it, comes in two stages.

Stage 1. A solicitation of an intuition to the effect that the proposition semantically expressed by an utterance of a sentence *S* (according to Semantic Minimalism) is incomplete, i.e., it is not the kind of thing that can take a truth value.

Stage 2. A solicitation of an intuition to the effect that utterances of *S* have a truth value, i.e., that they can express propositions, and hence, do have truth conditions, and so, can take truth value.

Conclusion. Something unaccounted for by Semantic Minimalism must be added in the context to the utterance in order for a complete proposition to be semantically expressed. (CL 2005: 59)

CL add (2005: 60) that we get a more comprehensive argument by adding context shifting intuitions to the conclusion that what is added for getting a complete proposition also shifts between contexts.

Examples of sentences for which incompleteness is claimed include (5) ('It's raining'). It is also claimed by Kent Bach for

(11) Steel isn't strong enough

(from Bach 1994*b*: 269; CL 2005: 34), since the intuition is that to express a full proposition by means of (11), it must be settled in the context of utterance what it is meant that steel isn't strong enough *for*. One naturally asks 'Strong enough for what?'

In a similar vein, Dan Sperber and Deirdre Wilson claim that

(12) Peter's bat is gray

is less than fully propositional, for the possessive construction might refer to many different relations between Peter and the bat; it might be the bat owned by Peter, the bat chosen by Peter, the bat killed by Peter and so on (CL 2005: 35; Sperber and Wilson 1992: 188).

CL deal with the Incompleteness Argument in a way analogous to the way they dealt with the CSAs, by arguing that Incompleteness is ubiquitous—if you allow it to affect the meaning anywhere, it must be allowed to affect it everywhere. And this once again leads to Radical Contextualism. First, the examples they provided for showing context shift in (4), (9), and (10) can also be used for

showing that those sentences are incomplete as they stand. For instance, one can ask the questions about (9)

Went to the gym how? Walked to the vicinity? Did something in the gym? Did what in the gym? For how long? [etc.] (CL 2005: 64–5)

So, these allegedly complete sentences seem to be incomplete by the same criterion that further questions can be asked and further information provided in the context of utterance.

Moreover, further questions can also be asked with respect to alleged completions of the incomplete sentences, such as

- (1b) Tipper is ready for the exam.
- (11b) Steel isn't strong enough to support the roof.
- (12b) The bat owned by Peter is gray.
- (5b) It's raining in Palo Alto.

(CL 2005: 62, different numbering). As an example, they take (11b), and claim that it fails to specify truth conditions for (11)

because it doesn't settle for how long the support must last. Do a few seconds suffice? More than three days? Many years? Why mustn't (11b) also settle whether (11) is false if steel fails to support the roof when placed in temperatures over 390° [etc.]? (CL 2005: 63)

In short, since it seems that further factors can always be added that can settle in context whether the utterance of a sentence is true or false, it appears that if the Incompleteness Argument is good, *any* sentence can with some ingenuity be seen as incomplete, and therefore also context-sensitive.

Before turning to the evaluation of CL's arguments, we shall sketch a framework for handling different kinds of pragmatic additions to semantic content.

3. RECANATI AND SYSTEMATIC THEORY

A natural methodology for justifying a semantic theory *S* of a natural language *L* is to see *S* as part of a more comprehensive theory *C* of communication by means of *L*. In each act of successful communicative exchange, one or more thought contents get conveyed by the speaker to the hearer. Speaker and hearer use their communicative abilities to achieve this communicative goal, and part of those abilities consist in making use of properties of the language *L* itself. A semantic theory of *L* will then be concerned with special properties of expressions of *L* by which certain expressions are apt for conveying contents of particular kinds. On such a general strategy, semantics plays a systematic role in a more comprehensive systematic theory of communicated content, or speech act content.

Such a methodology is completely rejected by CL, because of their belief, which we have already mentioned:

(N) *There can be no systematic theory of speech act content*

(CL 2005: 190). By means of a single utterance of a sentence *s*, a speaker says indefinitely many things, including things she is not aware of and does not even believe (CL 2005: 202–3). This is because a speaker says everything she can be correctly reported as saying, and what she can be correctly reported as saying depends on factors of the *reporter's* context that the speaker need not have knowledge of. If there can be no systematic theory of speech act content, then there can be no systematic account of how the meaning of sentences contribute to speech act content.

As we mentioned above, CL do believe, contrary to the claim (N), that there is a correct and systematic partial theory of speech act content, since they hold that the proposition semantically expressed by a sentence *s* is said by means of any utterance of *s*. The speaker *always* says the proposition semantically expressed by the sentence used. This is indeed a systematic theory of speech act content, but not one that is very easy to use to identify what is semantically expressed by a sentence.⁵ Rather, one gets the right result only if one has an independent knowledge of what the semantically expressed meaning is. So there is a challenge for CL to explain how the semanticist arrives at the meaning of any linguistic expression.⁶

Still, if the non-systematicity claim (N) is correct, the first mentioned methodology wouldn't work either, because there could not be a systematic theory *C* of communication. Is (N) correct? That question is a big one, and not easily answered. Moreover, the question is not even precise until it is made reasonably clear when to count a theory systematic. We shall here make the assumption

⁵ This is so, especially since in some cases what is intuitively conveyed by means of a sentence differs from what it semantically expresses. For instance, the literal meaning of the following sentence

- (13) If the elevator stops between two floors, press the alarm button for 20 seconds!

(from a Stockholm University elevator) is such that the injunction is complied with if the elevator stops between the fifth floor and the seventh floor, *on* the sixth floor, and one presses the alarm button for 20 seconds. Discerning the literal meaning from the conveyed meaning (*between two adjacent floors*) is extremely difficult, however. One rather needs an independent grasp of the literal meaning. Once one has that, on the other hand, one can easily see how the conveyed meaning is derived from the literal meaning by means of a free enrichment: the expansion of *two adjacent*, amounting to a restriction of the determiner meaning.

⁶ This is not just equivalent to explaining how one learns a first language: over and above understanding what people say by means of the language, the semanticist must determine what semantic concepts to use and how to apply them, and also be able to justify the choice. The alternative would be to baldly claim to have direct insight into theoretical semantic matters. We take this alternative to be unscientific, or unnaturalistic.

that for literal, non-context-sensitive meaning, the paradigm of systematicity⁷ is compositional semantics: a language has a compositional semantics just if

(PoC) The meaning of a syntactically defined whole is a function of the meanings of its syntactic parts and the mode of composition.⁸

The intuitive idea behind the claim that compositional semantics can explain communicative success is that speaker and hearer are able to construct new complexes by means of putting familiar parts together according to familiar patterns. This idea is beautifully expressed in the opening passage of Frege's 'Compound Thoughts':⁹

It is astonishing what language can do. With a few syllables it can express an incalculable number of thoughts, so that even a thought grasped by a human being for the very first time can be put into a form of words which will be understood by someone to whom the thought is entirely new. This would be impossible, were we not able to distinguish parts in the thought corresponding to the parts of a sentence, so that the structure of the sentence serves as an image of the structure of the thought. (Frege 1923)

The intuitive idea is fairly clear, even though much more is needed to spell it out in detail (cf. especially Pelletier 1994; Pagin 2003). It quickly becomes a lot less clear when we move from literal, context-insensitive meaning to uses of language where context plays an important role. Can the intuitive idea of compositionality be extended or generalized to cover utterance contents that do depend on context in various ways?

In order get a grip on that question we will return to the framework for describing semantic and pragmatic contributions to utterance content developed by François Recanati. As we remarked at the beginning of this chapter, Recanati starts his story with types of "pure literalism", theories of language that have no use for any kind of contextual information, not even indexicality. We need not pause over Recanati's account of why it had no use for indexicality,¹⁰ but instead

⁷ As the reader has already noticed, the participants in this discussion are using 'systematicity' in a non-technical sense, not that of e.g. Fodor 1987; Fodor and Pylyshyn 1988.

⁸ There are many fine points that are being glossed over here in this statement of compositionality. Some are discussed in Pelletier 1994, others in Pagin 2003, 2005*a*, and still others in Westerståhl 1998, 2004, and Hodges 1998, 2001. We will here mainly leave the exact force of compositionality unspecified, until we need to discuss some specific point.

⁹ The general sentiment about the understanding of new sentences and about the infinity of language was much in the air at the time and can also be found in Schlick 1985; Russell 1956; Wittgenstein 1981; all first published between 1918 and 1921. Although some have said that in this quote Frege is *not* trying to explain communication, but rather is using the presumed fact that communication is successful together with the considerations of this paragraph to conclude that Thoughts must be structured, to us it seems a more natural interpretation to say that he *is* concerned with communication—especially when taking into account that such a view is present in the contemporaneous works just cited.

¹⁰ Because, according to Recanati, these pure literalists thought that every statement containing an indexical element *said the same thing* as some other statement with no indexicals in it.

we will remark that this is the sort of language that accommodates semantic compositionality very easily. Here the idea is that each of the finitely many primitive elements of the language¹¹ has a meaning assigned to it, and each of the finitely many syntactic rules that can be used to combine simpler components into longer ones has some specified semantic effect on the meanings of these simpler components.

The next step is to start taking context into account.

3.1. Saturation

Literalism can be modified to allow “pieces of context”, according to Recanati. This can be done in two radically different ways, *saturation* and *modulation*, together called “primary pragmatic processes”. Saturation is the process of adding semantic values to various parameters associated with simple expressions, so as to get a full proposition, i.e. a truth evaluable entity. It is done in linguistically controlled and mandated ways (Recanati 2004: 7–10). An utterance of

(14) I am in pain

needs saturation by context to fix the reference of ‘I’ and the time associated with the present tense in order to have a proposition expressed that is true or false. Similarly,

(5) It is raining

needs a time value. According to some theorists, like Perry (1986), (5) also needs a location value, while others, including CL as well as Recanati, disagree.

It is clear that, by Recanati’s standards, all the indexical expressions in CL’s basic set induce the need for saturation. For instance, both CL and Recanati would count

(15) Yesterday, Phil met a foreigner

as in need of saturation both with respect to time, because of ‘yesterday’, and with respect to the implicit relatum of ‘foreigner’ (foreign to what?), as both ‘yesterday’ and ‘foreigner’ are in the basic set. CL disagree with Recanati and Bach (and us) about a sentence like

(11) Steel is strong enough

which Recanati and Bach but not CL think is in need of saturation.

¹¹ Let’s just call them ‘words’ for simplicity, and not worry about subparts of words that might be meaningful and longer phrases that might be primitively meaningful (i.e. idioms). We will consider some of these ideas later.

Setting aside the issue of which sentences are in need of saturation, the present question is whether context sensitivity in the sense of sensitivity to contextually determined saturation can be treated in a systematic manner, i.e. in general accordance with the paradigm that compositionality offers for non-context-sensitive literal meaning. The answer is clearly yes, and the general method, first introduced by Richard Montague (1970*a*, 1970*b*, 1973), and further developed in particular by David Kaplan (1989), is that of distinguishing between two levels of meaning: one a context independent level which is of a functional nature and takes arguments from context, and the other a context dependent level, to which the resulting values of the functions for those arguments belong. In Kaplan's case the higher, functional level is *character*, and the lower level is *content*. In short the *character* of a sentence like (14) is the meaning that belongs to it before saturation is performed, while its *content* is the lower level meaning that results from contextual saturation, giving a proposition about who is in pain when.

It is clear that such a semantics does or at least can conform to the paradigm of compositionality, since the semantics can have the general property (Principle of Contextual Compositionality) that

- (PCC) For any complex expression and context, the meaning-in-context of the complex expression is a function of the meaning-in-context of its parts and the mode of composition.¹²

It should be noted, though, that this is not just a trivial extension of the basic idea of compositionality. For (PCC) can fail in two ways: On the one hand, two expressions can have the same meaning-in-context in a context *c* while substituting the one for the other in a sentence *s* does not preserve the meaning-in-context of *s* in *c*. This would be the case e.g. if synonymy substitution isn't meaning-preserving in belief contexts:

- (16) a. Alfred believes he is a pediatrician
b. Alfred believes he is a child doctor

where, on some theories, (16a) can be true and (16b) false if Alfred believes that 'pediatrician' denotes something other than child doctors.

¹² One can imagine that the semantic effects of the syntactic rules might also contain variables that are assigned by context, so that they can have different outcomes in different contexts. While we think this is a live possibility, and might even be useful in accounting for hyperbole and sarcasm, for example, we will not follow this up in our discussion and will instead always talk about the effect of the context just on the meanings of the words and expressions, and not on the rules.

On the other hand, it may be that all the parts of a sentence s have the same meaning-in-context in contexts c and c' while s itself does not. Some theories assume that

(5) It is raining

has no part that specifies a location, but that the sentence as a whole does. In such a theory, two utterances of (5) at the same time t but at different locations l and l' are such that all the parts have the same contextual meaning but the uttered sentence as a whole has different contextual meanings in the two contexts. That would be because the difference in location isn't articulated in any part of the sentence, in accordance with a doctrine of unarticulated constituents.¹³ Because of these different variations, we have a distinction between two different notions of contextual compositionality that does not have any counterpart in the context independent case.

Recanati remarks that some of the semantic values of the indexical expressions are determined by speaker intentions, and are not antecedently given, as is suggested by the locution "determination by context". He mentions that the values of demonstratives are really set by what counts as salient and what the speaker intends. Even the value of 'here' and 'now', he says (Recanati 2005: 174), are claimed to be "highly sensitive to speaker's intent". But this, even if true, is not relevant to the particular point under discussion: whether a language with indexicals and demonstratives can in principle be given a compositional semantics. For that issue only concerns the dependence of complex-meanings-in-context on part-meanings-and-structure-in-context, not the mode of determination of part meanings, as we will further discuss in §6.

There is, however, a further question about the order of composition and value assignment. On one view, arriving at the meaning-in-context of a sentence is a bottom-up process where contextual values are assigned to simple expressions, and then the resulting meanings-in-context of simple expressions are combined to reach to meaning-in-context of the syntactic complexes.¹⁴

The immediate alternative to the bottom-up view of meaning-in-context is the view that we first combine *unsaturated* meanings of parts into a total unsaturated meaning of a complete sentence, and then in a *second* stage assign values to the parameters in the sentence to arrive at the meaning-in-context of the entire sentence.¹⁵

¹³ These matters are discussed in detail in Pagin 2005a. There the location sensitivity is assigned to the meaning of 'rain'. With respect to such an account, there is no context shift failure.

¹⁴ King and Stanley (2005) claim that this is the intuitively correct picture. See also Reimer (2002).

¹⁵ On some views, not all parameters need be given values for a complete proposition to be formed. On these views, such as *temporalism* with respect to time, complete propositions are true

This difference of order does not in itself have much significance, so long as all the operations are defined, as we discuss in §6. However, when Recanati criticizes the bottom–up view, he has a further difference in mind:

Contrary to what formal semanticists tend to assume, the (intuitive) truth-conditions of our utterances are not compositionally determined by the meanings of words and their semantic arrangement, in a strict bottom–up manner. They are shaped by contextual expectations and world-knowledge to a very large extent. That is true of all utterances, however ‘literal’ they are (in the ordinary sense). (Recanati 2004: 81–2)

Recanati’s alternative view is not just that we don’t start by assigning contextual values to simple expressions, but that the *way* a hearer does assign values to parameters proceeds by way of global considerations. The hearer aims at an overall interpretation that makes good sense, and selects contextual values so as to achieve that goal. For instance, to modify an example from Barbara Partee (1989: 275), consider

- (17) a. I admire central Europeans. They all speak foreign languages.
 b. I am completely lost among central Europeans. They all speak foreign languages.

The intuitively suggested interpretation of (17a) is that of assigning the central European as the relatum of ‘foreign’, while in (17b) it is the speaker of the discourse. Considerations of this kind seem to be included in what Recanati refers to as “top–down” processes.

However, from the point of view of the compositionality of semantics, it does not really matter whether this saturation process is bottom–up or top–down, or some kind of mixture, perhaps differing from case to case. Even when selecting values in a top–down fashion, the hearer can be seen as being guided by the goal of getting a reasonable outcome precisely *by the compositional semantics*. It is because selecting the discourse speaker as relatum to ‘foreign’ in (17a) results—by the compositional semantics—in a non-intuitive interpretation that this assignment is rejected. So, preferring a top–down view is not in conflict with compositionality, but rather trades on it. Compositionality extended to context dependence clearly contributes to a systematic account of speech act content.

The reason for this indifference is due to the fact that we are talking about *saturation* and therefore about fixing the value of some linguistic item, when the value is present in the (linguistic or non-linguistic) context. Clearly, if *all* values are assigned to words, then top–down and bottom–up will be the same. A difference might come to the fore were we to countenance saturation of items

or false relative to time points, but still constitute the contents of propositional attitudes. See e.g. Richard 1981.

larger than a word, where this saturation could not be accomplished by saturating some subpart of the larger item. For example, if there were a case where an entire verb phrase needed to be saturated but there was no relevant saturation of any of the contained words, then top–down and bottom–up would be different. And perhaps one might think of analyzing referential uses of definite descriptions in this manner also, thinking that there may be no relevant saturation of the descriptive material, but that this descriptive material plays some procedural role in helping the hearer make the correct identification.¹⁶ But these sorts of abstract possibilities have not been mentioned by those who favor such processes as saturation.

Matters are less clear when we switch over to modulation.

3.2. Modulation

Whereas saturation is mandatory, since it is needed to arrive at a full proposition, modulation is optional. Modulations operate on chunks of interpreted material that don't strictly need to be modulated in order to contribute properly. Recanati (2004) distinguishes between three different kinds of modulation: free enrichment, loosening, and semantic transfer.

In free enrichment, semantic material is optionally *added* to what is derived from word meaning and structure. Typical examples are

- (18) The table is covered with books
- (19) Mary took out her key and opened the door

A normal utterance of (18) would have content like that of

- (20) The table *of our living room* is covered with books

but in the case of (18), the addition, i.e. what is expressed by the italicized part of (18), is not articulated, but tacitly added (Recanati 2004: 10). The addition is not needed to have a proposition; it's just that the proposition literally expressed is, by ordinary conventional standards, absurd. The semantically expressed proposition gets freely enriched by the tacit material.

Similarly, in (19), a normal intuitive interpretation adds to the proposition semantically expressed that Mary opened the door *with the key that she took out*, although this is not necessary in order to have a full proposition, for with saturation added concerning 'her key', 'the door', the time etc., (19) expresses a proposition that can be true whether or not Mary used the contextually indicated key to open the contextually indicated door at the contextually indicated time.

¹⁶ A suggestion made by one of the referees of this chapter.

Another kind of modulation is *loosening* (Recanati 2004: 24), as in

- (21) The ATM swallowed my credit card

where the verb ‘to swallow’ has its application conditions *extended* to include the cash machine process referred to.

A third kind of modulation is *semantic transfer*, exemplified by

- (22) The ham sandwich left without paying

as said by one waiter to another at a restaurant (the example is due to Nunberg 1979). Here the phrase ‘the ham sandwich’ is used to refer to the guest who ordered the ham sandwich, rather than to the dish itself. The semantic value has been transferred from the latter to the former.

Modulation is characterized as a top–down process, where the hearer uses his general understanding of the situation to arrive at the interpretation, as opposed to a bottom–up process, where the additions are triggered by elements of the sentence used (Recanati 2004: 18).

Saturation and modulation are what Recanati calls *primary* pragmatic processes. *Secondary* pragmatic processes are distinguished by taking the result of the primary process as input. The secondary processes comprise various kinds of *implicature*. An example of Recanati’s combining both kinds is

- (23) I’ve had breakfast

said in reply to the question ‘Do you want something to eat?’ (Recanati 2004: 8). The speaker of (23) *implicates* she is not hungry and hence does not want anything to eat. Thereby it exemplifies implicature. It also exemplifies free enrichment, since she communicates not only that she has had breakfast at *some* time or other prior to the time of utterance, which is obviously true in its “literal meaning”, but over and above this that she has had breakfast on the very day of the utterance. It is important that the primary process of fixing the day being discussed (to the date of utterance of (23)) must take place *before* the secondary process takes place because the implicature inference cannot be performed without that information. (The hearer can’t infer that the speaker is not hungry unless she fixes the time that the speaker had breakfast.)

According to the position Recanati calls ‘minimalism’ there is a real communicative level of *what is said* that minimally deviates from the literal meaning of the sentence used. The minimal deviation comes from saturation.¹⁷

¹⁷ There is further view called ‘the syncretic view’ by Recanati, according to which there are two levels (or two notions) of what is said, the minimal one and a further level that results from modulations. Recanati rejects this position as well (2004: 64).

Recanati rejects minimalism, arguing that the only real level of *what is said* is that delivered by the *combination* of saturation and modulation (Recanati 2004: 21).

More crucially, the rejection of minimalism is backed by the claim that the so-called *minimal proposition*, the result of saturation, is *not*, or at least not in general, computed by the hearer. In the example (22), according to Recanati, the hearer does not first interpret the utterance as an assertion of the absurd proposition that the dish itself left without paying, as minimalism implies, but achieves the semantic transfer locally, i.e. pragmatically reinterprets the phrase ‘the ham sandwich’ before the combination with the interpretation of the verb phrase is performed (Recanati 2004: 30–1).

Accordingly, on Recanati’s picture, semantics, or literal meaning, normally contributes nothing but word meaning to the interpretation process. The interpretation of syntactic complexes takes place on a *pragmatic* level, with *modulated* contents. Modulated contents are composed into contents of larger syntactic units, and this is a purely pragmatic composition. The picture is one of a “bag of word meanings” that gets pragmatically modified and pieced together. In the even more radical picture of meaning eliminativism, words don’t really have any meanings at all but only a “semantic potential” (Recanati 2004: 152), and some sorts of “associations” that have been built up by the use of these words in past situations (p. 151). In both cases, semantic compositionality is irrelevant to the communication process, and in the latter case semantics has no role at all; communication rather resembles an “association of ideas” *à la* the old British Empiricists, or to its modern reincarnation as pattern recognition in a connectionist network.¹⁸

As we see, then, Recanati accepts “more context” in his minimal propositions than what CL accept as being semantically expressed (their minimal propositions). Nonetheless, he rejects the idea that even his minimal propositions play any role in the communication process, except in very unusual circumstances.

Recanati and CL thus agree that semantics cannot play any *systematic* role in accounting for speech act content. According to Recanati, so much the worse for semantics; according to CL, so much the worse for the theory of speech act content.

We disagree with both parties. There is still room for a systematic account of speech act content, where compositionality plays a central role. We shall now proceed to indicate how.

¹⁸ To make our position clear: we do agree with Recanati about the importance of the saturation/modulation distinction, which as he says is a distinction between what is and what is not linguistically controlled. For example, we would find it completely implausible that there be a *linguistic* parameter of ‘ham sandwich’ that takes *orderer* as value in (22), ‘The ham sandwich left without paying’.

4. A SEMANTIC–PRAGMATIC FRAMEWORK

We wish to lay out the structure of a theory that satisfies systematicity by obeying the intuitive content of the Principle of Compositionality and yet also allows for some general theory of language that Recanati would say was “between” pure literalism and radical contextualism (or even his less radical “full-bodied contextualism”). That is, we wish to outline a theory that CL would call a “moderate contextualism”. We take our task to have two parts. The first is simply to stake out such a theory, and the second to defend it against CL’s instability arguments.

We are moved to this attempted middle ground because we think that there is more to “what is said” than is allowed by CL’s minimalist theory, but that “what is said” cannot be so completely unstructured as Recanati’s “no meaning” theory (nor maybe even his “full-bodied contextualism”), because then there could be no account that would explain how conversational participants can produce or understand novel utterances.

The middle ground can be occupied by an account that integrates semantic composition and modulation. There is reason to think that the two *can* be integrated, for if we look at the examples given of modulations, it seems they are pretty much controlled by semantic structure. First, in the standard modulation examples given, the semantic *category* of the original meaning is preserved under modulation. That is, if the original meaning is a concept of an action, the modulated content will be a concept of an action. If the original meaning is a sortal concept, then the modulated concept will be a sortal concept as well, and so on.¹⁹

Second, the composition function, by which simpler contents are mapped onto more complex contents, is simply taken over from the semantics. If the speaker starts out with a subject-predicate construction, semantically interpreted as a concept applying to an object, the modulated combination will again correspond to the subject-predicate construction, with some concept (perhaps a different one) applying to some object (again perhaps a different one).

Third, the structural *role* of a modulated part in determining the modulated proposition will be the same as the structural role of the unmodulated part in determining the minimal proposition. The enriched interpretation of a

¹⁹ There are partial exceptions to this, as when you say, after Oscar has been accidentally squashed (‘Oscar’ was the proper name given to a tomato), that “There was Oscar all over the floor”. In that case the concept of an object is mapped onto the concept of a stuff; but in this case too, the composition function is taken from the linguistic construction, and the non-category-preserving modulation of an object concept is required by the semantic composition.

grammatical direct object will again be the interpretation of the grammatical direct object, and so on.

These regularities could be explained syntactically, if modulation were a syntactic phenomenon, where a syntactic operation is performed that does not leave any trace on the surface string, but introduces ellipsis at an underlying level. However, *pace* Stanley (2002), we do not find this syntactic hypothesis to be plausible. We see it rather, as on Recanati's picture, that modulations are content operations, and as such they should be seen as controlled by a *conceptual structure*, on which they operate. Those are the immediate reasons for the account below.

The main idea, then, is to combine a more basic immediate output of the semantics—here called a 'conceptual structure'—with pragmatic modulation operations that modify the transition from the conceptual structures to standard propositional representations. We shall operate with the familiar homomorphism format of compositionality from the Montague and Hodges traditions. We shall not set out the machinery with full precision, but only as far as is required to articulate the idea.

To start, we have a language L consisting of a set of atomic expressions A_L , a set of syntactic operations Σ_L , and a set of expressions E_L (we shall drop the subscripts from now on). E is the set that can be produced from A by means of the operations in Σ . Now we have a *meaning function* μ that maps structured syntactic entities, like syntactic trees or terms, onto meanings.²⁰

Together with the meaning function μ there is a *general composition function* ρ . ρ takes as argument a syntactic operation $\sigma \in \Sigma$, and gives as value a *particular composition function* $\rho(\sigma)$. The particular composition function $\rho(\sigma)$ in turn maps meanings of the parts of a complex expression onto the meaning of the complex expression itself.

This gives us the formalized counterpart to (PoC):

$$(PC) \quad \mu(\sigma(t_1, \dots, t_n)) = \rho(\sigma)(\mu(t_1), \dots, \mu(t_n))$$

where the t_1, \dots, t_n are the syntactic trees or terms that are the immediate syntactic constituents of the complex $\sigma(t_1, \dots, t_n)$. $\mu(t_1), \dots, \mu(t_n)$ are of course the meanings of the parts.

To this we can now add context dependence, by letting each constituent have a context argument. A context can be represented formally as a sequence of contextual values, like *time* or *place*, in the simplest cases. Certain operations on contexts, like $T(c)$ to pick out the time of context c , provide what is relevant to

²⁰ It is important that the meaning function is defined on expressions that are syntactically unambiguous. That is why it cannot in general be defined on ordinary expressions, since in natural language ordinary (surface) expressions are often ambiguous.

the expression in question, and for an atomic expression, this will be specified in its basic semantic clause.

With the addition of context (\mathcal{PC}) can be modified into

$$(\mathcal{PC}_c) \quad \mu(\sigma(t_1, \dots, t_n), c) = \rho(\sigma)(\mu(t_1, c), \dots, \mu(t_n, c))$$

where the contextual meaning of the complex depends on the contextual meanings of the parts and the mode of composition.²¹ To have a complete context semantics we would also have to specify how the meanings on the higher, context independent level relate to the contextual meanings, by the normal methods involving λ -abstraction and λ -conversion. Since our aim here is to give a way of handling modulation, we will leave this part out.

Normally, we would take the output of the meaning function to be a standard representation of states of affairs, of what things are like, say a proposition, where we think of a proposition as something without internal structure, a *flat* representation.²² If that is what we think the saturated content of an utterance is, and as well what the saturated and modulated content of an utterance is, then we would want the end-result of the formal interpretation process to be a proposition. But the immediate output of the semantic function need not be a proposition. Rather, for the purpose of providing material for further operations, a *structured output* is more suitable.

What we suggest, therefore, is an alternative meaning function μ' which, instead of delivering the result of applying the particular composition function onto the part meanings, delivers instead the ordered n -tuple of the particular composition function and the (alternative) part meanings. That is, we define μ' as follows:

$$(\mathcal{PC}'_c) \quad \begin{array}{l} \text{(i)} \quad \mu'(t, c) = \mu(t, c), \text{ if } t \text{ is atomic} \\ \text{(ii)} \quad \mu'(\sigma(t_1, \dots, t_n), c) = \langle \rho(\sigma), \mu'(t_1, c), \dots, \mu'(t_n, c) \rangle \end{array}$$

Note that the general composition function ρ' that is thereby defined for μ' is such that for a given n -place syntactic operation σ

$$(25) \quad \rho'(\sigma) = \langle \rho(\sigma), \xi_1, \dots, \xi_n \rangle$$

(with the ξ_i marking the argument places).

²¹ We get a different version of contextual compositionality by adding a context argument to the particular composition function:

$$(24) \quad \mu(\sigma(t_1, \dots, t_n), c) = \rho(\sigma)(\mu(t_1, c), \dots, \mu(t_n, c), c)$$

This weaker alternative corresponds to the idea that the sentence 'It is raining' (see above) can have different contextual meanings in two contexts even if all the parts have the same meaning in both contexts.

²² See Cresswell 2002 for an argument that propositions don't have structure.

If we have a complex expression with iterated applications of syntactic operators, the output of μ' will be a tree structure, where every constituent of an $n + 1$ -tuple, except the first, is the vertex of a sub-tree.

For a simple example, consider the sentence

- (26) a. Alfred loves Elsa
 b. [_S [_{NP} Alfred] [_{VP} [_{VT} loves] [_{NP} Elsa]]]

(where we employ a standard phrase structure as given in (26b)). Let σ_1 be the operation that combines a VT and an NP into a VP, and σ_2 the operation the combines an NP and a VP into an S. Then, given μ and ρ , the output of μ' will be

$$(27) \quad \mu'(\sigma_2(\text{Alfred}, \sigma_1(\text{loves}, \text{Elsa}))) = \\ \langle \rho(\sigma_2), \mu(\text{Alfred}), \langle \rho(\sigma_1), \mu(\text{loves}), \mu(\text{Elsa}) \rangle \rangle$$

(since μ' coincides with μ on atomic expressions). The third element of the top triple is itself a triple.

The output of μ' , a “conceptual structure”, is taken as the primary output of the semantics. The primary output is then transformed into a standard output by means of an *evaluation function* E that takes a conceptual structure as argument and gives a standard interpretation as value. The definition of E is straightforward:

$$(E) \quad (i) \quad E(\mu'(t)) = \mu'(t), \text{ for atomic } t \\ (ii) \quad E(\langle \rho(\sigma), \mu'(t_1, c), \dots, \mu'(t_n, c) \rangle) = \rho(\sigma)(E(\mu'(t_1, c)), \dots, \\ E(\mu'(t_n, c)))$$

So the application of E to an $(n + 1)$ -tuple simply consists in applying the first element of the tuple, which is an n -place function, to the E -values of remaining n elements. By an elementary induction proof it can be shown that

$$(28) \quad \mu(t) = E(\mu'(t))$$

for any term t for which μ is defined.

Although it seems that this is merely an alternative way to arrive at the same result that a standard semantic theory would generate, the point of introducing conceptual structures and their evaluations into ordinary interpretations is that we can identify the evaluation step as the step where pragmatic modulations enter the process. To illustrate, the *modulated evaluation*, E_m , of an utterance of (22)

(22) The ham sandwich left without paying

in a context c can be said to be achieved by means of applying the function *the orderer of* . . . to the part of the conceptual structure that is given by μ' (the ham sandwich). Skipping the syntactic details below the NP-VP structure, we have conceptual structure with the topmost triple

(29) $\langle \rho(\sigma_2), \mu'(\text{the ham sandwich}, c), \mu'(\text{left without paying}, c) \rangle$

We then have the contextual modulated evaluation:

(30) $E_m(\langle \rho(\sigma_2), \mu'(\text{the ham sandwich}, c), \mu'(\text{left without paying}, c) \rangle, c)$
 $= \rho(\sigma_2)(O(E_m(\mu'(\text{the ham sandwich}, c))), E_m(\mu'(\text{left without paying}, c)))$

where O is the *orderer of* function.

We can here see the modulation as an operation that *replaces* the standard evaluation function by a modified evaluation function E_m in the step from conceptual structure to standard interpretation. An alternative is to give a slightly more liberal definition of the standard evaluation function E so that modulation enters as a function from immediate conceptual structure to modulated conceptual structure. The idea is to add a third schema to the definition so that

- (E') (i) $E(\mu'(t)) = \mu'(t)$, for atomic t
 (ii) $E(\langle \rho(\sigma), \mu'(t_1, c), \dots, \mu'(t_n, c) \rangle) = \rho(\sigma)(E(\mu'(t_1, c)), \dots, E(\mu'(t_n, c)))$
 (iii) $E(f(\mu'(t))) = f(E(\mu'(t)))$

for arbitrary non-semantic function f . Then we can say instead that

(31) $E_m(\langle \rho(\sigma_2), \mu'(\text{the ham sandwich}, c), \mu'(\text{left without paying}, c) \rangle, c)$
 $= E(\langle \rho(\sigma_2), O(\mu'(\text{the ham sandwich}, c)), \mu'(\text{left without paying}, c) \rangle)$

If we suppose that E_m in context c gives us all the modulation information, i.e. the information of all the modulating operations performed on the conceptual structure, we can provide the general format as follows:

- (E_m) (i) $E_m(\mu'(t)) = \mathcal{M}(E(\mu'(t)))$, for atomic t
 (ii) $E_m(\langle \rho(\sigma), \mu'(t_1, c), \dots, \mu'(t_n, c) \rangle, c)$
 $= \mathcal{M}(E(\langle \rho(\sigma), E_m(\mu'(t_1, c), c), \dots, E_m(\mu'(t_n, c), c) \rangle))$

where \mathcal{M} is the modulation performed at the topmost level of the structure or substructure under consideration. \mathcal{M} might be the *null*-modulation, in case no non-trivial modulation is performed at that level.²³ Each application of E_m gives an output of the same format, with an initial modulation, possibly the null-modulation, followed by applying the basic evaluation function E to the following structure.

In this way a well-defined structure is the immediate output of the semantics, and the primary pragmatic processes (as opposed to implicatures and other secondary operations) can be defined as operations on this semantic output. This shows that acknowledging modulations as primary pragmatic processes isn't in conflict with accepting a central role for compositional semantics.²⁴

It is important to keep in mind that although the E_m function applies in the familiar recursive way, there is as yet no good reason to regard it as a *compositional* function in the ordinary sense. For this to be adequate, we would need to treat E_m as a *universal* function that delivers a particular modulation \mathcal{M} as value for particular conceptual + contextual arguments. So far, however, we have no conception of the *contextual elements* that are apt to trigger, for example, the semantic transfer modulation *the orderer of*. Rather, it seems to be a matter of global interpretational features of the context (i.e. what Recanati calls a top-down process) that cannot easily be coded as contextual elements. We will return in §6 to the question of whether this precludes a *systematic* account of the choice of selection of modulation function in the interpretation process.

5. THE INSTABILITY ARGUMENTS REVISITED

In the previous section we showed that modulations can be happily combined with a compositional semantics. We shall now try to show that a consistent application of the semantics/modulation distinction will defuse the CL instability arguments.

So consider first our man Smith

(4) Smith weighs 80 kg

²³ There is a complication that has to be accounted for somehow. Considering a sentence like

(32) The ham sandwich liked the ham sandwich

we would want an interpretation where the interpretation of the first occurrence of the definite description gets modulated by the *orderer of* function while the second doesn't. So it is not just the overall E_m -function and the argument that determines the modulation. Either we need a more fine-grained notion of context, or else a more fine-grained notion of a modulated evaluation function. We shall here leave the difficulty unresolved.

²⁴ That is the main point of Pagin 2005*b*, where this way of combining compositional semantics with modulations is first suggested.

and let's assume that 'weighs' is a genuinely present tensed transitive verb, true of a pair of an object and a pair of number and weight unit, at a particular time.²⁵ In that case we have a straightforward true application in the cases when Smith weighs

(4a) 80 kg when stripped in the morning

if the sentence is uttered in the morning and intended to concern the time of utterance. The same holds in the cases where Smith weighs

(4c) 80 kg after being force fed 4 liters of water

(4d) 80 kg four hours after having ingested a powerful diuretic

It might of course be the case that the speaker who makes the utterance of (4) at a time when Smith has been force fed 4 liters of water does intend Smith's normal morning weight. But that is easily treated as a case of modulation. For instance, the speaker might intend the addition of temporal quantification, together with a replacement of a reference to the present time by a morning-restricted time variable, perhaps with a further free enrichment concerning eating conditions:

(4a') [It is normally the case that] Smith weighs 80 kg [in the morning before breakfast]

where the added material is represented within square brackets. This result can be achieved by means free enrichment and perhaps also semantic transfer (for the time reference). In the case of

(4e) 80 kg after lunch adorned in heavy outer clothing

the result is again easily achieved by means of a modulation that can be classified in various ways:

(4e') Smith [together with his clothes] weighs 80 kg

And so on. The various propositions that might normally be intended by (4) are fairly easily accounted for by means of reference to modulations. These modulations account for context shift features, as well as for a possible impression

²⁵ We leave it open whether there is also a separate lexically coded *habitual* or *iterative* present tense, i.e. a coded aspect relating a pair not at a particular time but with respect to some time interval, and true of the pair if the number gives something like the average weight of the object during the interval.

of incompleteness that might be evoked by the context shifts. The extra material does not add to something that is inherently incomplete, but, as Recanati says, maps a (complete) proposition onto another proposition.

Let's consider the case of

- (9) John went to the gym

and assume that standardly, 'went to' is true of a pair consisting of a moving object and a place, at a time, just in case the the object at the time had traveled along some trajectory and reached the boundary of the place. Then, when the topic of conversation is John's walking habits, the content of the utterance would be straightforwardly accounted for by another enrichment

- (9') John went to [the vicinity of] the gym

In the other two cases, where we are concerned with John's exercise or construction work, then no modulation is needed. It is simply an implicature under the Gricean maxim 'Be relevant' (Grice 1989: 27). If the topic of conversation is John's work activities, and it is common knowledge that construction work is going on at the gym, then we have a violation of the relevance maxim if the speaker simply meant that John went to the gym to work out or admire the building. The speaker might have meant literally that John reached the boundary of the gym, without having any further information, in which case any stronger claim would violate the quality maxim 'Do not say that for which you lack adequate evidence' (Grice 1989: 27). Again, the further interpretations are easily derived. In this case, too, both context shift and the accompanying impression of incompleteness are accounted for.

The case of

- (10) That's a dangerous dog

is even easier, since the speaker need not even intend to inform the hearer, *by means of that particular utterance*, of the *explanation* of why the dog is dangerous. Similarly, when CL ask about (9) 'Went how to the gym?', this concerns further information about different ways (9) might be made true. No pragmatic process, neither primary nor secondary, is needed. It is not a case of propositional incompleteness at all.

These examples contrast strikingly with the proposed paradigmatic examples of incompleteness

- (1) Tipper is ready
(11) Steel isn't strong enough

where there simply is *no* candidate available for being the proposition expressed without modulation or implicature. We simply don't know what it is to be ready, or strong enough, *simpliciter*. It cannot be that (1) is true just in case a condition obtains that is common to all the contexts in which it is correct to assert (1). For what is common is that there is something for which Tipper is ready. Then, by parity of reasoning,

(33) Tipper is not ready

is true just if there is something for which Tipper is *not* ready. By these explanations it would hold that Tipper is both ready and not ready, in case there is something for which Tipper is ready and something for which Tipper is not ready.²⁶

CL's defense of the completeness claim is the instability argument itself, i.e. the argument that the intuitions by which (1) and (11) are deemed *incomplete* can be used equally well against the alleged completions (CL 2005: 62–3). We have by means of a number of examples tried to make it plausible that this strategy does not work, as the various cases differ in important respects.

Some of CL's examples are of course problematic. As regards the sentence

(11b) Steel isn't strong enough to support the roof

CL are right to say (2005: 62) that it is not so clear what the truth conditions of an utterance are meant to be. As they ask, for how long, and under what condition? Nevertheless, it is clear that at a time *t* such that there is some temporal neighborhood of *t* under which the roof is supported by a steel construction,²⁷ the sentence

(34) Steel supports the roof

would be literally true. By the same token, with respect to *t* it would be equally true to say

(35) Steel is strong enough to support the roof

No completion is necessary. As a matter of fact, we would not use (11) to make nor interpret an utterance of (11) as making such a weak claim, presumably because it would violate a relevance requirement. What is relevant is rather that there is support during a foreseeable future, and under fairly bad but not necessarily extreme conditions. Of course it is vague what to count as foreseeable

²⁶ A similar point is made in Borg 2006: §iii.

²⁷ That is, the construction undergoes only negligible decay during that time span.

and what to count as extreme conditions, but that is beside the point, since vagueness is a separate problem. The relevant issue is rather that the qualifications about time span and conditions can be seen as added by way of modulation.

All in all, the appeal to modulation, and sometimes even implicature, seems adequate to deal with the examples that are used by CL to show that any sentence can be seen as semantically incomplete and context dependent. As far as we can make out, the CL instability arguments fail.

6. CONCLUDING REMARKS

As we have remarked at various places in this chapter, there are further issues that are relevant to our attempt to employ semantic compositionality in our explication of how we might arrive at a moderate contextualism. Our professed strategy, recall, was to invoke systematicity, but to explain that notion by using compositionality as a prototypical case of systematicity. (Cf. n. 7 above.) It is time now to consider just how well we think compositionality fares on its own as an account of the systematicity that is required to counter the instability arguments that are used by both CL and Recanati in their denials of the viability of a moderate contextualism.²⁸

We start with some comments about compositionality, briefly surveying the extent to which it is successful by trying to give as strong a case as possible for employing compositionality in this realm. In the end we decide that compositionality, strictly understood, is not adequate for the job, and that we need to make an appeal to some wider notion of systematicity of the sort employed in §4. We follow this with a more general discussion on the prospects for a systematic theory. Further comments on this can be found in Pagin (2005*b*).

6.1. Reflections on Compositionality

One central topic concerning whether compositionality can be the entire content of systematicity is the issue concerning the appropriate manner to accommodate contextual information even when it is believed that the information *is* susceptible of a compositional account—a topic we considered in subsection 3.1. Consider for example the “first step” along Recanati’s “ever more context” pathway—explicit indexicals. In such a case we have a sentence containing some “ordinary” words that have their ordinary meanings and some indexical words that have a “meaning” but still need to have a semantic value specified for them. One way to proceed here, a way suggested by Kaplan’s character–content

²⁸ CL’s argumentative strategy was to endorse the radical contextualist’s employment of the Incompleteness and Context Shifting arguments. Both sides here deny that there is a stopping point that can give moderate contextualism a home.

distinction, would be to compute the meaning of the sentence in accordance with a compositional mechanism (giving us the character of the sentence) and then to specify the content of the sentence by letting context specify the values of the indexicals (by “saturation”). Compositionality here works on character, or as one might say, on the literal meaning of the words and phrases: the literal meaning of the sentence (its character) is compositionally derived from the literal meanings of its parts and the syntactic manner in which these parts are combined. Saturation then occurs at the sentence level to generate an occasion meaning.

But another way could be to saturate the indexical words—giving them their occasion meanings²⁹—and thereupon employing the compositional mechanism on these occasion meanings to generate the occasion meaning of the sentence. Here compositionality takes the occasion meanings of the parts and compositionally generates the occasion meaning of the sentence.

Although we discussed this sort of case in subsection 3.1, where we distinguished a Principle of Compositionality (PoC) from a Principle of Contextual Compositionality (PCC), we think that a few more words are in order. Even given the sort of simplicity we are imagining in the realm of indexicals (and demonstratives?), the two compositional methods are not equivalent. In subsection 3.1 we mentioned issues of substitution and of “unarticulated constituents”. A question for compositionality then is to determine whether these sorts of differences lead to any important differences in accounting for communication.

The example given above concerns explicit indexicals. One question is whether this can be extended to other words that a moderate contextualist might wish to embrace, such as *foreigner* and *enemy*.³⁰ As we explained in 3.1, there is an important difference between seeing the compositionality principles as *descriptive* characterizations of the language (as made by a linguist, say) versus as a *psychological* mechanism employed by the participants of a communicative act. In the former case, the character–content distinction could once again be invoked, with (it seems) pretty much the same results as we found for the indexicals. But the case is somewhat different when considered as a psychological mechanism employed by conversational participants. In this interpretation it is natural to think of the participants as engaging in Recanati’s “top–down processing”; and some theorists (including Recanati) seem to think that this cannot support compositionality, with its commitment to literal meaning. However, even here (as we noted in 3.1) we can view the psychological interpretation as being guided by trying to employ the descriptive compositional semantics.

And so a moderate contextualist who wished to rely solely on compositionality would wish to make the descriptive–psychological distinction in the case of *all* context-sensitive words (presumably even indexicals). As we said in 3.1, there is no barrier to a conversational participant using top–down processing or

²⁹ In this picture, the occasion meaning of non-indexical words would be their literal meaning.

³⁰ Or *different*, as in the Partee example we mentioned in 3.1.

bottom–up or a mixture of the two, so long as the descriptive compositional structure is in place.

As stated, the PCC method is only defined for cases where an individual *word* is to be saturated, as with the indexicals, or tense, or the other context-sensitive words. We earlier raised the question of whether saturation ever *needs* to be done on units that are larger than a word, where this effect cannot be achieved by saturating any individual words in the unit. For example, can there be noun phrases or verb phrases that require saturation to get their occasion meaning but where there is no subcomponent of that phrase which can be saturated to get the same effect. In 3.1 we gave the example of an utterance of *It is raining*, which some might argue needs to be saturated by a location, but also argue that there is no item in the sentence that can be thus saturated. (The question before us now is not a sentence saturation, but of a phrasal saturation that has no component word saturation.) It is not so clear to us that the *is raining* (or: *to rain*) really can accept a location saturation, but we acknowledge that this sort of idea remains a possibility that could be developed in a more complete theory.

The preceding was a defense of using “pure compositionality” as an account of how saturation might be dealt with. But a compositionalist might have something different to say concerning Recanati’s use of the various types of modulation, but especially the processes of *loosening* and *enrichment*. As a first pass, a committed compositionalist would deny that certain of the argumentative strategies that Recanati employs (and which CL then take advantage of in their Instability considerations) are even required. Consider, for example

- (18) The table is covered with books
- (19) Mary took out her key and opened the door
- (21) The ATM swallowed my credit card

where Recanati claims that “what is said” by (18) includes *table in our living room*, that “what is said” by (19) includes that *the door was opened by the key* (both cases of enrichment), and that the meaning of *swallowed* in (20) has been extended from its literal meaning. The committed compositionalist might dispute these presumed data, and claim that they should instead be seen as cases of *generality*. Why, the compositionalist might cry, should we think that ‘swallowed’ in (21) has been “loosened”? What has been added to the dictionary meaning *to envelop or take in*? If Recanati has at the forefront of his mind something about *taking through the mouth and esophagus into the stomach* when he utters ‘swallow’, the correct response by the compositionalist will be that this confuses generality with difference in meaning. The sentence

- (36) Kim Johnson is a fine person

simply does not have two different meanings, depending on whether the speaker “has a picture in mind” of a woman or of a man. The noun ‘person’ is *general* and applies to both. Why, a compositionalist asks, is the case with (21) any different? Similarly, why would one wish to say that ‘the table’ in (18) has been enriched? Although the speaker may “have in mind” the table in their living room, this can seem irrelevant to the compositionalist. (And in fact it does seem irrelevant to some such theorists, as in Lepore (2003).) The fact that sentences like (19) allow for cancellation of the *with the key* seems to show that this is not a case of enrichment (a primary pragmatic process, in Recanati’s terms) but rather an implicature (a secondary pragmatic process).

At any rate, those are the sort of things that a committed compositionalist might appeal to in an attempt to use compositionality as the touchstone for constructing a moderate contextualism. We think, however, that in the “strict” reading of compositionality there are some issues that need further embellishment, and it was for that reason we shifted our attention to the wider notion of systematicity. So, we should turn to the question of how well such a theory can be expected to perform in our search for a moderate contextualism.

6.2. The Prospects of a Systematic Theory

By showing how to integrate a compositional semantics of the sort indicated in §4 with primary pragmatic processes we have shown how it is possible to account for speech act content, up to secondary pragmatic processes, in a way that is systematic in one respect: it shows for a particular interpreted utterance how the interpretation is arrived at by means of the contextual meaning of parts, the mode of composition, and the modulation function. However, in another respect the account is not fully systematic.

Compare standard compositional semantics. The interpretation of a complex expression is determined completely by the syntactic build-up of the sentence and the meanings of the atomic expressions. There is scope for variation to the extent that there is semantic or syntactic ambiguity, for then the interpreter needs to choose between different syntactic/semantic readings. In this case, however, all the possible readings can be extracted from the syntax and semantics alone.

When we add indexicals, a further element of uncertainty is added. So-called ‘pure’ indexicals, in Kaplan’s sense, like ‘I’, ‘you’, ‘now’, normally get their contextual reference determined by a simple function from the context, for example, ‘I’ is assigned *the speaker of* the context. It is more difficult with demonstratives, including demonstrative uses of pronouns. For example, in case of an utterance of a sentence like

(37) She was late yesterday

the hearer needs to figure out which female the speaker has in mind. That will be the female that is most *salient* for the speaker at the time of utterance, which may and may not coincide with the female that is most salient for the *hearer* at that time.

In the case of possessive constructions like ‘Peter’s bat’ in (12) above, the hearer needs to figure out which *relation* between Peter and the bat is the most salient one.

As far as we are aware, there is so far no general and successful theory that tells us how to *predict* what will be most salient to a speaker or hearer on an occasion.³¹ We do have some simple rules of thumb, like the rule that things become salient by being mentioned, and by being visually or audibly different from their environment in conspicuous ways, but not much more than that. To the extent that we cannot predict the salience profile of a context from more overt elements, there is a limit to our ability to account for communicative success. We may know that usually the same things are salient to both speaker and hearer, but lack a general theory for predicting what it will be.

Still, when it comes to interpreting an utterance of (37), we can say in general, for example, the utterance is *true* provided the most speaker-salient female of the context of utterance was late to something or other on the day before the day of the utterance. That is, we know exactly how the interpretation depends on the salience profile of the context. The fact that Elsa (say) is the most speaker-salient female can be treated as one element of the context, similar to the fact that, say, Alfred, is the *speaker* of the context. Since we know how context-sensitive utterances like (37), or

(38) I was late yesterday

depend on the contextual parameters, we can specify the interpretation in terms of them.

In the case of modulations, the difficulties are greater. Given that we know the target interpretation of, for example,

(9') John went to [the vicinity of] the gym

and we know the saturated but non-modulated meaning of the sentence (9), we can account for how to get the former from the latter. But to have a fully

³¹ It does seem that, by default, what is salient to the speaker is what the speaker believes is salient to the hearer, and vice versa. In some cases, which deserve a closer study, a discrepancy arises. For instance, global considerations about reasonable interpretation might motivate a hearer to change her assumption about the identity of the female that is most salient to the speaker. This will make the new candidate (if there is one) most salient to the hearer.

systematic account, we should also be able to predict the interpretation from the context, i.e. predict which modulation to apply *without* making use of independent knowledge of the target interpretation.

To some extent, selecting the right modulation, given some contextual information, is like selecting the right referent of a demonstrative: in both cases some pragmatic skill, or representation of pragmatic skill, is required. But there is also a difference, for the appeal to salience in the demonstrative cases is an appeal to a contextual parameter that is independent of the interpretation itself. That is, to be the most speaker-salient female is a property that depends on speaker psychology, irrespective of interpretation.

By contrast, it is at least not obvious that there is some particular contextual parameter whose value in a context *c* determines the function *the vicinity of* as the right modulation of the interpretation of the NP 'the gym' in (9'). Rather, it seems to be something arrived at by means of general pragmatic skill, taking into account the topic of conversation, the current focus, the established conversational score, and perhaps further features of speaker psychology or speaker idiosyncrasies. It seems to depend on *global* interpretational considerations rather than individual parameters. If this is right, is there any hope of giving a correct, systematic account of the selection of modulations? Or must we ultimately agree with the CL thesis (N) that there cannot be a systematic theory of speech act content? If a modulation simply is a function mapping propositions onto propositions, without any further restriction, there is just no limit to how random and strange modulations can be.

Although this is not the place even to *start* an investigation into the systematic nature of modulations, let us express some optimism. We think that much of the work done in pragmatics in recent decades, for example, within Relevance Theory (Sperber and Wilson 1992) or within the theory of implicature (Bach 1994*a*), contributes to the understanding of these processes. The modulations that do occur seem to be fairly simple transformations of propositions or of parts of the conceptual structures that evaluate into propositions. They typically restrict, enlarge, or offset reference in small doses. They often add reference to assumed *causes* of events that are explicitly referred to.³² Probably, there is a fairly small set of normal *types* of modulation, even if we go beyond the division into loosening, free enrichment, and semantic transfer. Identifying modulations seems

³² This phenomenon was noted already by in Frege 1892, where Frege considers the sentence

- (39) Napoleon, who recognized the danger to his right flank, personally led his troops against the enemy's position

noting that what seems to be expressed includes a third proposition, that the recognition of the danger was the reason for the action.

to be a process motivated by standard considerations of relevance, clarity, and simplicity. There is reason to be optimistic about progress in mapping it out.³³

REFERENCES

- Bach, K. (1994a), 'Conversational Implicature', *Mind and Language*, 9: 124–62.
- ____ (1994b), 'Semantic Slack: What is Said and More', S. Tsohatzidis (ed), *Foundations of Speech Act Theory*, pp. 267–91, London: Routledge.
- Borg, E. (2006), 'Minimalism vs Contextualism in Semantics', in G. Preyer and G. Peter (eds.), *Content and Context: Essays on Semantics and Pragmatics*. Oxford: Oxford University Press.
- Cappelen, H., and E. Lepore (1997), 'On an Alleged Connection between Indirect Quotation and Semantic Theory', *Mind and Language*, 12: 278–96.
- ____ and ____ (2005), *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.
- Cresswell, J. M. (2002), 'Why Propositions have No Structure', *Nous*, 36: 643–62.
- Feigl, H., and W. Sellars (eds.) (1949), *Readings in Philosophical Analysis*. New York: Appleton-Century Croft.
- Fodor, J. (1987), *Psychosemantics*. Cambridge, Mass.: MIT Press.
- ____ and Z. Pylyshyn (1988), 'Connectionism and Cognitive Architecture', *Cognition*, 28: 3–71.
- Frege, G. (1892), 'Über Sinn und Bedeutung', *Zeitschrift für Philosophie und Philosophische Kritik*, 100: 22–50. Tr. Herbert Feigl as 'On Sense and Nominatum', in Feigl and Sellars 1949: 85–102.
- ____ (1923), 'Compound Thoughts', *Mind*, 72: 1–17, tr. R. Stoothof, and published in *Mind* (1963).
- Grice, H. P. (1989), *Studies in the Ways of Words*. Cambridge, Mass.: Harvard University Press.
- Hodges, W. (1998), 'Compositionality is Not the Problem', *Logic and Logical Philosophy*, 6: 7–33.
- ____ (2001), 'Formal Features of Compositionality', *Journal of Logic, Language and Information*, 10: 7–28.
- Kaplan, D. (1989), 'Demonstratives', in J. Almog, J. Perry, H. Wettstein (eds.), *Themes from Kaplan*. Oxford: Oxford University Press.
- King, J., and J. Stanley (2005), Semantics, Pragmatics, and the Role of Semantic Content, in Z. Szabó (ed.), *Semantics vs. Pragmatics*. Oxford: Oxford University Press.
- Lepore, E. (2003), 'An Abuse of Context in Semantics: The Case of Incomplete Definite Descriptions', in A. Bezuidenhout and M. Reimer (eds.), *Descriptions and Beyond: An Interdisciplinary Collection of Essays on Definite and Indefinite Descriptions and Other Related Phenomena*. Oxford: Oxford University Press.
- Montague, R. (1970a), 'Universal Grammar', *Theoria*, 36: 373–98.

³³ We are grateful to comments from two anonymous referees. We have also benefited from discussions of the themes of the paper with François Recanati and Mike Harnish. Work on this paper has been supported by a grant from the Swedish Research Council (for Pagin) and the Canadian NSERC (for Pelletier).

- (1970*b*), 'English as a Formal Language', in B. Visentini (ed.), *Linguaggi nella Società e nella Tecnica*, pp. 189–224. Milan: Edizioni di Comunità.
- (1973), 'The Proper Treatment of Quantification in Ordinary English', in J. Hintikka, J. Moravcsik, and P. Suppes (eds.), *Approaches to Natural Language*, pp. 221–42. Dordrecht: Reidel.
- Nunberg, G. (1979), 'The Non-Uniqueness of Semantic Solutions: Polysemy', *Linguistics and Philosophy*, 3: 143–84.
- Pagin, P. (2003), 'Communication and Strong Compositionality', *Journal of Philosophical Logic*, 32: 287–322.
- (2005*a*), 'Compositionality and Context', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*, pp. 303–48. Oxford: Oxford University Press.
- (2005*b*), 'Semantic and Pragmatic Composition', paper presented at the semantics–pragmatics workshop in Context 05.
- Partee, B. H. (1989), 'Binding Implicit Variables in Quantified Contexts', in C. Wiltshire, B. Music, and R. Graczyk (eds.), *Papers from CLS 25*, pp. 342–56. Chicago: Chicago Linguistics Society. Repr. in Partee 2004: 259–81. Page references to the reprint.
- (2004), *Compositionality in Formal Semantics: Selected Papers by Barbara H. Partee*. Oxford: Blackwell.
- Pelletier, F. J. (1994), 'The Principle of Semantic Compositionality', *Topoi*, 13: 11–24, expanded version reprinted in S. Davis and B. Gillon (eds.) (2004), *Semantics: A Reader*, pp. 133–57. New York: Oxford University Press.
- Perry, J. (1986), 'Thought without Representation', *Proceedings of the Aristotelian Society*, 60: 263–83.
- Recanati, F. (2004), *Literal Meaning*. Cambridge: Cambridge University Press.
- (2005), 'Literalism and Contextualism: Some Varieties', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*, pp. 171–96. Oxford: Oxford University Press.
- Reimer, M. (2002), 'Do Adjectives Conform to Compositionality?', *Nous*, 16: 183–98.
- Richard, M. (1981), 'Temporalism and Eternalism', *Philosophical Studies*, 39: 1–13.
- Russell, B. (1956), 'The Philosophy of Logical Atomism', in R. Marsh (ed.), *Bertrand Russell: Logic and Knowledge, Essays 1901–1950*, pp. 175–281. London: Unwin Hyman; 1st publ. 1918.
- Schlick, M. (1985), *General Theory of Knowledge*. Peru, Ill.: Open Court, 1st publ. 1918.
- Searle, J. (1978), 'Literal Meaning', *Erkenntnis*, 13: 207–24.
- (1980), 'The Background of Meaning', in J. Searle, F. Kiefer, and M. Bierwisch (eds.), *Speech Act Theory and Pragmatics*, pp. 221–32. Dordrecht: Reidel.
- (1983), *Intentionality*. Cambridge: Cambridge University Press.
- Sperber, D., and D. Wilson (1992), *Relevance: Communication and Cognition*, 2nd edn. Oxford: Blackwell.
- Stanley, J. (2002), 'Making it Articulated', *Mind and Language*, 17: 149–68.
- and Z. Szabó (2000), 'On Quantifier Domain Restriction', *Mind and Language*, 15: 219–61.
- Travis, C. (1985), 'On What is Strictly Speaking True', *Canadian Journal of Philosophy*, 15: 187–229.
- (1989), *The Uses of Sense*. Oxford: Oxford University Press.
- (1996), 'Meaning's Role in Truth', *Mind*, 100: 451–66.

- Westerståhl, D. (1998), 'On Mathematical Proofs of the Vacuity of Compositionality', *Linguistics and Philosophy*, 21: 635–43.
- (2004), 'On the Compositional Extension Problem', *Journal of Philosophical Logic*, 35: 549–82.
- Wittgenstein, L. (1981), *Tractatus Logico-Philosophicus*. London: Routledge & Kegan Paul; 1st publ. 1921.

EBSCOhost®