UNIVERSITAT DE BARCELONA

An approach to occasion-sensitivity

Claudia Picazo Jaque

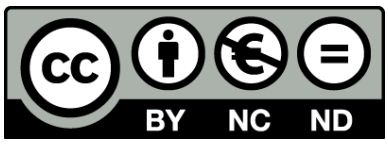

An approach to occasion-sensitivity

by Claudia Picazo Jaque

Ph.D. Dissertation
Written under the supervision of Manuel Pérez Otero

Programme: Cognitive Science and Language
Faculty of Philosophy
University of Barcelona

Abstract

In this dissertation, I explore and defend the hypothesis that language is occasion-sensitive. I understand occasion-sensitivity as a combination of two claims. The first claim has it that, for a large class of non-indexical sentences, linguistic meaning underdetermines truth-conditions. According to the second, sentence-tokens only determine a partial function from states of affairs to truth-values. Travis cases provide evidence for both claims. After discussing and dismissing minimalist and indexicalist explanations of Travis cases, I outline a situationalist approach in which the truth-value of a sentence-token partly depends on the constraints imposed by the activity that the token concerns. Finally, I discuss the compatibility of occasion-sensitivity with systematic theories of truth-conditions.

Acknowledgments

First and foremost, I would like to thank my supervisor, Manuel Pérez Otero. I have greatly benefited from his comments and criticisms on earlier drafts of this dissertation, as well as from his advice and support during these years. I am also indebted to François Recanati, my supervisor at the Jean Nicod Institute, and Patrick Greenough, my supervisor at the University of St Andrews, for very helpful discussions and constructive feedback. My gratitude goes also to all of the people with whom I have had the opportunity to discuss the topics of this dissertation. Many thanks to Delia Belleri, Laura Delgado, Inés Crespo, Manuel García-Carpintero, Johan Gebo, Andreas Heise, Max Kölbel, Palle Leth, Nasim Mahoozi, Marina Ortega, Giulia Pravato, David Rey, Charles Travis, Ignacio Vicario and Agustín Vicente, and to those I might have forgotten. Working at the LOGOS Research Group has been a very rewarding experience. I would also like to thank all the participants in the LOGOS GRG and the reading groups I have had the opportunity to attend. Special thanks to my friends and workmates at the office 4023 Carlos Jaén and Pilar Terrés for so many great times after–and also during–work; and to Jorge, for his support throughout the process.

This work has been funded by the Spanish Ministry of Economy and Competitiveness and the European Union, projects FFI2011-29560-C02-01 and FFI2015-63892-P (MINECO, AEI/FEDER, EU).

Contents

Introduction

It is common to conceive the content of a representation in terms of truth-conditions. However, this conception raises an interesting question: Should we think of representation types as the bearers of truth-conditions or is it rather representations in use that impose a condition on the world? This question has become specially pressing when it comes to linguistic representations (natural language sentences).

The notion of truth (and related notions such as reference or extension) has played a crucial role in the study of natural language. For many philosophers, including notably early Wittgenstein and Davidson, understanding a sentence S consists in grasping the conditions under which S would be true. This claim has motivated the identification of the meaning of S with its truth-conditions, or with something that determines truth-conditions. Following this idea, the meaning of the subsentential expressions can be seen as a contribution to the truth-conditions of the sentence–proper names contribute a referent, predicates contribute an extension or an intension, and so on. Let us call this approach the traditional view.

Let us suppose that we have a syntactic theory for a natural language or a fragment of it. We have a list of expressions, classified into categories or types (proper names, predicates, quantifiers, definite descriptions, and so on) and syntactic rules providing the means to form grammatically correct sentences out of them. The task of semantics is to assign meanings to those well-formed sentences. This is standardly achieved by assigning meanings (semantic values) to the simple expressions and compositional rules determining the meanings (semantic values) of complex expressions. Since Frege's writings, truth has played a central role in the task of assigning semantic values to sentences. Frege famously divided meaning into sense and reference (Frege, 2010). To different types of expressions correspond different types of referents. At the most basic level, the referent of a proper name is an object and that of a sentence is a truth-value (the True or the False). Other referents can be defined by using these two. For instance, the referent of a predicate is a function from objects to truth-values. The referent of complex expres-

sions, including sentences, can be obtained by compositional rules. These elements can be used to provide a semantic theory for natural language. A very simple example:

> (1) The referent of the predicate 'is a philosopher' is the function *f* from objects to truth-values such that the value of *f(x)* is the True if and only if x is a philosopher.
>
> (2) The referent of the proper name 'Frege' is Frege.
>
> Since Frege is a philosopher, the value of *f(Frege)* is the True and, consequently, the referent of the sentence 'Frege is a philosopher' is the True.

The traditional view comes with a certain answer to the question about the bearers of truth-conditions. According to it, non-indexical sentences express truth-conditional content outside a context of use. Having such-and-such truth-conditions, or expressing such-and-such truth-evaluable content, is a property of the sentence-type. Indexical sentences express different truth-conditions at different contexts. However, this content can be conceived as a structured proposition, and structured propositions are usually thought of as having truth-conditions independently of the context at which they are tokened.

There are reasons for calling into question the traditional view. Some philosophers have noted that, when it comes to actual communication, truth-conditional content is often the result of adjusting meaning to the context of use. Here is an oft-discussed example. The sentence 'I've had breakfast' only establishes that the speaker has had breakfast at a time t such that t is prior to the time of utterance. However, utterances of this sentence are typically interpreted as saying that the speaker has had breakfast the day of utterance–not a week before, etc. Similarly, an utterance of 'He eats rabbit' will very likely be interpreted as saying that the salient man eats rabbit meat, and not rabbit fur[1]. The relevant truth-conditions of an utterance seem to go beyond the properties of the sentence-type, even in absence of classical indexicals–they are the output of pragmatic interpretation.

In a more radical vein, some authors have called into question the thesis that meaning can be identified with reference or truth-conditions. Chomsky[2] seems to be sceptical that common-sense concepts, such as *book* or *house*, or proper names as 'London', are apt for figuring in a theory of meaning that specifies truth-conditions. His worry is that whether something can be rightly described as a 'book' or a 'house' depends on a number of things, including how we use the object. To what object the expressions 'the

[1] See Sperber and Wilson (1995, pp. 189-190) and Recanati (2004, p. 24).

[2] See specially his (2000).

book' or 'the house' refer depends on human interests. For example, 'book' can be used to talk about the physical object or the story, or both (as in 'The book he wrote weighs two pounds'). 'London' can also be used to refer to a number of things. Words seem to relate to objects in intricate ways.

One of the strongest criticisms against the traditional view has been that of Travis' on the basis that extension is sensitive to the occasion of use. Travis targets the thesis that the meaning of a declarative sentence can be identify with truth-conditions or with something that determines truth-conditions. He has identified a certain phenomenon of truth-value shifts across occasions and used it to argue that expressing a truth-evaluable content is a property that sentences have in use. What a sentence says, Travis holds, depends on the character of the occasion on which it is used.

The aim of this dissertation is to explore the hypothesis that language is occasion-sensitive. The first task is to examine what is meant by 'occasion-sensitivity' and the arguments in support of the claim that natural language is occasion-sensitive. This task is undertaken in chapter 1, an introduction to Travis' view. To hold that language is occasion-sensitive is to hold that the truth-conditions (or satisfaction-conditions) of most of our utterances depend on the occasion of use in a way that is not determined by meaning–or, in other words, that the semantics of most sentences underdetermine the truth-conditions of its tokens. In this chapter, I reconstruct what I take to be Travis' main arguments: Travis cases and the proliferation argument. Travis cases depict a sentence S and two different occasions of use. Whereas in the first occasion S seems to be true, in the second it seems to be false (or vice versa), despite the state of the world and the referents of indexicals and descriptions being the same. Travis concludes that the semantics of S (the properties of the type) do not determine the truth-conditions of its tokens. The proliferation argument is intended to cast doubt on the existence of representations that are immune to Travis cases on the basis that these can be created for more precise versions of S. Because of his endorsement of this argument, Travis holds that there are no truth-bearers outside a context of use. One of the aims of this chapter is to argue that a certain development of this claim–namely, the view that whether an utterance is true or false depends on what we judge reasonable–runs the risks of precluding the objectivity of content. This imposes some requirements on how to think of truth-conditional content on the assumption that language is occasion-sensitive.

In chapter 2, I set out to defend the claim that language is occasion-sensitive against minimalist and indexicalists replies to Travis' arguments. In order to do so, I distinguish the Principle of Compositionality from what I call Semantic Propositionalism. Ad-

vocates of occasion-sensitivity reject only the latter. Semantic Propositionalism states that the semantics of a well-formed sentence S determines a truth-evaluable content (at a context). I argue that neither minimalist nor indexicalist accounts succeed in their defence of Semantic Propositionalism vis-à-vis Travis cases. Minimalists find themselves in an unstable position. In order to secure minimal propositions, they need to dismiss common reactions to Travis cases. But this casts doubt on the possibility of finding out what the literal satisfaction conditions of the expressions used in Travis cases are. Indexicalism tries to secure Semantic Propositionalism by claiming that some predicates are context-dependent, usually because their syntactic form contains some context-sensitive variables. These proposals have a different problem. If they are to be a defence of Semantic Propositionalism, indexicalist theories must fulfil two conditions. First, they must provide a set of necessary and sufficient variables. Second, it must be the case that semantics determines, and not merely constrains, the extension of the predicate. I argue that current proposals do not succeed in fulfilling both conditions.

As a consequence of this discussion, I take it that there are good reasons to think that natural language is occasion-sensitive–namely, Travis cases together with the absence of good arguments to the contrary. At this point, two questions arise. First, one might wonder whether occasion-sensitivity applies to non-linguistic representations, such as mental representations. Second, it is interesting to assess to what extent formal frameworks that make use of double-indexing are compatible with occasion-sensitivity and whether these frameworks can be used to account for our capacity to grasp truth-conditional content.

Chapter 3 addresses the first question. Some pragmatist approaches hold that, in the case of natural language, semantics does not determine extension, but take it that the same is not true of mental representations. In this chapter, I focus on Fodor's arguments and on Carston's theory of *ad hoc* concepts. Fodor has provided two arguments to the effect that, in contrast with natural language, mental representations are not underdetermined. The productivity argument has it that the best explanation to the productivity of thought is compositionality, and mental representations being compositional prevents them from being underdetermined. The argument from equivocation is based on the idea that only a non-equivocal mental representation can resolve a linguistic equivocation. In order to properly assess the arguments, I distinguish Truth-conditional Compositionality from Meaning Compositionality. Following (Clapp, 2012b) and (Recanati, 2007), I argue that neither argument work. The productivity argument only establishes Meaning Compositionality, something compatible with (truth-conditional)

Underdeterminacy. As to equivocation, the context of use can solve the equivocation in absence of a non-equivocal mental representation. I also argue that if Carston's *ad hoc* concepts are created on line, then she cannot avail herself of the productivity argument. On the other hand, if *ad hoc* concepts are created out of pre-existing concepts that are similar to the ones encoded in natural language, as Carston's explanation seems to suggest, then there is no reason to think that mental representation are immune to Travis cases. An additional aim of chapter 3 is to distinguish Type-Underdeterminacy from Token-Underdeterminacy. I argue that there are reasons to think that even tokens suffer from some Underdeterminacy in the sense that they only determine a partial function from states of affairs to truth-values.

In chapter 3 I conclude that we are left with no reasons for positing occasion-insensitive representations. Instead of relying on occasion-insensitive mental representation or having recourse to occasion-insensitive structured propositions, occasion-sensitivity calls for a non-standard notion of utterance content. The aim of chapter 4 is to provide such a notion. I hold that Austinian propositions, conceived as including a lekton and an activity, can do the work. Travis cases suggest that the feature of the context that affects the extension of the predicate is the activity in place. To different activities correspond different criteria of applicability for words. Thus, adopting a situationalist framework[3], we can think of the truth-conditional content of an utterance as including not only the conventional meaning of the sentence uttered but also the activity against which it is evaluated. This notion of content is compatible with Token-Underdeterminacy. After having put forward this notion of utterance content, I discuss a potential problem for the approach. If activities are very finely individuated, as the possibility of creating complex Travis cases recommends, then sharing content across contexts will be difficult to achieve. However, sharing content, for example by reporting what someone has said, is usually easy. I argue that this problem can be solved by having Austinian propositions with different granularities, thus adopting a form of multi-propositionalism. I further argue that this approach to occasion-sensitivity escapes the potential loss of objectivity introduced in chapter 1.

In chapter 5 I address the question whether phenomenon that Travis has detected is compatible with standard semantic[4] theories. Semantic theories have been seen as

[3]Situationalism can be understood as a form of relativism or non-indexical contextualism in which the truth-value of the proposition expressed by an utterance is evaluated with respect to a situation (and not only with respect to more standard parameters as possible world, time and place).

[4]The term 'semantics' can be understood in two ways. First, it can be understood as the study of meaning or of the properties of expression-types. Second, semantics can be understood as the study of

an explanation of our ability to interpret speech. Advocates of occasion-sensitivity and similar pragmatic views are under pressure to show that their rejection of certain theories is compatible with a plausible account of our ability to grasp truth-conditions. I argue that occasion-sensitivity, and in particular the notion of truth-conditional content I introduce in chapter 4, is compatible with there being systematic connections between activities and truth-conditions, which can be used to account for our ability to interpret speech. Moreover, following (Predelli, 2004, 2005b,a) I argue that occasion-sensitivity is compatible with double-indexed theories. Double-indexed frameworks are equipped to deal with different forms of context-sensitivity (broadly understood). For example, a semantic theory in which indices include standards of precision can deal with Austin's well-known example that whether an utterance of the sentence 'France is hexagonal' is true or false depends, roughly, on what is at stake. Similarly, indices can include activities. I further argue that occasion-sensitivity is compatible with semantic theories that specify truth-conditions on the assumption that these specifications bear default understandings.

truth-conditions. Here I am using the term in this second sense.

Chapter 1

Occasion-sensitivity

Charles Travis has provided a number of examples aimed at showing that semantics[1] is compatible with variation in truth-conditions (keeping the value of indexicals fixed). According to Travis, this shows that language is occasion-sensitive. The aim of this chapter is to introduce occasion-sensitivity and Travis' arguments in support of his pragmatist position. In order to do so, I reconstruct what I take to be Travis' main arguments (Travis cases, the proliferation argument). After this, I argue that a certain reconstruction of some of Travis' positive remarks about truth-conditions runs the risk of precluding objectivity and that therefore occasion-sensitivity calls for a non-standard notion of content that escapes the threat. In the last section I briefly present the current positions in the contextualist debate and, follwoing Borg, argue that Travis' approach is different from other pragmatists approaches usually labelled as 'contextualism'.

1.1 A phenomenon in need of explanation

Here is a phenomenon exhibited by natural languages. Take a declarative sentence S. Make sure that S contains no indexicals or, if it does, make sure that their values are kept fixed throughout the process. Imagine two different scenarios in which S is used. Keep

[1] In what follows I mean by 'semantics' the properties of the expression type or the properties determined by the expression type at a context of use. In this sense, it is a substantial question whether semantics must deal with truth-conditions

the state of the world fixed in the relevant respects across these scenarios. For a large class of sentences, it can be the case that the sentences is true in one of those scenarios and false in the other. This phenomenon of variation is the core of occasion-sensitivity.

Charles Travis, the philosopher who coined the expression 'occasion-sensitivity' and the main advocate of the claim that natural language is occasion-sensitive[2], has come up with many different examples in which the phenomenon is deployed. Let me quote at length some of them[3].

> ['The kettle is black']
> Consider the following two contrasting cases:
>
> A. Max fills his shiny new aluminum kettle with the makings of a stew, and sets it over the campfire. An hour later, he informs Sam that he has done this. 'That was pretty stupid', Sam replies, and rushes out to the fire. He returns holding a soot-blackened pot and says (speaking truly), 'Look. The kettle is black.'
>
> B. Everard and Clothilde are acquiring their first common batterie de cuisine. For many reasons, including tradition and presumed heat-retaining properties, they want only black pots. (Though what sort of black pot happens not to matter much.) Coincidentally, Max's soot-blackened pot has come to rest precisely in the shop window into which they are now staring. Everard says, 'Look. There's a nice black pot.' But Clothilde is more observant. 'No it isn't black', she replies, 'it's only covered with soot. How careless of them to let that get in their window.' And off they go elsewhere, with, to all appearances, Clothilde having spoken the truth. (Travis, 2008b, p. 26)

> ['The ball is round']
> Consider the sentence 'The ball is round', and two cases of its use. Case A: What shape do squash balls assume on rebound? Pia hits a decent stroke; Jones watches. 'The ball is round,' she says at the crucial moment. Wrong. It has deformed into an ovoid. Jones did not say the ball to be as it was, so spoke falsely. Case B: Fiona has never seen squash played. From her present vantage point the ball seems a constant blur. 'What shape is that ball?,' she asks. 'The ball is round,' Alf replies; truly, since it is the sort of ball a squash ball (and this one) is. It is not, e.g., like a

[2] Despite the originality of his examples, Travis is not the first philosopher to have detected the phenomenon of occasion-sensitivity. As he often acknowledges, Austin (1975) had already expressed similar worries about the relation between sentences and truth-values and put forward a famous Travis case involving the sentence 'France is hexagonal'. Searle (1978; 1980) also presents similar examples and argues that satisfaction conditions are relative to a background of implicit assumptions.

[3] I will call this kind of example a Travis case.

very small rugby ball. (Travis, 1996, p. 97)

['The leaves are green']
A story. Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, 'That's better. The leaves are green now.' She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those.' But now Pia speaks falsehood.

If the story is right, then there are two distinguishable things to be said in speaking (1) with the stipulated semantics. One is true; one false; so each would be true under different conditions. That semantics is, then, compatible with semantic variety, and with variety in truth-involving properties. So what the words of (1) mean is compatible with various distinct conditions for its truth. (Travis, 1997, p. 111-112)

['There's milk in the refrigerator']
Suppose that the refrigerator is devoid of milk except for a puddle of milk at the bottom of it. Now consider two possible speakings, by Odile, of the words, 'There's milk in the refrigerator.' For the first, Hugo is seated at the breakfast table, reading the paper, and from time to time looking dejectedly (butmeaningfully) at his cup of black coffee, which he is idly stirring with a spoon. Odile volunteers, 'There's milk in the refrigerator.' For the second, Hugo has been given the task of cleaning the refrigerator. He has just changed out of his house-cleaning garb, and is settling with satisfaction into his armchair, book and beverage in hand. Odile opens the refrigerator, looks in, closes it and sternly utters the above words. I claim that the example bears at least the following description: though there is no ambiguity in the English words 'There is milk in the refrigerator', or none relevant to the differences between the two speakings, Odile's words in the first case said what was false, while in the second case they said what was true. (Travis, 1989, pp. 18-19)

['Hugo weighs 79 kilos']
Consider the words, 'Hugo weighs 79 kilos' and the following situation: when Hugo steps on the scale in the morning, it reads 79 kilos, and that is a stable result. However, it is now after lunch; fully dressed (in winter clothing), Hugo would register 81 kilos on the scales. Now consider two speakings of the words. For the first, Hugo must weigh 79 kilos, and no more, to qualify for some sporting event. There is a discussion as to whether he does qualify. Odile, who has seen him step

> on the scale, tries to settle the matter by revealing her information on the subject. She says, 'Hugo weighs 79 kilos.' For the second, Hugo is about to step on to a very delicate trestle bridge across a ravine which can take a maximum of 80 kilos without snapping. Or Hugo is placed on a balance scale to weigh out 79 kilos of gold (for some weighty purpose). The question is whether Hugo ought to step on to the bridge, or whether that really is 79 kilos of gold, and not more. Odile volunteers, 'Hugo weighs 79 kilos.' Again, the claim is that Odile spoke truth in the first case and falsehood in the second. (Travis, 1989, pp. 19-20)

> ['The student has a desk']
> Pia, concerned for an impoverished student, asks Max whether the student has a desk. Max replies, 'It depends on what you mean by a desk. If you count a door over two stacks of milk crates as a desk, then yes. If not, then no'. Max has given a clear enough reply, making a natural and familiar enough use of English. (Travis, 2000, p. 3)[4].

These examples reveal a phenomenon of truth-value changes across uses, with semantics and the relevant state of affairs being fixed. They motivate the claim that semantics, understood as the properties of expression-types, is compatible with variation in extension across uses (keeping the state of the world fixed). Throughout this work, I will focus on predicates. Concerning predicates, the examples motivate the claim that, for a large class of predicates, the meaning of a predicate F, or its semantics, is compatible with different uses of F determining different extensions–not because the state of the world changes from one use to the other, but because different uses of F impose different criteria on what counts as F. In this sense, these examples are supposed to reveal something about the relation between semantics, truth and use. In particular, the lesson is that the occasion in which a well-formed declarative sentence is used matters

[4] As an example of an actual Travis case, see the following quote from an interview to Varoufakis:
'HL: You must have been thinking about a Grexit from day one...
YV: Yes, absolutely.
HL: ...have preparations been made?
YV: The answer is yes and no. We had a small group, a 'war cabinet' within the ministry, of about five people that were doing this: so we worked out in theory, on paper, everything that had to be done [to prepare for/in the event of a Grexit]. But it's one thing to do that at the level of 4-5 people, it's quite another to prepare the country for it. To prepare the country an executive decision had to be taken, and that decision was never taken.' (Newstatesman. 2015. Yanis Varoufakis full transcript: our battle to save Greece. [ONLINE] Available at: http://www.newstatesman.com/world-affairs/2015/07/yanis-varoufakis-full-transcript-our-battle-save-greece. [Accessed 1 June 2017].)

to the truth-conditions of the utterance. As Travis puts it:

> This is pretty much the core idea of occasion-sensitivity. One specifies some things to speak of–being blue, say, the sky–such that in speaking of them in a certain structured way–saying the sky to be blue–one might, if things go well, say something to be so. Then in speaking of those things, in that way, and saying something to be so, one might say any of many distinguishable things. I mean the statement to be as general as just stated. Or rather, I mean it to hold for any sublunary things to speak of. For the nonce, I leave mathematics to one side. The core idea, naturally enough, has a corollary for semantics–by which I mean here a theory of what expressions of a language mean. Expressions of a language identify things to talk about–as 'being blue' identifies being blue, 'The North Sea' identifies the North Sea, and 'The North Sea is blue' identifies the North Sea's being blue. By the thesis, in talking about those things (in a given structured way) one might say any of many things. So what the expressions mean cannot fix some one of these as that which is thus said. So what they mean cannot fix any one condition as the condition for 'their' truth. Meaning cannot connect to truth like that. (Travis, 2008a, p. 4)

The variation in truth-conditions deployed in Travis cases is supposed to be a pragmatic matter. To hold that language is occasion-sensitive is to hold that the variation in truth-conditions is a pragmatic matter–rather than being driven by linguistic meaning[5] . In this sense, occasion-sensitivity is equivalent to what other philosophers call 'semantic underdeterminacy'. I think that the variety of examples presented and the intuitions about truth-values they trigger make the phenomenon worthy of interest. It is important to ask what is going on in the examples and how, in view of the variation they involve, we should think about language.

This last point becomes salient when one considers the consequences of the claim that language is occasion-sensitive. It has been common to identify meaning either with truth-conditions or with something that determines truth-conditions. Here is a simplified version of a traditional approach to language. First, the meaning of non-indexical declarative sentences is thought of as being both truth-conditional and systematic, that is, it can be generated by a theory consisting in a set of axioms (the meanings

[5] I will use the term 'occasion-sensitivity' not as a synonym for 'variation in truth-conditions', but to refer to a pragmatist view about the phenomenon. To hold that language is occasion-sensitive is to hold that truth-conditions depend on the occasion of use in a way that is not determined by semantics. See (Davies, 2011) for a similar use. In chapter 3 I will introduce Token-Underdeterminacy. I think Token-Underdeterminacy is also part of the notion of occasion-sensitivity as Travis presents it.

of simple expressions) and rules of composition. Second, it is sometimes assumed that most context-sensitivity is linguistically driven. Even though features about the uses of context-sensitive sentences help determine their truth-conditions (at a context of use), it is still the meaning of the sentence that determines that this is so. Third, there is a division of labour between semantics and pragmatics. According to the traditional divide, semantics and pragmatics come in at different stages. First, semantics determines the truth-conditions of the sentence in question, thus determining what is literally said by the sentence. Second, on an occasion of use, a sentence can be used to convey something different from what it, strictly speaking, means (says). This is where pragmatics enters the picture.

Travis cases cast doubts on the three claims. Against the first, Travis cases provide reasons to think that meaning does not connect to truth outside a context of use, that is, that truth-conditions depend on uses of sentences, and not merely on what the sentences mean. Against the second, it seems that a great deal of context-sensitivity (broadly understood) is not linguistically driven. Finally, against the third, pragmatics seems to enter the picture at the first step. Since truth-conditions are affected by features of the context of use, pragmatics cannot operate at a second, post truth-conditional level.

Travis is not the only philosopher to have called into question this traditional approach. During the last decades, several philosophers have argued for a general phenomenon of underdeterminacy in natural language on the basis of the existence of non-indexical, well-formed declarative sentences that seem unfit to be evaluated as true or false or, otherwise, that seem to express different truth-conditions at different contexts. I will only mention a couple of examples. Bach (1994) argues that some sentences express incomplete propositions. For example, he argues that the sentence 'Tipper is ready' does not express a complete proposition, but only a propositional radical. The proposition, of itself, does not specify what Tipper is ready for. However, Bach claims, this piece of information is needed in order to get a truth-value. As a result, the linguistic meaning of 'Tipper is ready' underdetermines the truth-conditions of an utterance of this sentence. Similarly, Recanati (2004) argues for the existence of strong pragmatics effects. According to Recanati, the conventional meaning of an expression can be pragmatically modified even when the syntactic or logical form of the expression does not mandate it. 'Modulation' is the cover term for these strong pragmatic effects. Words without context-sensitive variables can be modulated in certain contexts, i. e., their meaning is adjusted to the context of use. This adjusted meaning should be re-

garded, according to Recanati, as what is said by the utterance (as opposed to what is implicated).

Travis' position is more radical. First, he generalizes occasion-sensitivity to any sentence (excepts mathematics). Second, as I will show later, he takes it that what goes for sentences goes for any other representation. Because of this, he calls into question the notion of proposition as a structured item with intrinsic truth-conditions. This view can be considered similar to Searle's ((1978) and (1980)), who claims that truth-conditions, in general, are relative to a background of non-explicit assumptions, some of which might not even be representational.

In what follows I examine Travis' view on occasion-sensitivity. In sections 2 and 3, I present what I take to be Travis' main targets and arguments. I distinguish two arguments about the relation between semantics and truth. First, we have the well-known Travis cases. It is important to note that Travis cases can be iterated. This motivates a second argument, one that casts doubts on the possibility of domesticating first-order Travis cases. Section 4 briefly presents the Wittgensteinian roots of Travis' philosophy and sketches his positive proposal. The next section examines whether Travis' positive remarks are tenable. I finish by placing occasion-sensitivity within the contextualist debate.

1.2 Meaning and truth (I): Travis cases

Travis cases undermine the thesis that semantics is in the business of delivering truth-evaluable content. In particular, Travis seems to target the following principle[6] :

> *Semantic Propositionalism (Travis):* The meaning of a well-formed declarative sentence S is identical to or determines a truth-evaluable content (a proposition, a truth-condition).

This principle seems to underlie Grice's philosophy. Grice writes that 'In the sense in which I am using the word *say*, I intend what someone has said to be closely related

[6] Bach (2006) calls this 'Propositionalism'. In chap. 2 I will use slightly different version of this principle.

to the conventional meaning of the words (the sentence) he has uttered' (Grice, 1975, p. 44). If what is said is also assumed to be truth-conditional, the resulting picture will be an implementation of Semantic Propositionalism. This principle is also on the basis of Davidson's thesis that one can specify the meaning of a sentence by specifying the conditions under which it is true (Davidson, 1967).

I will start with Travis' argument against the identification of meaning with truth-conditions. Travis' argument against this principle is usually presented as concerning meaning, although he sometimes refers to semantic properties. It goes as follows. Whatever the meaning of the sentences involved in Travis cases is, it is something that plausibly remains constant across uses[7]. However, the truth-evaluable content that S expresses vary across contexts. So the meaning of S is compatible with S's tokens expressing different truth-evaluable contents. So meaning cannot be identified with any of these truth-evaluable contents. Meaning is something remains fixed across uses, whereas truth-evaluable content is something that can vary[8] .

What does 'meaning' means here? Travis is not putting forward a theory of meaning. For example, in his (1981), he simply takes the meaning of an expression to be whatever we describe when we explain the meaning of a word—whatever the m is in 'E means m'. Now, whatever that might be, it is not to be identified with the extension or the satisfaction conditions of the expression, since that m will be something that remains constant across uses, whereas truth-conditions can shift.

In other texts, Travis talks about the semantics of a sentence instead of its meaning. What about semantics? 'Semantics is concerned with properties that expressions of a language such as English have. It is not per se concerned with properties that only some occurrences of them—if any—have.' (Travis, 2006a, p. 151). Semantics is concerned with properties of expressions types. In the previous examples, the same sentence is used twice. Since it is the same expression type, its semantics must remain constant. Still, its truth-conditions vary.

If the argument is sound, what a sentence means cannot be identified with when it would be true. What a sentence means is compatible with expressing different truth-conditions on different occasions. It seems that truth-value partly depends on the specifics of the occasion. However, it could still be the case that what a sentence means determines what a use of it would say at a context of use. Travis reacts against this second

[7] I am not yet considering the possibility that the predicates involved are indexical. If so, their character would remain constant, but the content expressed would shift across contexts. The referents of other expressions are kept fixed in the examples.

[8] See the previous quote for a statement of the argument.

possibility:

> A sentence (nearly any) may, on one speaking or another, say any of indefinitely many distinct things, each true under different conditions. Nearly any part–a simple predicate like 'is red', say–may make any of many contributions to what it thus says, specifically to the conditions for its truth. All these contributions are ones those words would sometimes make given what they mean; all are compatible with their meaning that. Nor does their meaning provide the means for deriving when they would make which contribution. Rather, seeing what words did, or would, say on a given occasion is a matter of properly appreciating the circumstances of that speaking, and correctly perceiving which of their many possible contributions they are most reasonably taken to have made in those circumstances. (Travis, 1991, p. 68)

The quote is about what a sentence on an occasion says. Travis sometimes goes from truth-conditions to what is said and vice-versa. There is an intuitive sense of 'saying' in which what a sentence says in one of the scenarios depicted in a Travis case is not equivalent to what the same sentence says in the other scenario. If we were to paraphrase them, we would be using different words. Moreover, this point is reinforce by the usual idea that content can be conceived as truth conditions. Travis is disputing that meaning, or semantics, determine truth-conditions, not that grasping what someone says has to do when knowing in what conditions his utterance would be true.

Travis considers several ways in which his defence of the need of a pragmatic approach could be rejected (what he sometimes calls 'domestications'). First, one can look for an ambiguity in the relevant predicates, so that this is not a case of one sentence having different truth-values, but of two different sentences having different truth-values–something trivial (Travis, 1981, 1997). Second, one can claim that the sentences for which Travis cases arise are elliptical (Travis, 1981, 1997). Third, the meaning of a sentence S could be said to determine (instead of being identical to) the truth-conditions of utterances of S. One can look for an indexical or somehow apply the model provided by indexicals or other forms of context-sensitivity to these cases (Travis, 1997).

Moreover, one can deny that there is a change in truth-value in the examples proposed and account for the intuitions of differences in what is said by appealing to implicatures or some other kind of communicated, not strictly speaking said, content (Travis, 2008b, 1997). Following Grice[9] , one could argue that there is here some implicature-like

[9]See his (1975) and also the Prolegomena to (1991).

content that is responsible for the truth-value shift intuition. The problem with this option is that, according to Travis, this solution only works assuming that the meaning of the words 'is black' decides which of the utterances (if any) is literally true. However, it is not at all clear that meaning is able to do this, for it seems perfectly compatible with what 'is black' means that it can be sometimes used to describe the surface of an object at a determinate moment, whereas other times it is used to describe the object's original colour (Travis, 2008b, pp. 28-30). The word 'black' simply refers to a colour[10] .

What about ambiguity, ellipsis and indexicality? An interesting feature of Travis cases make these options less appealing than it might firstly appear: Travis cases can be iterated, i.e., we can created higher-order Travis cases. Take a sentence for which a Travis case has been proposed. Make explicit the aspect that shifts from one scenario to the other or keep the dimension that has shifted fixed (select one of the readings of the sentence). Given ingenuity enough, it is very likely that someone will come up with a Travis case for this new sentence or the reading of the old sentence that was selected. This is the core of what I will call *the proliferation argument*.

1.3 Meaning and truth (II): The proliferation argument

It could be argued that the variation exhibited in Travis cases does not motivate a pragmatic view. In this sense, it could be 'domesticated'. In this section, I reconstruct what I take to be Travis' argument against this option. Before that, let me note that Searle has a very similar argument. He claims that:

> [F]or a large number of cases the notion of the literal meaning of a sentence only has application relative to a set of background assumptions, and furthermore these background assumptions are not all and could not all be realized in the semantic structure of the sentence in the way that presuppositions and indexically dependent elements of the sentence's truth conditions are realized in the semantic structure of the sentence. (Searle, 1978, p. 210)

His method of argumentation consists in showing how simple declarative sentences, such as 'The cat is on the mat', have truth-conditions only relative to a number of as-

[10] In chap. 2 I examine current applications of the Gricean strategy.

sumptions that often go unnoticed–what he calls 'the background'. In this case, for example, the sentence only has application on the assumption that there is gravitational force[11]. Could we make gravitational force somehow part of the semantics of the sentence, so as to get rid of background-dependence? Searle's negative answer is supported by three considerations ((1978) and (1980)). First, if we started adding assumption, we would never know where to stop. Second, in spelling out the assumptions, we need to use sentences, and so we would be bringing in further assumptions. And third, it is not clear that the background is representational, that is, that all the assumptions are representational in content. The reason why we understand certain sentences as we do has to do with our implicit knowledge of the things involved and the practices in which we use them.

Travis' occasion-sensitivity is relevantly similar to Searle's background dependence. Travis cases can be seen as scenarios where the assumptions vary. For the sake of clarity, we can establish a distinction between cases in which the application conditions of a sentence are relative to an assumption that remains constant across all our actual uses and cases in which the relevant assumption is only present in some ordinary uses of the sentence. For example, the assumption that an object's colour is its original colour is in place only on some occasions. The assumption that there is gravitational force was an example of the first kind before human beings came to know about its effects and travelled to outer space. Any time a speaker spoke of cats on mats, his utterance's satisfaction conditions included the fact that this expression could only be true (satisfied) in a situation in which the cat and the mat where subject to the Earth's gravitational force. Nowadays it could be an example of the second. Imagine that astronauts in outer space use 'x on the mat' as applying to objects being in contact with mats, or that, in some conversations, the relevant gravitational force is that of Mars. If that is so, then the assumption concerning gravitational force is dropped or modified in some uses. Travis focuses on scenarios of this second kind. Although there is no essential difference between them–a constant assumption necessary for applying an expression E could stop being constant if a new practice in which to use E was developed– I will mostly use the term 'occasion-sensitivity' for examples of the second kind. This is compatible with saying that occasion-sensitivity shows that truth-conditions depend on a number of implicit assumptions. These assumptions vary across occasions.

The proliferation argument bears important similarities with Searle's line of reason-

[11] Waismann's (1951) thesis that words have an open texture is similar to Searle's background-dependence

ing. I distinguish three versions of the argument.

1.3.1 Ambiguity and ellipsis

Travis argues that occasion-sensitivity cannot be eliminated by distinguishing different senses for the predicates involved. The main reason is that for any sense that we can specify a further Travis case might be created. We can call these further Travis cases 'higher-order Travis cases'. This speaks against the following domestications of Travis cases.

If one takes it that the two utterances of 'The kettle is black' in the example have different truth-values but wants to keep the claim that non-indexical sentences express stable truth-conditions, one can look for an ambiguity. If some word there is ambiguous, then we have different expressions after all and the claim that meaning fixes truth-conditions is not threatened. According to this option, 'is black' ('is round', 'is my grandmother's', 'is green', etc.) would have various distinct meanings. Now, if this option is going to work, the distinct meanings must be immune to further Travis cases (the disambiguated expressions must be occasion-insensitive), for distinguishing two or more meanings counts as a rejection of occasion-sensitivity only on the assumption that the resulting meanings are not occasion-sensitive[12].

The problem, as I mentioned, is that Travis cases can be created for a variety of words, including 'superficially black', 'painted green', 'is usually used by my grandmother', 'is round in normal conditions':

> [About the possibility of substituting 'The leaves are green' by 'The leaves are painted green'] Or suppose they are painted, but in pointillist style: from a decent distance they look green, but up close they look mottled. Is that a way of painting leaves green? It might sometimes, but only sometimes, so count. So there would be two distinct things to be said in the presumed 'paint counts' sense of 'is green'. (Travis, 1997, p. 112)

There are, to be sure, several intuitive paraphrases of the content of the speech acts depicted in Travis cases. We can substitute 'The leaves are green' by 'The leaves are painted green', 'The leaves are naturally green', 'The leaves are superficially green', etc. The problem with these paraphrases is that we have not been told what kind of objects count as 'painted green' (what about pointillism?), 'naturally green' (what about

[12] Analogous considerations apply to approaches based on polysemy instead of ambiguity.

the presence of a bacteria that would make the leaf look green?), or 'superficially green' (how much of the surface? Where exactly?). If the disambiguations resemble these paraphrases, occasion-sensitivity has not been ruled out. So the proponent of the ambiguity view needs to show that the disambiguated meanings are unlike 'painted green', etc.

Recanati, following Searle, gives an analogous argument:

> The contextualist emphasizes the unending potential for variation in order to point out that the (modulated) meaning of an expression always depends upon the context and cannot be fixed simply by complexifying the expression and 'making everything explicit'. Thus the contextualist gives the following sort of example in support of the irreducibly contextual character of the interpretation process. 'John took out his key and opened the door' is interpreted in such a way that John is understood to have opened the door with the key ; this we get through modulation of 'open the door' which is understood via the contextual provision of a specific manner of opening. Can we make that explicit in the sentence, so as to get rid of the context- dependence ? Not quite : If we say 'He opened the door with the key' the new material gives rise to new underdeterminacies because it, too, can be variously modulated. The key may have been used as an axe to break the door open as well as inserted into the keyhole (Searle 1992 :182). And if we make the way of using the key explicit, further indeterminacies will arise, and different meanings will emerge through modulation. However, when language is actually used and something is said, there is a definite context (both linguistic and extralinguistic) and it is finite. In virtue of the context, the various expressions used in it get a definite meaning. (Recanati, 2010, p. 47)

A similar problem arises if one takes the predicate to be elliptic (or incomplete in Bach's sense). For any completion, or any explicit form, new Travis cases can be generated. So substituting 'is black' for 'is originally black', etc. is not going to work[13].

1.3.2 Parameters

In view of the context-sensitivity of natural language one could claim that, in the kind of cases Travis has put forward, meaning identifies a set of parameters that, given a context, determine a referent/extension/intension. Let us call this the parameter approach.This approach maintains the traditional view by assimilating Travis cases to semantically

[13] This point is stressed by Cappelen and Lepore (2005) in his argument against the stability of moderate contextualism. According to them, if one accepts some Travis cases, he is thereby committed to accept an indefinite number of them.

driven context-sensitivity. According to it, meaning determines truth-evaluable content (at a context of use).

Again, Travis argues that for any set of parameters that we can specify it is to be expected that a new Travis case can be generated. If so, then that set of parameters didn't fix truth-conditions after all:

> But pick any set of further factors you like. It will always be the case that, on some speakings where values of those factors are fixed,yet further considerations may dictate taking yet further factors into account, or considering the given factors in any of many new ways. [...] Since this is true for any packet, throw in the further factors and the situation repeats itself. (Travis, 1981, p. 53)[14]

If the factors that one would need to include are similar to the dimensions that we could make explicit–origin of the colour, part, and the like–, then higher-order Travis cases can be created for the new semantic item. This line of reasoning would work by putting forward examples. The proponent of the parameter approach needs to put forward a set of parameters. Then, the advocate of occasion-sensitivity will argue that there are further Travis cases. But Travis is skeptical that one can find a satisfactory set of parameters. He argues that the predicates he uses are unlike other cases of context-sensitivity for which it seems plausible to find a parameter that determines referent–indexicals, mainly (Travis, 1997, pp. 115-117). According to Travis, whereas it is part of the meaning of an indexical as 'I' that this word is a device for referring to the speaker, it is not part of the meaning of, for example, 'green', that this word comes with a set of parameters–green simply refers to a colour. This point is reinforced with an argument by Recanati (2010, p. 57). Recanati imagines a language containing a predicate 'red' that does not involve any context-sensitive variable. This language is used to describe the colour of different objects. Since the objects have parts and their surfaces do not always have a homogeneous colour, the question will often arise as to what part of the object is relevant for the colour attribution. So, the fact that colour attributions can describe different parts of the object does not imply that they are similar to indexicals. Even if they simply referred to a colour they would behave the same.

Because of Travis' rejection of the parameter approach, Davies (2014) holds that Travis thesis' is that meaning is neither a content nor a character [15] . However, it is im-

[14] Travis does not provide examples for predicates, but only for the description 'the boy'.

[15] I think that this is equivalent to the version of Semantic Propositionalism that I have presented in the previous section. Davies' idea is that, according to Travis, meaning can neither be identified with a truth-evaluable content nor with something that determines a truth-evaluable content.

portant to note that Travis' argument against the parameter approach is intended to apply whether the role of the parameter is to determine a content (a proposition, as happens with indexicals) or a truth-value (in which case they would be included in the index). Let us suppose that we have a context sensitive predicate F. On the first kind of account, the relevant set of parameters is plugged-into the meaning of F. A sentence as 'a is F' is conceived as a function from contexts to contents. F expresses different contents at different contexts, with these contents being determined by the set of parameters. By contrast, on the second kind of account, the set of parameters is part of the index of evaluation. There is no intermediate step. In a context of use, the (constant) proposition expressed by 'a is F' is evaluated against an index of evaluation including the relevant set of parameters.

The proliferation argument is directed against both accounts at once. Regardless of whether the set of parameters is plugged-into the meaning or the index of evaluation, this argument calls into question the assumption that there is a definite set of parameters such that, keeping their values fixed, the truth-conditions of all utterances of the sentence 'a is F' will be identical (leaving time aside), no matter what else goes on[16].

1.3.3 Properties

At this point, it can be useful to distinguish two questions. First, there is the question about the determination of truth-evaluable contents. The previous arguments provide reasons for claiming that the semantics of a well-formed declarative sentence S does not determine the truth-evaluable content that S expresses on an occasion of use. Second, there is a further question about the existence of occasion-insensitive truth-evaluable items fit to be the contents of our utterances. One might agree with Travis on the first point but claim that there are occasion-insensitive propositions to be had[17]. On this view, which proposition a given speech act expresses simply is not a semantic matter. But the idea that our utterances express propositions–in the sense of Fregean, complete thoughts–remains untouched.

A way of fleshing out this possibility consists in distinguishing predicate meaning from genuine properties. We have reasons to think that the linguistic meaning of some predicates is compatible with variation in the predicate's contributions to the truth-

[16] In chap. 2 I will examine current versions of this approach (indexicalism), as well as an ambiguity view. To my knowledge, Travis does not apply this argument to any actual account dealing with predicates.

[17] According to Borg (2012), this point distinguishes contextualism and occasionalism.

conditions of sentences in which it appears. If properties are thought to make constant contributions–to always have the same content, to always establish the same partition among objects–, then it seems that the same predicate must be expressing different properties across uses.

Travis' view is that this option is problematic. In a book about Wittgenstein, he presents what can be considered a version of the proliferation argument concerning properties:

> If 'weighs 79 kilos' or 'contains milk' refers to a family of S-use insensitive properties, the question is what the members of this family might be. Suppose for the sake of argument that Odile spoke of some one of these properties in speaking truly of the refrigerator. Then we may speak of 'the property Odile (then) attributed to the refrigerator'. (Or we might prefer to speak of 'the property of the refrigerator in virtue of which what Odile said was true'.) Call this property Q. We now have the means to attribute that property to other items, and to consider whether other items, or the refrigerator in other states, have it. Will Q be an S-use sensitive property? The key point is this. In deciding that Odile spoke truly of the refrigerator, we solved one problem, or a few, about how to sort things into those containing milk and those not. But in principle there may always be more. In fact, we can easily think of countless more: suppose that what was in the refrigerator was a slice of cheese or cheesecake, or a vial of secretions of rabbit mammary glands, or a pint bottle of thoroughly coagulated (or very sour) milk, or synthetic milk with the same molecular structure as milk but which had never seen a cow; and so on. (Similarly, consider the property we decided Hugo has, even after lunch, if—but only if— his morning weighing showed—or suitable ones would show—79 kilos. Would he have that property if the earth's gravitational force were to be halved overnight? Or if, to take a Wittgensteinian example, he began occasionally to grow and shrink spectacularly and for no apparent reason?).
>
> Most of these problems, and others, are problems for Q as well. (Travis, 1989, p. 23)

Let's assume that, in order to explain away Travis cases, we postulate that two different properties have been attributed to the object in question. Then, we can ask whether the property attributed in the first case is one that some other object has. But now, we might find ourselves unable to answer such a question, for new questions arise, questions that will be solved in different ways on different occasions. Travis' conclusion is that this shows that the alleged property was not occasion-insensitive. This argument

calls into question the viability of an approach in which there are occasion-insensitive properties to be had, for now it is on the proponent of such an account to show that his properties are unlike the ones we ordinarily seem to attribute–green, weighs 79 kilos, milk, and so on–, yet that they are the ones we in fact attribute[18].

I think that this argument only works on the assumption that the metaphysics of meaning cannot outrun its epistemology. The idea here is that P and Q are no better than 'milk' or 'weighs 79 kilos', and this is shown by showing that we are unable to answer certain questions. However, the conclusion only follows on the assumption that our being unable to answer certain questions is proof that these questions lack an answer. And this is true only if the epistemology of meaning imposes a limit on its metaphysics[19] Nonetheless, rejecting the assumption has the unpalatable consequence that the posited properties are not fully transparent to us [20]

1.4 Assessing the argument

I think we can reconstruct Travis' most general critique as targeting the following principle[21] :

> *Truth-Conditional Compositionality for Structured Representations:* The truth-conditions of a token of a structured representation S are determined (exhausted) by features of S's type.

[18] Travis' notion of property (a way for things to be, in his terminology) is a sui generis but common sense one. A property simply is a way a thing can be. They are what we usually talk and think about. We have words for them (or, if we don't, we can always coin new words). As a consequence of this view on properties, properties are occasion-sensitive. If F is a property and 'F' the word we have for that property, then it is occasion-sensitive whether an object can be rightly described as 'F' because it is an occasion-sensitive matter whether the object is F . Objects are green, round, etc. on occasions, depending on the purposes of the classification. Because of this, Travis' view is not only about words but about what words mean (as he explicitly says in (Travis, 2008a, p. 178)).

[19] This view is not unprecedented. According to Glüer (2011), Davidson held a similar view on the relation between the metaphysics of meaning and its epistemology.

[20] I will come back to this argument in chap. 3.

[21] Travis (2000)targets a notion of content according to which contents can be individuated both in terms of structure and in terms of truth-conditions. He argues that items that are identified by structure (as sentences) are typically such that their tokens can express different truth-conditions.

This principle generalizes Semantic Propositionalism. Semantic Propositionalism applies to sentences. But one could hold that more complex logical forms (ones to which, for example, a time, a place, etc. has been added) are such that their semantics determines their truth-conditions, perhaps at a context of use. Or that structured propositions do. Against this, Travis holds a general underdeterminacy claim.

Travis' explicit target is Frege's notion of a thought (Frege, 1956). Thoughts, according to Frege, differ from natural language sentences in that they are complete. Frege considers as a an example the present tense. In order to apprehend the thought expressed by a sentence in the present tense, one needs to know when it was uttered. Because of this, Frege holds that the time is part of the thought–even though it is not always part of the sentence that the speaker uses to convey it.

In its most general form, the proliferation argument is about the prospects of either finding disambiguations of the sentence–classical ambiguity, distinguishing properties– or complexifying syntactic structures –by adding some additional structure or context-sensitive parameters. Both options rest on the assumption that there is some set of features that a (non-indexical) syntactic structure or a structured proposition can have such that no matter the tokening context it will always express the same truth-conditions. Travis targets this assumption:

> What [this assumption] says is that for each statement there is some specifiable set of representational features which decides all that is decided as to when the statement would be true. [...] The idea would be this: take any statement and ascribe to it any set of specifiable representational features you like; then two or more statements might all share those features, yet differ in what they said, and hence in when they would be true. We start to find features we must ascribe to a given statement, S, by contrasting it with other possible ones. We find a statement that differs from S in when it would be true; we thereby see the need to ascribe to S a feature of a certain sort–F, say. We then find a statement with F which still differs from S in when it would be true. So we assign S another feature, F'. And so on. But in the envisioned situation, no matter how we start, or how we continue this process, there is no way of bringing to a halt the sequence of statements which, sharing more and more representational structure with S, nonetheless differ in content from S. (Travis, 2000, pp. 35-36).

We can reconstruct the argument as follows:

1. The semantics of item I (for example, a sentence) is compatible with

variation in truth-conditions. (First-order Travis cases justify this premise).

2. I is substituted by I*, where the semantics of I* explicitly solves the indeterminacy responsible for the previous variation. Alternatively, a parameter solving the doubt is added to the semantics of I.

3. Again, the new semantics involve indeterminacies concerning truth-conditions. (Higher-order Travis cases provide justification for this premise).

Conclusion. There are no reasons to think that there is a point at which we will get an item whose tokens always express the same truth-conditions.

There can be different formulations of premise 2. I have already mentioned several of them. Here is a summary focusing on predicates:

- The predicate involved in a Travis case can be substituted by a more precise predicate. For example, 'is green' could be substituted by 'is superficially green'.

- The semantics of the predicate can be made more precise. For example, instead of a simple rule determining an intension ('the semantic value of 'is green' is a function from possible worlds to the the set of green objects at those worlds'), we could have a more sofisticated semantic theory involving context-sensitive semantic rules ('the predicate 'is green' is true of object x if and only if, for a contextually determined part P of x, x is green at P').

- The predicate is taken to refer to several distinct properties. For example, the predicate 'is green' could be taken to refer to two different properties, F and G.

The underlying problem concerns the relation between structures, or better, representational forms identified by structure, and satisfaction-conditions. Thus viewed, the moral is that no matter how much one complexifies a given structure, this structure will always be compatible with variation is truth-conditions.

The inferences from 2 to 3 and from 3 to conclusion depend on semantics being specifiable, on the possibility of having a sentence encoding the item for which Travis cases can be created. Let me take as an example the parameter approach. The idea that there is a set of parameters fixing the extension of a predicate in a context of use usually motivates the claim that we can have semantic theories matching sets of parameters and extensions. The set of parameters can be thought of as something that we can specify.

Moreover, we could find a sentence encoding its value at a context of use ('The leaf is superficially green in such part'). For contents we could not in principle specify we could not create a Travis case.

Taking this feature about specifiability into account, the outcome of the proliferation argument can be paraphrased as follows:

> *The proliferation argument:* No specifiable content (no content that can be articulated in natural language) has intrinsic truth-condition.

Now, it could be argued that the non-specifiability of occasion-insensitive items is purely contingent. In order to get a paraphrase not open to new Travis cases we would need to specify all the relevant details, something that outruns our cognitive capacities. The problem with taking non-specifiability to be purely contingent, however, is that what gives ground for creating further Travis cases are the new expressions one brings in in the paraphrases. Thus, if one was to specify how much of the surface is relevant ('green on at least 90% of its surface'), by using measures one would be introducing the need for a standard of precision[22] .

How does the proliferation argument fare against Truth-Conditional Compositionality for Structured Representations? I think that the possibility of iterating Travis cases–i.e., of imagining Travis cases for sentences that overcome the underdeterminacy of 'The leaves are green' in the example previously presented–raises serious doubts that Truth-Conditional Compositionality is true for natural language sentences. It also raises some doubts that adding standard parameters, such as time and place, to the logical form of the sentence, or to the structured proposition expressed by the sentence at a context of use will result in a logical form that respects Truth-Conditional Compositionality. Regarding colour predicates, adding a part parameter is insufficient, for we can create Travis cases having to do with the illumination conditions and the observation distance, origin of the colour... and more parameters will be needed as more examples are imagined.

Now, as I have already mentioned, the argument does not apply to propositions involving properties that are non-encodable in natural language and not fully transparent

[22] Arguably, the source of occasion-sensitivity is not lack of precision but the relation between descriptive contents and its application conditions. We start with a descriptive content. We try to apply it. And we see that there are different ways of applying it. Adding more descriptive content doesn't seem like a good option, for then we have the problem of how to apply this new descriptive content. This line of reasoning, echoing some Wittgensteinian worries, is briefly discussed in chap. 3.

to us. I think the argument is also insufficient against more general, or non-standard, parameters. Although it seems reasonable to conclude that positing, let's say, a parameter for part and another for illumination conditions in the logical form of colour terms will not get us a content that is immune to further Travis cases, it is difficult to see whether the same would happen in case we added more generic parameters such as 'what is at stake'. Moreover, one could go for less standard notions of content, such as Austinian propositions including situations (or occasions)[23] . Similarly, but in a more Searlian spirit, one could define utterance content as including not only a structured proposition but also a background of implicit assumptions (some of which, as Searle claims, might be non-representational).

1.5 Wittgensteinian themes

Travis' view on language is shaped by his interpretation of Wittgenstein and Austin. From Austin he took the claim that sentences are not, per se, true or false, the core idea behind occasion-sensitivity. As to Wittgenstein, the proliferation argument bears important similarities with the discussion on rules in the *Philosophical investigations*[24] : 'A rule stands there like a signpost. Does the signpost leave no doubt about the way I have to go? Does it show which direction I am to take when I have passed it, whether along the road or the footpath or cross-country? But where does it say which way I am to follow it; whether in the direction of its finger or (for example) in the opposite one?' (Wittgenstein, 2009, 85). Rules, and in Travis discussion this means semantic rules, are open to different interpretations, and can be followed in different ways. However, as Wittgenstein shows, adding interpretations (disambiguations, for Travis) of the rules does not solve the problem, for the new interpretations are themselves open to different interpretations. We, nonetheless, when trained in the use of language and situated in the proper circumstances, have no problem in following rules. Let me develop the parallelism with some more detail.

Signposts, of themselves, as the objects they are, do not tell us what the correct way

[23] One could read Travis' claim that speech is situated as calling for such a notion.

[24] Travis sometimes makes reference to this Wittgensteinian origins. See for instance Travis (2000, p. 25)

of following them is. We can give alternative interpretations. The same happens with the rule governing the use of 'green' (its meaning). We are presented with different scenarios, one in which the rule is naturally followed in way A, and another in which it is naturally followed in way B. The rule 'green' stands there like a signpost, and leaves doubts about how to follow it–does it apply to an object in case green is its natural colour? Or in case most of its surface is green? But then perhaps we could add something to the signpost so to remove the doubt about how to follow it. The proposal might be that that 'something' is a rule for interpreting the rule or a sharper meaning. In the case of 'green', perhaps we should have a second rule such as 'green on the surface'. To put it more explicitly, the rule could be something like: 'In this context, apply 'green' only to objects that are superficially green'. This (contextually determined) rule has removed a doubt: natural colour does not matter, pay attention to the surface. But it has not removed all possible doubts: how much of the surface? Is dust part of the surface? So we need a third rule that explains how to apply the second one: 'green on most part of the object's surface, excluding dust'. And so on. So, this way of thinking seems to lead to a regress.

The moral of the rule-following considerations, for Wittgenstein, is the rejection of the idea that understanding a word, and acting in accordance to a rule, is a case of interpretation. In Travis' philosophy, our parochial ways of thinking decide how it is more reasonable to follow the rule in particular circumstances. Their role is indispensable.

As a result, when presenting occasion-sensitivity Travis doesn't speak of one sentence expressing different propositions or contents in different occasions of use. Instead, he often speaks of words bearing different understandings. Understandings are not further representations. Although Travis does not provide a precise characterization, understanding an utterance seems to mean being able to act in accordance to it. The idea that understandings are abilities is explicit in Travis' discussion of what it takes to understand an utterance (Travis, 2000, pp. 207-211). When someone understands some words, he thereby gains an ability to see when they would be true. This ability needs not, according to Travis, be further explained as grasping a proposition. Being able to deploy the ability depends on two things: wordliness (or knowledge about the world and social practices) and reasonableness.

Reasonableness plays an important role. According to Travis, in particular in his discussions about Wittgenstein, the correct understanding is the one that seems reasonable–to thinkers like us. Whether some painted leaves count as 'green' on a certain occasion depends on whether this is a reasonable way of taking them, given the point of the ut-

terance and for thinkers like us. Here are some examples:

> The standards of correctness (truth where relevant) governing a given use of 'chaussure' in saying something are those it would be most reasonable to see as imposed, given the circumstances of that use (Travis, 2006b, p. 29).

> Content is inseparable from point. What is communicated in our words lies, inseparably, in what we would expect of them. How our words represent things is a matter of, and not detachable from, their (recognizable) import for our lives. Calling something (such as my car) blue places it (on most uses) within one or another system of categories: blue, and not red, or green; blue, and not turquoise or chartreuse; etc. If I call my car blue, the question arises what the point would be, on that occasion, of so placing it; or, again, what one might reasonably expect the point to be; what ought one to be able to do with the information that the car so classifies. What I in fact said in then calling my car blue is not fixed independent of the answers to such questions (Travis, 2006b, p. 33).

> Were Pia said Sid to grunt, things count as being as she thus said just where, on our shared (and parochial) understanding of reasonable, that is the reasonable view of things being as they are (Travis, 2006b, p. 107)

Travis' idea of reasonableness involves two components: (i) how it is reasonable to understand an utterance depends on the point of the utterance, on what we are doing with it, on the possibilities of action it affords, and (ii) what is reasonable is relative to specific ways of thinking–what he calls 'parochial ways of thinking'.

1.6 Is occasionalism a tenable approach?

On a reading consistent with some o his texts, Travis' view can be understood as involving the claim that the truth-value of our utterances depends on what thinkers like us find reasonable, given the surroundings–for example, given the point of uttering the sentence. In this section, I will argue that occasionalism thus conceived runs the risk of precluding the objectivity of content[25].

[25] I will not try to identify or assess Travis' answer to this problem. My aim here is to argue that there is a potential problem for occasionalism.

The problematic view is suggested in passages as the following, together with the ones quoted in the previous section:

> If circumstances of a stating matter to what was stated along the lines indicated here—lines on which what was stated is fixed by what one then had a right to expect of things being as stated (of a cat's having mange, say) — then there is in principle no end of opportunities for circumstances of a stating to matter to what was stated. There is no point at which circumstances choose for us some truth-evaluable item which is itself immune in principle to admitting of different further understandings — no point at which, through appeal to circumstance, we arrive at the sort of invisible, intangible truth-bearer (what Frege called a 'thought') which, Frege held, was the only thing that could really make a determinate question of truth arise. (Travis, 2008a, p. 6)

And, contrasting Frege's and Wittgenstein's view:

> What Frege leaves out on Wittgenstein's view–in fact, explicitly denies–is that 'true' and 'false' are in the first instance evaluations of particular historical events–speakings of words on particular occasions, in particular circumstances–and of the fittingness of the words for those circumstances. (Travis, 2008a, p. 254)

Or:

> On the present idea of reason and responsibility, when, in given circumstances, one speaks of a thing (my car, say) in given terms (one calls it blue, say), what one thus says as to how things are, so when one would have spoken truly, is determined by what it is reasonable to expect of that particular describing — reasonable to hold one thus responsible for — given the circumstances of its giving; what expectations would, in those circumstances, reasonably be aroused. (Travis, 2006b, pp. 30-31)

Travis rejects Fregean thoughts–contents whose truth-conditions are intrinsic, that do not admit of different understandings. The proper bearers of truth-conditions are utterances, not thoughts or propositions. The problem now is that whether an utterance is true or false of a certain state of affairs seems to depend, for him, on whether it fits the occasion and on what we find reasonable to expect. But if that is so, our psychology might end up playing a too crucial role.

I will be concerned here with the relation between there being meanings that determine correct application and the objectivity of judgement, and assess occasionalism in light of Wright's reflections about Wittgenstein's rule-following considerations[26] .

[26] See Wright (2001a).

Adapting McDowell's metaphor (McDowell, 1984) to occasion-sensitivity, the problem is how to steer a course between Scylla and Charybdis. Scylla would be the idea that words, of themselves, impose a condition on the world. Charybdis would be the idea that no item imposes a condition, that it all depends on how we happen to react. Once we have seen what is wrong with Scylla (sentences impose different conditions on different occasions) we risk steering too close to Charybdis. By abandoning the Fregean notion of content, occasionalism runs the risk of ending up claiming that whether a speech act is true or not depends, at least to a great extent, on how we happen to react to it. And this would undermine the objectivity of judgement.

Many of our judgements and utterances aim to objectively describe how things are. Judgements like 'The Earth revolves around the Sun', 'Water is colourless' or, in more mundane examples, 'There is a chair in the living-room', 'The leaves are green', etc., seem to reflect how things are independently of us—of our best opinion or the state of our investigation. As Wright (2001b) holds, investigation-independence characterises objectivity: the objectivity of decidable statements consists in the possession of investigation-independent truth-values. The notion of objectivity marks a distinction between our (perhaps justified) opinion and the true facts of the matter (Wright, 2001b, pp. 33). As way of comparison, the common assumption that taste or comic discourse is non-objective is motivated by the intuition that there are no facts of the matter in these domains but only opinions, which abolish the distinction between how things are in themselves and how we take them to be.

Now, Wright has argued that possession of objective truth-values requires possession of objective meaning (Wright, 1984). If that is so, the objectivity of the truth-value of the sentences mentioned in the previous examples requires the objectivity of the meaning of the terms 'water', 'revolves', 'chair', 'green', etc. What is the objectivity of meaning? According to Wright (2001b), objective meaning can be conceived as a pattern that determines correct uses in unconsidered cases, independently of our opinion. "The pattern is thus thought to of as an extending of itself to cases which we have yet to confront" (Wright, 2001b, pp. 34). Objective meaning is meaning that extends of itself to unconsidered cases. Since it extends of itself, it is independent of our opinion or the state of our investigation. It is judgement-independent. An example: let's assume that the meaning of 'chair' is objective. Then whether this word can be correctly applied to an object in room 4010 of the Faculty of Philosophy —whether 'That is a chair', used to describe that object is true or false— is independent of our opinion. We might never enter that room and be unaware of the objects inside. Nonetheless, if meaning is objective,

then it extends to this case. Meaning determines application conditions independently of us.

The identification of meaning with functions provides a good model for objective meaning. Let us assume that we identify the meaning of a predicate with a function from objects to truth-values. Thus, the meaning of 'is a chair' maps objects onto the truth-value True (those that are chairs) of the truth-value False (those that are, e. g., tables, couches...). It covers past and future uses alike, considered and unconsidered cases. For every object in the domain, it yields one, and only one, verdict, one that is independent of our judgements. It extends of itself, independently of our judgements.

Here is Wright's argument:

> [The truth-meaning platitude] If the truth value of S is determined by its meaning and the state of the world in relevant respects, then non- factuality in one of the determinants can be expected to induce non- factuality in the outcome.[...] A projectivist view of meaning is thus, it appears, going to enjoin a projectivist view of what is for a statement to be true. Whence, unless it is, mysteriously, possible for a projective statement to sustain a biconditional with a genuinely factual statement, the disquotational schema "'P', is true if and only if P" will churn out the result that all statements are projective. (Wright, 1984, p. 769)

The argument can be reconstructed as follows:
First, from the non-objectivity of meaning to the non-objectivity of truth:

P1. The truth-value of a statement S is a function of the meaning of S and the state of the world. (Meaning-Truth Platitude)

P2. Non-factuality in one of the determinants (meaning or state of the world) induces non-factuality in the outcome (truth-value). (MP 2, 1)

P3. The meaning of S is non-factual. (Hypothesis)

$\therefore$ The truth-value of S is non-factual.

Second, from the non-objectivity of truth to the non-objectivity of the statement:

P1. The truth-value of S is non-factual ("'S' is true" is non-factual).

P2. 'S' is true if and only if S. (Disquotational schema)

P3. The non-factuality of one statement of a biconditional entails non-factuality in the other statement of the biconditional.

∴ S is non-factual.

What about occasion-sensitivity? The outcome of Travis' critique is that the traditional notion of content is ill-conceived. So we'd better do without it. This would be unproblematic if the role of the occasion were that of identifying a further semantic item–such as a property–that extended of itself to unconsidered cases. But the proliferation argument is intended to cast doubts on the existence of such items. Now, if there is a gap between our representational items (sentences, propositions) and truth-evaluability, it seems that either we are unjustified in taking some items as true (and others as false) or there is some extrinsic source of normativity. But what could play this role? We might intuitively take some speech acts as true, others as false, but our so reacting is insufficient to bridge the gap–if truth is partially determined by our reactions, then it looses its objectivity.

At this point, there are two possibilities. First, one can abandon the idea that there is some level of content that extends of itself and try to ground objectivity on the reasonableness of our reactions. According to this option, uttering a sentence would not be a matter of expressing a structured content with intrinsic truth-conditions, or anything similar. Some speech acts are reasonably taken as true, others are reasonably taken as false. Second, one can reject the traditional approach to content but keep the idea that we arrive at something that is similar to a Fregean thought in that it is only to the world to decide whether it is true or false.

Travis' way of writing sometimes suggest the first option. However, instead of trying to ground objectivity on our reactions or judgements and on the notion of reasonableness, an an advocate of occasion-sensitivity can follow the second option and revise the notion of content. In chapter 4 I will go for this second option. The main idea of the outline sketched there is that the practice, or the activity, in which a sentence is used provides the constraints needed to reach something truth-evaluable. Truth-evaluable content can be seen as a pair including a structured proposition and an activity. Thus understood, content can extend of itself.

1.7 The contextualist debate

During the last decades there has been a growing interest in the semantics-pragmatics divide. On the one hand, a number of philosophers have provided philosophical and linguistic arguments showing that natural language contains more context-dependent expressions than it was thought. It has been claimed that quantifiers (Stanley and Szabó, 2000), knowledge attributions (DeRose, 1992), and epistemic modals (Egan et al., 2005), to cite just a few examples, are context-sensitive in some sense or other[27] . Semantic frameworks deal now with a great amount of context-dependence. On the other hand, the role of pragmatics in the determination of truth-conditional content has been the focus of much debate. Once it is acknowledged that an expression has different senses, or expresses different contents, at different contexts of use, it remains to be determined whether this variation is semantically mandated.

Borg (2012) distinguishes five current positions about how to draw the semantics-pragmatics divide. These five positions have explained Travis cases in different ways.

Two of these views can be considered semantic. Minimalism, the first of them, is the current representative of the traditional view. According to this view, the semantic content of a well-formed declarative sentence is, in absence of indexicals, truth-evaluable. Thus, minimalists[28] take the shift in truth-value in Travis cases to be merely apparent. By contrast, indexicalism[29] takes it that there is a shift in truth-value, but revises the logical form of the predicate involved so as to assimilate it to an indexical by positing some context-sensitive hidden variables.

The other three views can be considered pragmatic[30]. According to these views, the underdeterminacy of natural language calls for truth-conditional pragmatics[31] . Contextualism disagrees with minimalism in the assumption that truth-conditional content can be recovered via semantics alone. According to contextualists[32] , at least in some

[27] The discussion is often about whether the linguistic data are better account for by indexical or non-indexical contextualism. I take it that both options agree that the expression in question is context-sensitive in a broad sense.

[28] See Sainsbury (2001), Berg (2002), Cappelen and Lepore (2005), and Borg (2004, 2012).

[29] See Stanley (2000), Szabó (2001) and Hansen (2011)

[30] Relativism can also be seen as a semantic view. I am classifying it as pragmatic because, according to it, the truth-conditions of utterances depend on something beyond meaning.

[31] Truth-conditional pragmatics is the view that the meaning of a sentence does not determine the truth-conditional content of an utterance of it (even after fixing the referent of indexicals), for truth-conditions need to be pragmatically supplemented.

[32] See Sperber and Wilson (1995), Recanati (1989, 2001, 2004), and Carston (1988, 2002).

cases, the recovery of the proposition expressed by an utterance, what is said, even in absence of indexicals, must go via pragmatic interpretation, where pragmatics usually has to do with speaker meaning. Typically, a contextualist claims that, in a Travis case, the sentence expresses different propositions or, in a fairly common terminology, that what is said in the first utterance is different from what is said in the second. It is common to distinguish moderate and radical contextualism. According to radical contextualism, the recovery of a truth-evaluable content must always go via pragmatic interpretation. According to moderate contextualism, this is only sometimes the case.

Relativism, also called non-indexical contextualism, agrees that the semantic content of a sentence might not be truth-evaluable, but doesn't make use of a further kind of content. Instead, relativism takes it that truth is relative to a variety of parameters besides possible world. A relativist[33] , or a nonindexical contextualist, would account for Travis cases by claiming that the truth of the sentence is relative to some parameter (as the purposes of the conversation) whose value shifts across contexts.

The fifth view is occasionalism. According to Borg, occasionalism claims that there is no determinate content outside a context of use. Despite of being very close to radical contextualism in its general context-dependence claim, occasionalism rejects the assumption that there is something deserving being called 'the proposition expressed'– unless 'proposition' just means truth-conditions[34].

Although I find this classification very useful, I have some doubts that occasionalism is on a par with the other approaches, for Travis has mainly put forward a negative, not a positive, proposal. As Borg acknowledges, Travis has sustained an attack on a traditional notion of content that seems to be shared by minimalists and contextualists. Despite the fact that most people in the debate have used Travis cases to argue that semantics doesn't determine truth-conditions (or have disputed that the examples support such a claim), occasion-sensitivity, as it is presented by Travis, is intended to undermine standard notions of content such as structured propositions or Fregean thoughts. In this sense, occasion-sensitivity motivates a rejection of what can be called Semantic propositionalism (the thesis that the semantics of a sentence determines the proposition this sentence expresses, relative perhaps to a context), but also of Pragmatic propositionalism (the thesis that the features of the use of a sentence determine the proposition that sentence, on that occasion, expresses), for what is problematic is the notion of (struc-

[33] See Predelli (2004, 2005b,a) and MacFarlane (2007, 2009)

[34] Besides Searle's and Travis' positions, I think Recanati's meaning eliminativism could be considered an occasionalist approach to language.

tured) proposition.[35]. Occasion-sensitivity is supposed to show that, in the case of natural language, individuating structures is not a way of individuating truth-conditional content. Additionally, assuming that language is occasion-sensitive and therefore that we cannot specify contents with intrinsic truth-conditions, occasion-sensitivity is intended to cast doubt on the existence of a further level of representation for which Travis cases could not arise. As Wittgenstein would have it, 'You have no model of this inordinate fact, but you are seduced into using a super-expression' (Wittgenstein, 2009, 192). The very notion of there being structured propositions with context-independent truth-conditions to be had is called into question[36] . For Travis, this motivates a generalization: any structure (including propositions and mental representations) is compatible with its tokens having different truth-conditions.

So far Travis has put forward a negative view. However, it is not clear how a positive occasionalist view should look like. Travis holds that being true and false are properties of tokens. However, he also speaks of words bearing 'understandings' and of truth being relative to the character of the occasion–something very close to certain forms of non-indexical contextualism. So it is difficult to identify what his version of occasionalism amounts to.

I think that an advocate of occasion-sensitivity can opt for two kinds of approaches: he can stick to the claim that being true and false are properties of our utterances having to do with what it is reasonable to expect, or he can put forward a non-standard notion of content, as it is suggested by the idea that truth is relative to the character of the occasion. I will opt for the second. However, in Borg's classification, my proposal, as well as Travis' remarks on how truth depends on the character of the occasion, could be considered a form of relativism.

[35]It is also common to identify truth-conditional content not with structured propositions but with sets of possible worlds. Travis' criticism does not apply to this notion of proposition.

[36]Travis often speaks of what an utterance says. However, the locution 'what is said' is not equivalent to 'the proposition expressed'. Travis calls propositions 'shadows' and aims at 'unshadow' thought.

1.8 Conclusions

Occasion-sensitivity is a phenomenon of truth-conditional variability not driven by semantics. In this chapter I have presented Travis' arguments to the effect that natural language is occasion-sensitive. The first argument is provided by so-called Travis cases. I have called the second 'the proliferation argument'. This argument is directed against several ways of making the truth-value shift exhibited in Travis cases compatible with truth-conditional semantics. The argument casts doubts on the possibility of specifying an item for which no Travis case can be created.

After reconstructing Travis' negative proposal, I have related Travis' occasionalism to the contextualist debate. I have argued that Travis' position is more radical than standard contextualism, for Travis casts doubt on the assumptions that there is a stock of structured propositions with context-independent truth-conditions. Nonetheless, I have noted the difficulty of identifying Travis' positive proposal.

I have raised a problem for a possible reconstruction of some of Travis' passages. The reconstruction would have it that truth-value depends on what thinkers like us find reasonable. Following Wright, I have understood objectivity as judgement independence. The problem with the reconstruction is that it makes truth-conditions dependent on our judgements, which calls into question objectivity. However, I have noted that occasionalism need not be reconstructed along those lines, for occasion-sensitivity calls into question a standard notion of content, but other notions are possible.

Chapter 2

Truth-evaluable content and the role of semantics

Travis cases challenge the claim, often assumed in truth-conditional semantics, that semantics determines truth-evaluable content (call this principle 'Semantic Propositionalism'). After characterising Semantic Propositionalism and distinguishing it from the Principle of Compositionality I will assess two families of arguments intended to deactivate the challenge. On the one hand, minimalists call into question the assumption that Travis cases involve a shift in truth-value. I will argue that the minimalist arguments found in the literature are defective for various reasons. On the other hand, proponents of context-dependent analyses argue that Travis cases can be accounted for within the framework of truth-conditional semantics. I will argue that these accounts, when properly scrutinised, either are insufficient to explain the variety of Travis cases found in the literature or do not respect Semantic Propositionalism and therefore do not touch upon the claim that semantics does not determine truth-evaluable content.

2.1 Semantic Propositionalism

It has been common to assume a semantic principle along the following lines:

> *Truth-Conditional Compositionality:* The truth-conditions of a well-formed

> declarative sentence S are determined by the semantics (or the meaning) of the expressions in S and the syntactic structure of S.

As an illustration of the implicit use of this principle, here are some lines from the opening pages of Heim and Kratzer's *Semantics in Generative Grammar:*

> A theory of meaning, then, pairs sentences with their truth-conditions. The results are statements of the following form:
>
> Truth-conditions
> The sentence 'There is a bag of potatoes in my pantry' is true if and only if there is a bag of potatoes in my pantry. [...]
>
> A theory that produces such schemata would indeed be trivial if there wasn't another property of natural language that it has to capture: namely, that we understand sentences we have never heard before. We are able to compute the meaning of sentences from the meanings of their parts'. (Heim and Kratzer, 1998, pp. 1-2)

Heim and Kratzer take the project of semantics to be that of attributing truth-conditions to sentences in a way that it explains the productivity of natural language (our ability to understand new sentences). The aim is to give an account of meaning in terms of truth-conditions that models linguistic competence.

In case S contains some indexical expression E, one should distinguish the character (linguistic meaning) of E from its content at a context. The linguistic meaning of indexicals is usually conceived, following Kaplan (1989), as a function from contexts to contents (referents). Given a context C, the content of E will be its referent at C. Although indexicals motivate a revision of the principle, its spirit can be maintained—the semantics of an indexical sentence S are supposed to determine, in a context of use C, the truth-conditions of S at C. In the first half of the paper, I will focus on a version of the principle for indexical-free sentences. I will come back to indexicality in section 4.

Truth-Conditional Compositionality combines two other principles[1] : Semantic Compositionality and what I will call Semantic Propositionalism. Semantic Compositionality is a principle about the relation between the semantic content of a (declarative) sentence and its parts:

[1] See (Searle, 1980), (Clapp, 2012b), (Davies, 2014), for similar distinctions.

> *Semantic compositionality:* The semantic content of a well-formed declarative sentence S is determined by the semantic content of the expressions in S and its syntactic structure[2] .

By contrast, Semantic Propositionalism concerns the relation between semantic content and truth-evaluable content. It is this principle I will be concerned with:

> *Semantic Propositionalism:* The semantics of a well-formed declarative sentence S determines a truth-evaluable content (a proposition, a truth-condition)[3] .

This truth-evaluable, semantically determined content has been traditionally identified with the literal content of an utterance of the sentence.

Semantic Compositionality and Semantic Propositionalism are different principles. One could be true with the other being false. For example, it could be case that the semantic content of a well-formed declarative sentence, determined by the semantics of its components and its structure, is a time-neutral content. This content would not be truth-evaluable—so not a proposition in the classical sense—unless a time of evaluation was provided. In this scenario, Semantic Propositionalism would be false, but Semantic Compositionality would be true[4] .

It has been fairly standard to think of semantic content as truth-conditional content or, in the case of sub-sentential expressions, as a contribution to truth-conditional content[5]. What else, if not 'truth-conditional content' could one mean by 'semantic content'? As an answer, semantic content is content determined by semantics, and the semantics of a given expression E are restricted to the properties E has qua type. Semantic properties are shared by all uses, or tokens, of the expression. By contrast, some properties are properties of the token, not of the type. For example, being written in black or

[2] One could have instead a principle of pragmatic compositionality. If so, the composition would be of pragmatically determined values.

[3] As I define Semantic Propositionalism, it is a principle about meaning determining truth-conditions in the sense of determining a content that can be evaluated as true or false, not about semantics determining an intension (a function from possible worlds and other features such as standards of precision, etc, to truth-values). See footnote 10. It could also be doubted that meaning can be identified with an intension.

[4] Bach's propositional radicals (Bach, 1994) are plausibly semantic compositional, but they are not truth-evaluable, for they are incomplete.

[5] Thus Lewis writes: 'Semantics with no treatment of truth-conditions is not semantics' (Lewis, 1970, p. 18).

being uttered in Madrid, being inappropriate, etc. In this sense, semantic content just is linguistic meaning. In the case of indexical sentences, semantic content can also be understood as the content determined by the semantics (character, linguistic meaning) of the sentence at a context of use. In this sense, the expression 'semantic content' is not short for 'whatever goes into composition and delivers truth-conditions': it refers to the content of the expression type or, for indexicals, to the value determined by the value of the expression type.

According to Semantic Propositionalism, truth-conditional content is a semantic matter. If this principle is true, being truth-evaluable, or determining a truth-evaluable content, is a property of sentence-types. Fix the semantics of S and you will have fixed a truth-evaluable content—or, for indexical sentences, you will have fixed something that fixes a truth-evaluable content at a context of use. Assuming the world is cooperative enough, you will have fixed the truth-value of S as well.

Semantic contents thus conceived, as being both semantically determined and truth-evaluable, are sometimes called 'minimal propositions'[6] . Minimal propositions have been the focus of much debate. On the one hand, it has been argued that minimal propositions don't play any role in communication or even in our cognitive lives, for the contents relevant for communication and the contents interlocutors are aware of, often, and perhaps always, go beyond semantics—they are pragmatically adjusted propositions. Consider again this often discussed example[7] . Imagine that, as a reply to the question 'Would you like something to eat?', the addressee says 'I've had breakfast'. Let us assume that the semantic content of this sentence is, roughly, that the speaker has had breakfast before the time of utterance. It seems that this semantic content would be true if he had had breakfast one week ago. Now, this sentence, in this context, communicates something else: that the speaker has had breakfast that morning. This is, arguably, the content relevant for communication and the one interlocutors are aware of.

On the other hand, the very notion of a minimal proposition has been challenged. Travis, Searle, and others[8] have argued that semantic content is not truth-evaluable con-

[6] I will reserve the term 'minimal proposition' for contents that are truth-evaluable and semantically determined where the semantics of all the expressions in the sentence, except for classical indexicals, demonstratives and descriptions is taken to be context-insensitive. Specifically, a sentence expresses a minimal proposition only if its predicate is context-insensitive.

[7] See Sperber and Wilson (1995, pp. 189-190) and Recanati (2004, p. 8), among others.

[8] See Travis (2008a) and Searle (1978) for a radical pragmatist attack on minimal propositions. Related views are held or explored, among others, by Bezuidenhout (2002), Carston (2002) and Recanati (2004).

tent[9] . One of the main arguments they have used consists in a series of examples, as Travis cases, in which two uses of a well-formed declarative sentence S have different truth-values—despite the referent of indexicals and the relevant state of affairs being fixed. Since the truth-value shifts, the two uses of the sentence must be true under different conditions. However, being tokens of the same type, they share their semantic properties. So the semantics of S doesn't seem to be determining the truth-conditions of an utterance of S. If the argument works, and assuming the class of sentences for which we can generate such examples is large enough, we would have reasons to think that minimal propositions are a chimera and that semantics is not in the business of delivering truth-evaluable content [10] .

Here I will focus on the second criticism[11] . In particular, I will focus on Travis cases and on how the argument based on them has been contested by philosophers who (often implicitly) endorse Semantic Propositionalism. These philosophers can be classified into two groups. Minimalists (as Sainsbury (2001) and (2008), Cappelen and Lepore (2005), and Borg (2004) and (2012)) typically claim that semantic content can be recovered without pragmatic interpretation and keep context-sensitivity to a minimum. By contrast, indexicalists[12] (Szabó (2001), Hansen (2011)) aim at securing truth-conditional semantics by positing context-sensitivity beyond the list of classical indexicals (following the model provided by indexicals). I will argue that the minismalist and indexicalist responses to the pragmatist challenge found in the literature are defective for various reasons. Hence, in view of Travis cases, we have no reasons to maintain Semantic Propositionalism. This does not mean that there is no such a thing as semantic content. Words and sentences have fixed linguistic meanings. We can keep the term 'semantic content'

[9] For a large class of sentences.

[10] Predelli (2005b) argues that Travis argument relies on a false premise, namely, that if two utterances of S, u and v, seem to have different truth-conditions, then an adequate system ought to assign to u and v different intensions. Against this, Predelli, an advocate of relativist semantics, argues that u and v might have the same intension but be evaluated at different points of evaluation. This, however, is of no help to the advocate of Semantic Propositionalism, for the output of Predelli's system (a function from clause-index pairs to points of evaluation) is not, of itself, truth-evaluable. on an occasion of use, it remains to be determined which is the point of evaluation. In a sense, u and v do have different truth-conditions: u is true if and only if S is true at the point of evaluation determined by the context of u, whereas v is true if and only if S is true at the point of evaluation determined by v. I will not discuss relativism here. For the sake of clarity we can distinguish intensions from truth-conditions.

[11] Note that the second criticism is stronger than the first, for it is still compatible with Semantic Propositionalism that semantic contents don't play any relevant role in communication.

[12] Together with indexicalism I will discuss other related views (sec. 4).

to refer to it. It just means that semantic content is not truth-evaluable content.

The plan is the following. In sec. 2 I will present some Travis cases for colour predicates and the argument they motivate. In sec. 3 I will present and reject some minimalist arguments intended to deactivate the challenge by denying that Travis cases involve a truth-value shift. In sec. 4 I will assess some theories that take colour predicates to be context-sensitive. I will conclude that neither of these strategies works.

2.2 Semantic Propositionalism vis-à-vis Travis cases

Travis cases have been generated for a variety of words. However, the most discussed examples involve colour terms. Because of this, I will focus on those. Here are some of them:

'The leaves are green':

> Pia's Japanese maple has russet leaves; she paints them green. Addressing her neighbor, a photographer looking for a green subject, she says, apparently truly:
>
> (1) the leaves are green.
>
> Imagine now that Pia's botanist friend is interested in green leaves for her dissertation and that, in reply, Pia utters (1) again. This time, her utterance seems intuitively false. (Predelli (2005b), modifying Travis' example)

'The painting is purple':

> Max buys a red painting and hangs it on a blacklight room. When he turns on the blacklight, he says: 'The painting is purple'. This seems true. When he turns the light off, he says: 'The painting is not purple'. If he is talking about how the painting looks in normal illumination conditions, then it is true. If he talks about things changing colour, then it is false. (Travis, quoted in (Recanati, 2010))

'Swatch 27 is green':

> Suppose we are scientists studying color-blindness, and each week is devoted to studying a particular subject. During the first week we are studying subject #1. He is shown various swatches, and when one looks green to him he responds affirmatively. During this first week, we adopt our use of 'green' to match subject #1's perceptual judgments; so, during the first week, a swatch counts as green if and only if it is green for subject #1. The following week we are studying a different subject, subject #2, and during this second week we adopt our use of 'green' to match her perceptual judgments. And of course it might happen that a particular swatch, swatch 27 say, is green for subject #1, but not for subject #2. And thus uses of 'Swatch 27 is green' might count as true during the first week, but not during the second week, and the intuitive dimension of 'contextual incompleteness' concerns the judge. (Clapp, 2012a, p. 79)

And a last example, not involving colour terms (inspired on an example by Searle):

'The cat is on the mat':

> Roberto arrives home after work and, wanting to play for a while with his cat, asks François: 'Where's the cat?' François answers: '[The cat is] on the mat'. As it happens, the cat is sitting on a piece of paper on the mat. The piece of paper, however, is irrelevant for the purposes of locating the cat, and so the utterance seems true. Now, imagine the next scenario. Roberto and François have recently got all the mats in the house cleaned. On rainy days, their cat often goes outside and comes to the house with mud on his legs. They don't want the mat the cat usually sits on to get dirty, so Roberto has covered it with a piece of paper. François, worried about the mud, asks: 'Where is the cat?' François replies: 'Don't worry, it's not on the mat, I covered it with a paper, the cat is sitting on the paper'.

The argument against Semantic Propositionalism can be reconstructed as follows (assuming S is free of indexicals[13] and assuming the relevant state of the world remains fixed):

> P1. If Semantic Propositionalism is true, then tokens of S will all express the same truth-evaluable content.

[13] The sentences involved in Travis are not free of indexicals, but their value is kept fixed. I am not yet considering the option that predicates contain hidden context-sensitive variables.

P2. In a Travis case, tokens of S have different truth-values.

P3. If two tokens of a same sentence, assessed against the same state of affairs, have different truth-values, then they must have different truth-conditions (express different truth-evaluable contents).

P4. Given that the two tokens of S are being assessed against the same state of affairs, it follows that not all tokens of S express the same truth-evaluable content.

C. Semantic Propositionalism is not true.

Note that with the examples, the advocate of truth-conditional pragmatics[14] is not merely claiming that there can be shifts of truth-value of the discussed kind, but also that in absence of a context of use the sentence lacks a truth-value. Since uses of it can be true and can be false (again, keeping the relevant state of affairs fixed), the type is compatible with truth and falsity, and so is neither true nor false.

Those with a Gricean spirit might want to dismiss the argument by claiming that it is based on an equivocation between an utterance being true/false and it being appropriate/inappropriate. For example, it certainly is inappropriate to tell a botanist that you have green leaves when all you have are painted green leaves. But this, the thought goes, doesn't mean that the utterance is false. It might be simply misleading (given the context).

However, this idea, of itself, is not yet a counterargument. Some work needs to be done in explaining how to apply the Gricean insight to Travis cases and isolate what is strictly speaking said in both scenarios, i.e., one needs to justify why the utterance should be considered true instead of false (or vice versa). It seems that the meaning of the predicate is of no help here[15] . 'Green', for example, refers to a colour; 'is green' is a means for describing things as having that colour. But the meaning of 'green' doesn't seem to decide whether painted things count as having that colour. However, if Semantic Propositionalism is true, the semantic features of the predicate should be doing the work.

The challenge, in short, is that semantics doesn't seem to be fit to do the work Semantic Propositionalism assigns to it. Advocates of minimal propositions need either a good argument for choosing one truth-value over the other, or, in absence of such an argument, a good explanation why we are unable to choose—with semantics, by

[14] I include here both contextualists and occasionalists, in Borg's terminology.

[15] See Travis (Travis, 1997, p. 90).

contrast, being able to determine a unique truth-value. Hence, those who endorse Semantic Propositionalism typically reject premise 2. In section 3 I assess four minimalist strategies to this end. Alternatively, advocates of Semantic Propositionalism can argue that, whereas the semantics of the sentence remains fix, its content varies across contexts. Semantic content is content determined by semantics and this, following Kaplan's analysis on indexicals, can vary. I will explore the prospects of this way of securing Semantic Propositionalism in section 4.

2.3 Rejecting P2

In this section I will review four arguments that attempt to rebut the truth-value shift intuition. If they were successful in rejecting premise 2 in the argument, then they would thereby secure minimal propositions. However, I will argue that they are defective for various reasons.

2.3.1 Rejecting P2: Unspecific meanings

How can one make sense of the claim that Travis cases don't involve a truth-value shift? According to Sainsbury (2001, 2008), the meaning of 'is green', etc. is unspecific. The satisfaction conditions of this expression are unspecific relative to different ways of being green. Thus, the two utterances of 'The leaves are green' are true. 'The leaves are green' can be made true by virtue of the leaves being painted green, or by the leaves being naturally green by the effect of chlorophyll, etc. These are simply more specific ways of being green. But there being different ways of being green is not equivalent, according to Sainsbury, to there being different satisfaction-conditions for 'is green'.

The general idea Sainsbury is putting forward is that some predicates can be satisfied in different ways. For example, one can dance in different ways. Paul dancing in some particular way makes an utterance of the sentence 'Paul is dancing' true, regardless of the specific way in which he is dancing and of the possibility of dancing in other ways.

Borg, another proponent of this view, calls this 'liberal truth-conditions'—'liberal' because they 'admit of satisfaction by a range of more specific states of affairs' (Borg, 2004, p. 230).

The problem with this proposal is that it relies on there being context-insensitive ways of classifying things as green. But this is precisely what Travis cases call into question and what needs to be argued. Sainsbury's argument depends on the following two assumptions. First, if an object o is F in some way, then 'o is F' is literally true (bracketing time). Second, whether o is F in some way or not can be decided in a context-insensitive manner. This second assumption is problematic. We can ask: Is being painted green a way of being green? Is looking green in non-standard conditions a way of being green? Is being perceived as green in such-and-such illumination conditions a way of being green? What about being perceived by someone wearing green glasses? Travis cases suggest that there are no purpose-independent answers to these questions—change what is at stake at the conversation and the answer varies. Because of this, assuming that being painted green is (purpose-independently) a way of being 'green' is question begging—this is precisely the claim that would need to be argued for.

Moreover, when one moves from the green leaves example to the purple paint example or the swatch 27 example Sainsbury's strategy sounds less convincing. If one accepts that being naturally green and being painted green are more specific ways of being green, should one also accept that being perceived as green in such-and-such illumination conditions, or when wearing certain glasses, is a way of being green? After all, we can create Travis cases involving these circumstances. In general, many objects that are not perceived as green in normal illumination conditions (daylight, no glasses, let's say) would appear as green if we changed the light, or wore special glasses. Is that a way of being green? Or, many objects have green paint in what, for most purposes, are irrelevant parts—think, for example, of a book with one tiny word printed in green. Is the sentence 'This book is green', when 'this book' refers to the described book, literally and purpose-independently true? There is no clear answer to this question, unless we embed the sentence in a conversation where a topic of discussion has been fixed. However, extending Sainsbury's answer to these cases, one should conclude that all these ways of being are simply more specific ways of being green.

As a consequence, the advocate of unspecific meanings faces a difficult position. On the one hand, he can insist that all the sentences in the previous examples are literally true, i.e., he can extend Sainsbury's analysis to other examples and take these examples to involve different specific ways of being green. But this, as I have argued, has the un-

palatable consequence that it would automatically make many sentences of the form 'x is green' true, for many objects are perceived as green in some condition or other, by some observer or other, or have some very tiny green parts. In particular, it would make true sentences that are best seen as false in most ordinary situations. For example, the sentence 'This sheet is green' (talking about a normal white sheet), would be true if being seen green by someone wearing green glasses is a way of being green. However, in most occasions in which this sentence is used to describe the sheet it seems like a blatant falsity.

On the other hand, he can insist that there is a criterion distinguishing literally true and literally false colour-sentences in Travis cases, i.e., he can try to find a robust criterion establishing which ways of being are ways of being green. However, the prospects of finding such a criterion seem rather dim. We know, for specific purposes, or in specific circumstances, which ways of being are ways of being green. But besides this, it is not at all clear that we know, independently of any purpose, which ways of being are *strictly speaking* ways of being green—unless we stipulate it or explicate the concept (see sec. 3.3).

2.3.2 Undermining the justification for P2: Privileged intuitions

The justification for premise 2 (the truth-value shift) is given by appealing to the intuitions triggered by Travis cases. Consequently, one can reject the argument by undermining the assumption that our intuitions support the existence of truth-value shifts. Borg suggests that intuitions do not clearly justify premise 2, for some intuitions are in line with the prediction of insensitive or minimal semantics (a semantics that would respect Semantic Propositionalism). She appeals to two different kinds of intuitions:

> [I]f we are really interested in finding out what the words and sentences we use literally mean, then we have no need to look further than the meaning which can be recovered via sensitivity to formal features alone. Recall also that there is evidence that speakers can grasp this very liberal, literal meaning when they want to: children and philosophers, I often find, have a very acute sense of what they have literally committed themselves to by a given utterance and this fits entirely with the liberal truth-conditions specified above. (Borg, 2004, p. 243)

> I think it is not right to portray minimalism (as it sometimes is) as running counter to all pre-theoretical intuitions about utterance-level content. Given the right context (i.e. one where subjects are asked to reflect on 'literal' or 'strict' meaning) ordinary interlocutors can and do grasp exactly the kinds of contents minimalism predicts. (Borg, 2012, p. 14)

According to Borg, we could focus on the intuitions of particular groups: children and philosophers. Or, we could ask ordinary speakers to reflect on literal meaning. Now, would this secure Semantic Propositionalism? It is important to note that these two tests could deliver inconsistent verdicts, and that the results could also be in contradiction with Borg's (and Sainsbury's) liberal truth-conditions. For example, it could be the case that children and philosophers took painted green objects to satisfy 'is green' (thus to coincide with Borg's liberal truth-conditions) but that ordinary speakers, when asked to evaluate the literal or strict meaning of 'The leaves are green' would answer that the leaves are not strictly speaking green, but only painted green[16] . Let's bracket this issue and focus on the proposals.

I am not sure that philosophers' intuitions will be of much help for the advocates of minimal propositions. After all, many philosophers of language, impressed by Travis cases and other similar examples, have put forward different versions of contextualism (in a broad sense) and relativism. What about children? Children are precisely those who are being trained in the use of language. Their linguistic knowledge is incomplete. Why should we trust their intuitions more than adults' intuitions?

The third option (to ask speakers to reflect on literal meaning) seems like a more reasonable proposal. But here we have two problems. First (I have already mentioned this problem), it is not at all clear that this test will be in line with Borg's version of Semantic Propositionalism (perhaps painted green things are not, for ordinary speakers, literally, or strictly speaking, green). But more importantly, there is no reason why we should focus on explicit judgements, in very specific contexts (including the words 'strict' and 'literal') instead of focusing on implicit judgement. Borg is assuming that ordinary speakers can reflect upon literal meaning when asked—but ordinary speakers are people who master the use of language, not necessarily people who will achieve, during a test, relevant discoveries about meaning. To begin with, it is not at all clear that 'literal meaning' is an ordinary term—it seems rather like a technical term. If one wanted to put to the test the linguistic knowledge of ordinary speakers, why not put

[16] This is the reply I sometimes get when I present the green leaves case. Some people complain that the leaves are not green, because they are *painted* green.

truth-conditional semantics to the test by observing which utterances ordinary speakers hold true, as Davidson suggested?

Davidson (1973) puts two constraints on a theory of interpretation for a language L. First, the theory must be finite but deliver truth-conditions for any sentence in L, i.e., for an infinite set of sentences. Second, the theory must be empirically correct, i.e., the truth-conditions it attributes to a sentence S must coincide with the truth-conditions speakers of L attribute to S. How can we check that a semantic theory is correct in this second sense? For Davidson, the evidence justifying the formal theory can be obtained by observing the behaviour of speakers. A starting point for the justification of the formal theory is provided by the attitude of holding a sentence true. The outputs of the theory must coincide with the sentences speakers hold true (assuming there is no reason to think they are mistaken). Following this idea, a good way to check whether advocates of Semantic Propositionalism as Borg are right would consist in checking whether the liberal truth-conditions they attribute are in line with the sentences speakers hold true (in context). The problem for Borg is that it is likely that we will find that there are variations in the sentences speakers hold true across contexts (keeping the relevant state of affairs fixed). A quite ordinary example: on some occasions speakers seem to hold true the description 'is blond' as applied to someone who has dyed his hair ('Jess is the blond guy'), whereas on others they seem to hold not true the same description as applied to a similar (or even the same) person ('He's not blond, he dyed hid hair').

2.3.3 Rejecting the role of intuitions (I): Experts

Cappelen and Lepore suggest a way of discarding the truth-value shift intuition (Cappelen and Lepore, 2005, p. 164). According to them, whereas it might be part of the job of the semanticist to tell us that the proposition semantically expressed by 'a is green' is that a is green, and that an utterance of 'a is green' is true if and only if a is green, it is to the metaphysician to tell us what it is for a to be green. So it is to the metaphysician to tell us what the truth-value of Pia's utterance is.

The problem with this answer is that, given that semantics is neither in the business of telling us what it is for an object to be green, nor in the business of telling us which sentences are true, how can it be in the business of telling us that two utterances of a

sentence must share a truth-value? The metaphysician could very well take properties to be occasion-sensitive or, alternatively, he could take the predicate 'is green' to express different properties in different occasions.

Let me explore a related line of reasoning. Perhaps the idea is that, if everything goes well, and after philosophical inquiry, philosophers will tell us what the literal satisfaction conditions of 'is green' are (whether this expression is literally true of painted leaves). Is this line of reasoning promising? I think it won't work as a defence of minimal propositions. If the advocate of minimal propositions is interested in finding out more about the property ordinary speakers attribute when they use the expression 'is green', he will have to pay attention to ordinary uses. However, again, chances are that he will have to give up minimal propositions. Alternatively, he can opt for explicating the term in Carnap's sense[17] . Explicating a concept consists in taking an ordinary concept, more or less inexact (in this case, one whose satisfaction-conditions seem to shift across contexts) and transforming it in a new, more precise concept (in this case, one whose satisfaction-conditions don't shift across contexts). Thus, the philosopher who opts for explicating 'green' would be coining a new concept–the philosophical 'is green'. In this second case, he would not be illuminating the discussion about our predicate 'is green', but about a new, more philosophical one. But what we are interested here in is natural language.

2.3.4 Rejecting the role of intuitions (II): epistemicism

It might be the case, as for example Borg (2012) holds, that some of the examples used against Semantic Propositionalism arise from misplaced intuitions: our intuitions might be telling us something about some property of the utterance or the sentence (maybe about a presupposition, about its triviality…) but not about its literal truth-value. However, finding out where exactly our intuitions have gone wrong might be a difficult task for those who endorse Semantic Propositionalism. Nonetheless, there is another

[17] 'The task of making more exact a vague or not quite exact concept used in everyday life or in an earlier stage of scientific or logical development, or rather of replacing it by a newly constructed, more exact concept, belongs among the most important tasks of logical analysis and logical construction. We call this the task of explicating, or of giving an *explication* for, the earlier concept' (Carnap, 1947, pp. 8-9)

strategy they might follow. Semantic Propositionalism is a principle about semantics, or about meaning, not about us and what we are able to do. In this sense, it is about the metaphysics of meaning. It states that meaning determines a truth-evaluable content, not that we are able to ascertain, given the actual world, what the literal truth-value of the sentence is. It could be argued that our inability to grasp the literal truth-value of a given utterance reveals nothing about the metaphysics of meaning.

One can read some parts of Borg (2012) as an implementation of this line of reasoning. According to Borg, the condition that semantic contents seem to fail to meet vis-à-vis the existence of Travis cases is that of, in Borg's terminology, sorting worlds. Travis cases call into question the assumption that semantically determined contents are truth-evaluable. They do so by showing that we, normal speakers, interpret utterances of the same sentence in different ways in different occasions. Because of this, if we are asked whether Pia's leaves are, strictly speaking, green, abstracting from the occasion of use, we are very likely to be puzzled. Now, it could be objected that one should distinguish semantic contents sorting worlds from us knowing how is the actual world sorted. Perhaps we should not use our intuitions about truth-values in discussing the metaphysics of meaning.

Borg presents a view that I will call, because of its similarities with Williamson's view on vagueness, Borg's epistemicism:

> [T]he current stance holds that there is a fact of the matter about whether or not, say, the sentence 'That cat is on that mat', relative to a context c, is true or false, but the approach also allows that this fact of the matter may be unknown to an agent who both fully comprehends the literal meaning of the sentence and is in a position to assess the relevant state of affairs in the world. (Borg, 2012, p. 109)

> So, according to the minimalist, a sentence like 'the cat is on the mat' expresses a complete proposition on the basis of lexico-syntactic constituents alone and this proposition is what a subject needs to grasp to understand the literal meaning of the sentence, yet grasp of such a proposition may not, in and of itself, allow a subject to immediately tell for all possible worlds whether they are worlds which satisfy or fail to satisfy the propositional content (even though there is a fact of the matter to be discovered here). (Borg, 2012, pp. 109-110)

'The cat is on the mat' makes reference to a famous example used by Searle in his argument that truth-conditions are relative to a background of implicit, perhaps non-representational, assumptions. Searle's main idea is that the expression 'is on the mat'

only has application conditions assuming there is gravitational force—of cats and mats floating in outer space we wouldn't know what to say. To be sure, we could have a practice in which this expression would apply in absence of gravitational force in case, for example, the cat and the mat were in touch. Astronauts could use the expression in this way. If so, we could use these two understandings and create a Travis case for the sentence. Or we can imagine a cat on a paper on a mat and create a different Travis case (see the example presented in sec. 2). Or we can imagine a cat with three legs on the mat and one on the floor. Borg focuses on this last example, but the threat to minimalism is best seen when we realise that the problem that the advocate of truth-conditional pragmatics is putting forward has little to do with vagueness[18] .

Borg claims that a subject who knows the literal meaning of the sentence (plus the referent of indexicals, let's say)—i.e., a subject who grasps the minimal proposition—and who is in a position to assess the object or state of affairs the minimal proposition is about might be incapable of telling whether the utterance is true or false. Pia, and all of us, might be a competent English speaker (understand the sentence 'The leaves are green') and be looking at the relevant leaves, but fail to know whether the sentence is true or false of precisely those leaves. According to Borg, this is compatible with minimal propositions determining (in a metaphysic sense) a truth-value. To put it in other words: minimal propositions impose a condition of the world, but whether the world satisfies that condition might be epistemically closed to well-positioned[19] competent speakers.

In this picture, the fact of the matter determining the truth-value of the minimal proposition is, allegedly, only contingently epistemically closed to us. We could come to know it. As Borg puts it, we could discover what the fact of the matter is 'by further investigating the meaning of the expressions involved (e.g. setting out to discover whether the English word 'on' denotes a property true of cats and mats when cats have only three out of four legs resting on the mat, where this investigation would involve probing, among other things, ordinary speakers usage of the term).' (Borg, 2012, p. 109) This is important because, without this, Borg's epistemicism would imply a rejection of

[18] Another example Borg often considers is 'The apple is red'. 'is red' sometimes describes the colour of the skin, sometimes its interior. Moreover, Borg consider that indexicalism and contextualism are unstable positions, and so she takes it that it is either occasionalism (i.e., Travis' view) or minimalism.

[19] By 'well-positioned' I mean subjects with access to the relevant state of affairs, paying attention and with fully operative cognitive capacities. Note that we are talking about green leaves, red apples and cats on mats and that the source of our inability to assess the proposition is related to meaning—not to the metaphysics of colour, etc.

the interpretability of natural language.

Taking language to be interpretable means taking the satisfaction-conditions of words to be something knowable. There probably are domains in which truth is unknowable. But meaning doesn't seem to be one of those. Why? Because language is a social affair. The meaning a certain word has depends on how we happen to use it. Meaning facts are not out there, independently of our use of language. In the end, as seems to be implicit in the previous quote, whether the expression 'on the mat' is true of a certain cat depends on how competent speakers use that expression. And this is something we can observe

Before going to the reasons why the kind of situation Borg is describing is problematic, let me make two remarks. First, we don't have here new reasons to take Semantic Propositionalism to be true. Semantic Propositionalism is taken as the default position. This stance tries to deactivate some criticism by claiming that our not knowing whether 'The cat is on the mat' is true or false of certain states of affairs is compatible with Semantic Propositionalism. So it works on the assumption that we don't know whether that sentence, in a certain context, is true or false, despite the fact that when reading a Travis case speakers usually have the impression that one utterance is true and the other false, not that they don't know whether the utterances are true or false. Claiming that we don't know the literal truth-value means conceding no weight to intuitions triggered by Travis cases. Borg's epistemicism works on the assumption that our having shifting intuitions shows that we are blind to the literal truth-value.

Second, it is important to note that we very often find ourselves interacting with objects and in situations for which Travis-style cases could be created. Nearly no object is green in every respect. Few cats sit perfectly on the mat. We often sit on cushions on chairs (is that a way of sitting *on* a chair?). As a result, if Semantic Propositionalism is true, the situation Borg is describing must be a quite ordinary one. It is not that there are a few strange situations in which our linguistic competence doesn't enable us to see the literal truth-value (despite the fact that there is a literal truth-value). Rather, for a considerable number of ordinary objects and a considerable number of ordinary predicates we are unable to see if the object literally is as the predicate describes it. Let us call this class of objects that are not F (green, etc.) in every respect 'O'.

Is Borg's epistemicism a tenable position? Two considerations tip the scale towards a negative answer[20] . The first problem with epistemicism is that it makes our capacity to

[20]Recanati (2004) provides a response to Borg's (2004) claim that the contextualist is guilty of confusing truth-conditions and verification conditions, for the contextualist (according to Borg) confuses

speak literally and convey information mysterious. Let us assume Borg's epistemicism and imagine the following situation. I have a shirt that belongs to O for the predicate 'is yellow', because it has a yellow label but is not yellow all over. Although I might sometimes describe it as 'yellow', in most occasions, or for most conversations, I take it not to be 'yellow'. According to Borg's epistemicism, and given the fact that if a Travis cases was to be presented, I would have the truth-value shift intuition, I am ignorant about whether the description 'is yellow' is literally true or false of my shirt. Despite this, when a friend asks if I have a yellow shirt that she can borrow for a party because, she tells me, she thinks a yellow shirt would fit nicely with her black trousers, I simply say 'I don't have a yellow shirt' and manage to convey some useful information. However, given that I don't know how the actual world is sorted by the sentence 'My shirt is yellow', and related sentences, it is a mystery how I manage to say 'I don't have any yellow shirt' and thereby convey some useful information, i.e., it is mysterious how I and my friend manage to coordinate ourselves in reducing the set of possibilities (that I have a yellow shirt, that I don't) to how things really are. In my case, it is mysterious how I can find a sentence that operates the reduction of possibilities to the actual one, given that I am unable to assess how sentences referring to my shirt and describing it as 'yellow' sort the actual world. Strictly speaking, if epistemicism is true and it can be put to work in explaining away Travis cases I don't know how is my utterance literally sorting the actual world.

Perhaps the explanation is that my utterance expresses, besides the minimal proposition whose truth-value I don't know, a pragmatically enriched one—something equivalent to 'I don't have any shirt whose predominant colour is yellow'. However, if one explains the case by saying that the proposition we are aware of expressing is one pragmatically enriched, then one undermines our capacity to speak literally, that is, to use the contents that, allegedly, are semantically expressed by the sentences we use. Securing

knowing the truth-conditions of a sentence (given by disquotational schema) and being able to verify whether a given state of affairs satisfy them. Recanati argues (rightly in my view) that the 'central idea of truth-conditional semantics (as opposed to mere 'translational semantics') is the idea that, via truth, we connect words and the world. If we know the truth-conditions of a sentence, we know which state of affairs must hold for the sentence to be true. T-sentences display knowledge of truth-conditions in that sense only if the right-hand side of the biconditional is used' (Recanati, 2004, pp. 92-93). This suggests that the situation Borg (2012) describes is not tenable: if a speaker knows truth-conditions, then he knows 'which state of affairs must hold for the sentence to be true' and, in particular, whether the actual state of affairs is such a state of affairs (assuming he is a position to asses it). I will not pursue this line of reasoning further. Instead, I will argue that, even granting that the situation described is intelligible, it is problematic.

minimal propositions amounts to preventing them from having a use[21] .

The second problem has to do with interpretability. Borg's view is that one could make more work and find out what the literal truth-value of the sentence is. We could do some work and find out whether the predicate 'is green' is one that is true of Pia's leaves, whether a cat sat on piece of paper on a mat makes 'The cat is on the mat' true, etc. The way to do it is to probe ordinary speakers' usage of the term. The problem now is that Borg's position makes ordinary speakers' usage unreliable. We have been told that knowing a minimal proposition is compatible with not knowing its truth-value relative to a state of affairs we are in a position to assess, and that this explains our (allegedly) confusion in assessing Travis cases. But then, why should we trust ordinary speakers? If we, competent speakers, don't know how the actual world is sorted in the examples discussed, and have shifting intuitions, why should we think other speakers do? Are their intuitions about speech act content neither misplaced nor confused? We haven't been given any reasons to rule out that they, just like us, are ignorant about literal truth-values.

2.3.5 Conclusions

I have argued that the arguments offered in support of minimal propositions are problematic. First, assuming that there are purpose-independent ways of being green is question begging, for this is the claim that Travis cases target. Second, relying on the intuitions of a specific group of people (children, or those who are asked to reflect on literal meaning) is unwarranted. Third, explicating the concept would not be fair to the goal at stake, i.e., finding out whether Semantic Propositional is true of natural language. Fourth, epistemicism makes our capacity to convey information mysterious, or relies on pragmatic adjustment to the extent of casting doubts that minimal propositions are playing any role. Moreover, it is at odds with interpretability.

As a result, advocates of minimal propositions find themselves in an unstable position. In order to secure minimal propositions, they need minimal propositions to be immune to common reactions to Travis cases. So they need to reject the intuitions

[21] A further problem would be to justify why epistemicism doesn't hold for pragmatically enriched propositions.

triggered by these cases, either by convincing us that both utterances in a Travis case are literally true (or false) or by ruling out the role of our intuitions. However, ordinary speakers' behaviour and intuitions are the best (if not the only) way of finding out what the meanings of the words in a language are and what their extensions are on an occasion of use. So they are led to a dilemma. Either they prevent ordinary speakers' behaviour and intuitions from having any role, which might preclude the interpretability of natural language; or they allow ordinary speakers' behaviour and intuitions to be used in deciding which is the literal truth-value of utterances in Travis cases, in which case they run the risk of being forced to renounce to minimal propositions.

2.4 Rejecting the non-indexicality assumption

So far I have assumed that advocates of Semantic Propositionalism take the predicates involved in Travis cases to be context-insensitive. However, some semanticists have argued that colour predicates contain hidden context-sensitive variables, or that they are ambiguous, or that they involve some other form of context-dependence. If this is so, there is a chance that Semantic Propositionalism can be made compatible with the variation we find in Travis cases, at least when it comes to colour predicates. In this section I examine these proposals.

2.4.1 Semantic Propositionalism, ambiguity and context-sensitivity

How would Semantic Propositionalism work for ambiguous and context-sensitive sentences? Let us recall the principle of Semantic Propositionalism:

> *Semantic Propositionalism:* The semantic content of a well-formed declarative sentence S determines a truth-evaluable content (a proposition, a truth-condition).

The principle, as it is stated here, is about sentences. If 'is green' in 'The leaves are green' is ambiguous, then we can take it that this inscription can express two different sentences. The principle would apply to these disambiguated sentences.

In the case of context-sensitivity, one should understand 'semantic content' as 'character':

> *Semantic Propositionalism (context-sensitivity):* The character of a well-formed declarative sentence S determines a truth-evaluable content (a proposition, a truth-condition) at a context of use.

If a predicate is conceived as exhibiting context-sensitivity, its semantics will need to be revised. One might, for example, posit some hidden (aphonic) context-sensitive variable in the logical form of the predicate, following Stanley's indexicalism (Stanley, 2000). This procedure implies a syntactic revision with semantic consequences. Alternatively, one could argue that the semantics of an expression involves some form of context-sensitivity, because, let's say, it is part of its meaning that some contextual contribution is needed. Although this second option does not seem to imply a revision of the syntax, it is not clearly distinguishable from indexicalism. What is relevant here is that these options work on the assumption that we can identify the sources of the possible variations (the dimensions that can provoke a change in satisfaction-conditions) and translate them into a set of parameters—that we will include either in the syntax of the predicate or otherwise in its semantics. The route to the identification of these parameters often goes via Travis cases[22] .

In order to be a proper defence of Semantic Propositionalism, an account having recourse to context-sensitivity must fulfil two conditions. First, it must provide a set of necessary and sufficient parameters[23] . The motivation for this condition is not difficult to see. If the set of parameters that are relevant to account for the satisfaction conditions of a use of a predicate were to be given by the context, they would not be semantic after all. Semantics would be compatible with variation in truth-conditions. By contrast, if semantics is to determine the satisfaction-conditions of a predicate on an occasion of use, it must encode all the possible dimensions of variation. Any possible Travis

[22] The fact that the reasons for revising the syntax are merely pragmatic is already a problem for some advocates of Semantic Propositionalism, as Borg (Borg, 2012). Stanley (2000) provides a syntactic argument in support of positing certain hidden variables. However, indexicalists dealing with Travis cases often lack independent arguments supporting the claim that colour predicates contain hidden variables.

[23] See Davies (2014) for the use of a similar desideratum against some indexicalist proposals.

case must be generated by shifting the value of a parameter encoded in the syntax, or otherwise in the semantics, of the predicate.

Not all the accounts to be described in the next section deal with parameters. The spirit of the condition, however, is that the semantic analysis must be able to handle all sources of indeterminacy, that is, it must account for all Travis cases. This means that the semantics of the expression must fix all the possible satisfaction-conditions the expression can have[24] .

Second, it must be an instance of what I will call semantic determination. It is important to note that the principle of Semantic Propositionalism is about semantics determining a truth-evaluable content. Applying the character/content distinction does not automatically get us a framework in which semantics is sufficient for delivering truth-evaluable content. As Recanati (2017) argues, the idea of indexicals having characters suggests that the semantics of these expressions involve some kind of procedure whereby a speaker could access the truth-conditional content of an utterance[25] . Linguistic competence is supposed to be sufficient to grasp the referent of an indexical on an occasion of use. However, many indexicals don't work like this. Let us compare the sentences 'I am in Barcelona' and 'She is in Barcelona'. In the first, the semantics of 'I' establishes that the referent of this expression is the person uttering the sentence. Given a context (in the intuitive, not the technical sense), the referent can be said to be automatically determined. That piece of information ('utterer of this sentence') is sufficient for the interpreter to pick up an individual. By contrast, it is plausible to take the semantics of ''she' to be something like 'salient female'. Given a context of use, this semantic information is insufficient to pick up a referent. The interpreter needs to know in advance who is salient in the conversation in order to be able to interpret the utterance. Whereas the semantics of 'I am in Barcelona' can be said to respect Semantic Propositionalism there are reasonable doubts that 'She is in Barcelona' does. The semantics of 'she' constrains the referent of this expression on an occasion of use, but it is not at all clear that it determines it.

Despite this intuitive difference between 'I' and 'she', 'context' is often a technical notion. A context can be conceived as a sequence of objects. When we conceive contexts in this or related ways, the referent of a demonstrative is fed into the context. Once this is done, the character of the sentence can be said to determine the content at that context.

[24] Alternatively, one could claim that some cases not accounted for by his theory are not really cases of truth-value shifts and combine indexicalism with some version of minimalism.

[25] Perry (2001) distinguishes automatic from discretionary indexicals.

But one should question whether this kind of mechanism is worth calling 'semantic determination'. It does not capture a procedure a speaker could use in order to access the referent of an indexical. The character is not telling him which object the indexical refers to. It works on the assumption that he already knows that. Hence, if a theory based on context-sensitivity is not a genuine case of semantic determination, then it is not really semantics that is doing the work and the theory will not count as a defence of Semantic Propositionalism.

As a result, a semantic analysis of an expression E will count as an respecting Semantic Propositionalism only in case the two following conditions are respected:

> *Condition I:* The semantics of E fixes the ways in which E is indeterminate (for example, by providing a set of necessary and sufficient parameters).

Whether an account satisfies this condition can be checked by applying the account to a range of Travis cases. The idea would be the following. Let us suppose that we have a theory T identifying a set of parameters as part of the semantics of expression E. If T is to be a defence of Semantic Propositionalism, then that set of parameters must exhaust the sources of indeterminacy, and with it, the possibility of generating Travis cases. Now, imagine that we keep the value of those parameters fixed and, nonetheless, are able to create a further Travis case. That would mean that the semantics T attributes to E are compatible with variation of truth-conditions. So T is of no use as a defence of Semantic Propositionalism[26] [27] .

> *Condition II:* Semantics determines, and not merely constrains, the referent or extension of E on an occasion of use.

We can check that an account fulfils this condition either by taking recourse to the kind of knowledge a speaker would need to use in order to interpret an utterance including the expression, or by assessing whether the information encoded in the semantics of E is enough to select its referent or extension.

I will only discuss Travis cases involving colour predicates, for these are the examples often considered by those who intend to meet Travis' challenge by positing some

[26] Travis often uses this kind of argument ((Travis, 1981, p. 53), (Travis, 2000, pp. 35-36)).

[27] Strictly speaking, there could always be further Travis cases we have not yet imagined that are not explained by the theory, but if we find ourselves unable to provide such a case, this can be seen as a reason to take the account to satisfy this condition.

form of context-dependence. However, one might ask whether indexicalism, or similar views, could be extended to other Travis cases. I think that approaches based on context-sensitivity don't seem very appealing when it comes to other words. Thus, even assuming the indexicalist succeeds in his analysis of colour predicates, the advocate of Semantic Propositionalism would still need to do much work. As Clapp puts it 'As we stroll through the transportation museum, the extension of our uses of 'airplane' will be negotiated and accommodated, and thereby sharpened, but I trust that nobody maintains that 'airplane' is an indexical.' (Clapp, 2012a, p. 86).

The theories I will discuss are instances of what is sometimes called indexical contextualism[28] . Indexical contextualists aim at explaining truth-value shifts by positing different propositions. By contrast, according to relativism or nonindexical contextualism, the proposition expressed in the two occasions described in a Travis case remains constant. The shift is due to it being evaluated relative to parameters that get different values. Because of linguistic reasons, one might find non-indexical contextualism preferable (or the other way round)[29] . I will not address this issue here. If a relativist was to posit the same parameters as the indexicalists I will discuss, my criticisms could be applied to this proposal as well [30] .

[28] See MacFarlane (2009). MacFarlane introduces the distinction and advocates for nonindexical contextualism.

[29] Suppose that we have reasons to think that the predicate 'is hexagonal' involves a parameter standing for standards of precision. A proponent of indexical contextualism would take the proposition expressed by the sentence 'France is hexagonal' to change from context to context. A proponent of non-indexical contextualism could argue as follows. Imagine this conversation:
(A) Italy is boot shaped.
(B) France is hexagonal.
(C) Well, strictly speaking, that is not true.
Here, what (C) says refers to the proposition expressed by (B) and is true. But if indexical contextualism is true, it cannot be the case, for the proposition expressed by (B) is something like 'France is roughly hexagonal', and this proposition, strictly speaking, is true. So what (C) says should be false. Because of this, the context-dependence of 'hexagonal' seems to be better explained by locating the parameter standing for standards of precision at the index. In this way, the proposition expressed by different utterances of 'France is hexagonal' is kept fixed but its truth-value is allowed to vary. However, both proposals assume that positing a parameter standing for standards of precision is sufficient to account for the variations in extension of this predicate across uses (keeping the state of the world fixed in the relevant respects).

[30] As I have noted, Davies (2014) presents a related critique against Szabó, Kennedy and McNally, and Hansen. His argument is based on the claim that these approaches don't model the meaning of colour predicates as a character, buy merely as something that varies with context. It is difficult to assess whether Davies takes semantic determination to be an intrinsic feature of what he calls 'characters'.

2.4.2 Available approaches

In this section I will present six approaches that purport to assimilate Travis cases to the framework of truth-conditional semantics by appealing to context-sensitivity, broadly understood ((Szabó, 2001), (Vicente, 2012, 2015) (Giberman, 2016), (Kennedy and McNally, 2010), (Hansen, 2011), (Rothschild and Segal, 2009)). In the next section I will assess whether they respect conditions I and II.

Szabó (2001) presents an analysis of colour predicates according to which the logical form of these expressions contains hidden context-sensitive variables. The proposal is based on the idea that 'different dimensions of incompleteness correspond to different variables' (Szabó, 2001, p. 136). Szabó finds two dimensions of incompleteness, comparison class and part. The motivation for the first dimension is that different objects are, for example, green in different respects—it is not the same to be green for a laptop than for a dress. Regarding the second, objects of the same class can be green in different parts. A dress can be green all over, in most of its surface, etc. This provides an easy explanation of the green leaves example. In the first utterance, the variable standing for the class of reference gets the value *leaves* and the variable standing for part gets the value *surface*, whereas in the second, this last variable gets as value *inside*.

Szabó's analysis locates the shifting dimension in the predicate. By contrast, Vicente (Vicente, 2012, 2015) locates it in the subject. This proposal is similar to Szabó in that it also takes parts to be relevant. Vicente is an advocate of conceptual semantics. According to this approach, lexical entries for nouns consist in complex conceptual representations including rich information about the kind of entity the noun designates[31] . Thus, the lexical entry corresponding to 'leaf' includes information about it being a physical object. The entry corresponding to 'physical object', in turn, (or the entry for 'leaf' itself) contains the information that physical objects can have surface properties that are not their original properties. Although it is not a form of indexicalism, for no hidden context-sensitive variables are postulated, it is similar to indexicalism in the following sense. According to Vicente, the meaning of the word 'leaf' includes the information that leaves are physical objects and that, as such, they have apparent surfaces and hidden, original surfaces. This is semantically equivalent to theories that establish that the word 'leaf' contains a variable standing for surface that can take two values, original surface or apparent surface.

In Vicente's account, a rule such as 'GREEN modifies LEAVES' leaves it indetermi-

[31] See Belleri (2014a) for a similar view.

nate which part of the conceptual structure is activated on an occasion of use. However, the meaning of 'leaves' determines the possible meanings of this rule, that is, more specific rules as 'GREEN modifies the apparent surface of LEAVES' or 'GREEN modifies the original surface of LEAVES'. Contextual factors decide which of them is selected on an occasion of use. Vicente concludes that 'when we have the relevant information, we are not only supplied with the meaning, but also with the *truth-conditional* meaning: which means that meaning does determine truth-conditions after all' (Vicente, 2012, p. 14)[32] .

Giberman (2016) also locates the shift in the subject. However, he takes mereological relations to be the key to accommodating Travis cases within truth-conditional semantics. On Giberman's analysis, the logical form of the subject of a sentence of the form 'a is F' is an ordered triple <object$_1$, pseudo-mereological operation, object$_2$ >. There are two possible pseudo-mereological operations, addition and subtraction. Their results are pseudo-mereological fusions of object$_1$ and object$_2$. Giberman takes contextual factors to determine the objects that occupy the slots, as well as the operation. Once this is settled, the pseudo-mereological operation operates on the objects and delivers the semantic value of the expression of which the property is predicated.

Giberman's explanation of the green leaves example goes as follows. The content of Pia's first utterance is:

> [<THE LEAVES at t , +, paint at t >, greenness]

whereas the content of the second utterance is:

> [<THE LEAVES at t , +, leaf parts spatiotemporally continuous with THE LEAVES at t that constitute their histories and extend back to a seed of a green-leafed plant species >, greenness]

[32]It can be a bit unfair to describe Vicente's analysis as a defence of Semantic Propositionalism for, in his analysis, the semantics of S doesn't deliver *one* truth-evaluable content. Vicente writes: 'It is up to pragmatics to select one of these possible meanings, i.e., to pick out some specific (but pre-existent) truth-conditions.' (Vicente, 2012, p. 14). As a consequence, the truth-conditions of an utterance of a sentence S are partly a pragmatic matter. Semantics delivers a set of possible readings for S. Which is the one the utterance expresses depends on pragmatic factors. However, I think the analysis is similar to accounts that assume Semantic Propositionalism in this sense: according to Vicente's analysis, semantics can be seen as determining a set of truth-evaluable contents–in this sense semantics delivers truth-evaluable content. Moreover, Vicente's proposal can be translated into a form of indexicalism by posting hidden context-sensitive variables corresponding to the conceptual information of the noun.

The fact that two different contents have been expressed, true under different conditions, explains the shift in truth-value. The analysis is intended to secure truth-conditional semantics vis-à-vis Travis cases.

In spite of the merits of these analyses, it has been repeatedly noted that parts are not the only source of shiftiness. Some Travis cases involve changes in the illumination conditions (the purple painting). Others cases involve changes in who is the relevant judge (swatch 27). Szabó's theory is unable to deal with these examples, for here we have variation in truth-conditions even though the values of the class of comparison and part parameters are kept fixed. Giberman's theory is also unable to account for these examples, for in them the result of the mereological fusion may remain constant ('is purple' refers to exactly the same object: the painting). In the case of Vicente's theory, we are not told how much information the lexical entry for 'leaf' includes, and we are not told anything about the lexical entry for 'painting'. However, the variation in the purple painting example is due to a property of colours: the colour we perceive an object as having depends on the observation conditions. It wouldn't seem very plausible to claim that the lexical entries for 'leaf', 'painting', 'swatch', etc. contain information about the metaphysics of colour.

Because of the existence of multiple sources of variation, more complex analyses have been provided. Hansen's (2011) approach, a modification on Kennedy and McNally's ambiguity view, is the most complete so far. Let me begin by introducing Kennedy and McNally (2010).

Kennedy and McNally (2010) present a solution to what they call Travis' green leaves puzzle based on the idea that 'green' is ambiguous. Their proposal is not intended as a reply to Travis' occasionalism, but only as an explanation of what is going on in the green leaves scenario. In order to do this, they provide a semantics for colour predicates. Roughly, their theory establishes that[33] colour predicates are ambiguous between three readings: (1) a gradable quantity reading, (2) a gradable quality reading, and (3) a classificatory reading, in which the predicate is used to distinguish objects on the basis of why they are green. The gradable quantity reading is the one we get when using expressions such as 'completely green' or 'half green', and substitutes Szabó's Part variable. The predicate is interpreted with respect to a scale of greenness[34]. How much green is

[33] Kennedy and McNally provide linguistic evidence supporting the ambiguity claim.

[34] 'Different degree morphemes introduce different kinds of standards that determine whether the property in question is held in sufficient degree for the predicate to truthfully apply to its argument. For example, the (unmodified) positive form involves a null degree morpheme pos that introduces a relation to a contextual standard of comparison' (Kennedy and McNally, 2010, p. 95).

green enough depends on the context (the disambiguated sentence is context-sensitive). The gradable quality reading works similarly, but here what matters are qualitative aspects (hue, etc.). Again, this gradable reading is context-sensitive. In the classificatory reading, 'having the property denoted by the color adjective is crucially correlated with having some other property or properties which are relevant for some purpose or other' (Kennedy and McNally, 2010, p. 88).

A difference in truth-value in two utterances of 'The leaves are green' might be thus explained as a difference in the sentence uttered. Pia's first utterance expresses a gradable-quality reading of 'green' (what matters is how the leaves look), whereas the second utterance expresses a non-gradable reading.

Despite the merits of Kennedy and McNally's semantics, their theory is insufficient as a defence of Semantic Propositionalism for colour predicates. Again, the analysis cannot deal with examples involving changes in observation conditions. Relying on a relational metaphysics of colour according to which colours are relational properties, Hansen (2011) adds three variables to the syntactic form of colour terms: frame of reference, observation conditions, and observer. These parameters, common to both the gradable and the non-gradable readings, allow Hansen to handle a variety of cases, including the purple painting example. The parameter standing for observation conditions can do the required work: its value shifts from the first utterance (daylight) to the second (blacklight). The analysis can also explain cases involving other forms of variation. Imagine someone saying 'Patch A is greener than patch B' while looking at two patches of grass on a shiny day. 'Greener' here could be used to describe, how the two patches look that day, or a more stable property. The variable standing for the frame of reference explains this variation. This variable can get two values, stimulus or object. Imagine that patch A looks greener on a shiny day, but that that difference is not stable. Most of the time, the two patches look exactly the same. If the speaker is talking about how the patch looks to he and his addressee (stimulus) on a shiny day, an utterance of 'Patch A is greener than patch B' will be true. If they are discussing about how those two patches are in general (object), it will be false.

What about Pia's leaves? Following Kennedy and McNally, Hansen claims that an ambiguity gradable/non-gradable explains the difference in truth-conditions. It is worth quoting Hansen at length ((1) is the sentence 'The leaves are green'):

> The painted leaves are projected by the color adjective onto a scale of greenness that is sensitive to the relevant frame of reference, conditions of observation and standard observer in the context. The standard value (contributed by pos in

> (1)) that the salient parts of the leaves have to meet or exceed in order to count as green is determined by features of the context. In C1, the leaves are projected to a point on the contextually determined scale of greenness that meets or exceeds the contextually determined standard value, so when Pia utters (1), she thereby says something true. In C2, the falsity of Pia's utterance of (1) is due to her using green to go proxy for some contextually relevant property (like naturally green). The difference in the truth conditions of what Pia says in the two contexts is explained by a form of ambiguity, and the central commitment of CTCS [Compositional Truth-Conditional Semantics] is not threatened. (Hansen, 2011, p. 219)

In the first utterance, 'green' is gradable. The context determines a scale and the parts that must be green in order for an object to count as 'green'. When Pia first says 'The leaves are green', the leaves are projected to a point on the scale that exceeds (or at least meets) what is contextually required to count as 'green'. In the second utterance, 'green' is non-gradable and, consequently, is proxy for another property. Hansen mentions the property naturally green. This explains the truth-value shift.

Instead of having recourse to hidden context-sensitive variables, Rothschild and Segal (2009) treat predicates as indexicals, like 'I' or 'that'. Their analysis goes as follows. They group together all tokens of 'green' in a context. All tokens of 'green' that occur in j-th context are grouped as being instances of the syntactic type red_j. Tokens of this type can only occur in j-th context. These context-bound types are subtypes of larger types. The semantics of the larger types is conceived as a function from contexts to extensions. In this framework, extension is determined by the conversational standards of the context: 'an object satisfies a token of 'red_j' in a context, if it counts as red by the standards of the context' (Rothschild and Segal, 2009, p. 472). This theory can easily accommodate Travis cases: what counts as red varies from context to context. Therefore, different objects satisfy 'red' on different occasions.

2.5 Assessment

Do the previous analysis respect conditions I and II? As I have already noted, it is doubtful that Szabó, Vicente and Giberman have identified a set of sufficient parameters. Observation and illumination conditions bring problems for these approaches. Moreover,

as Kennedy and McNally (2010) argue, colour terms are sometimes used to classify objects on the basis of their possession of other correlated properties. Vicente and Szabó's accounts might have problems with this. I focus on Vicente's. Vicente's account deals with the property of being originally green, but there are others. For example, one could use 'green' to describe an old wooden box that was painted green in the 13th century but which doesn't exhibit the colour anymore (in a conversation about arts and crafts during the middle age, in order to distinguish that box from other boxes that were, at that time, painted red, on the basis of the presence of some invisible chemical compound: 'The box we own is a green one! The chemical analysis has confirmed it.'). In order to account for this example, one would need to add more information to the lexical entry 'box' (something like: 'boxes have properties that apply to the original and apparent surface, where the apparent surface can be the current one or the one it had in the 13th century, etc.').

Moreover, it is still possible to create a further Travis case on the top of that. For example, a box could have been superficially green by accident (some paint fell on it) and not as a result of an artisan's work. Some people might be interested on the colour the box had at that time, other people might be interested only in artisans' works. So one would need to complexify again the information contained in the lexical entry ('boxes have properties that apply to the original and apparent surface, where the apparent surface can be the current one or the one it had in the 13th century, and where the apparent surface it had in the 13th century can be the result of an intentional or of an unintentional action'.) But as we complexify the rule it becomes less and less plausible that the information be in fact lexical, i.e., that this information is part of the meaning of 'box'.

As to Giberman's account, even if it could account for cases where there are changes in illumination conditions, it is not clear that the account respects condition II. Giberman writes: '[F]or subject term 'a' and predicate 'F', the "rule" for understanding the content of 'a' in an utterance of 'a is F' is roughly: use previously established conversational interests—including interests in the property expressed by 'F'—to select some objects to mereologically add to or subtract from the metaphysically perspicuous designatum of 'a' and select the resulting object as the content.' (Giberman, 2016, p. 113). Semantics is merely telling the interpreter to look for two objects, but it's not giving any clue as to which those objects are. Identifying them is a wholly pragmatic affair. The determination of a truth-evaluable content goes via pragmatic interpretation.

Since Kennedy and McNally's analysis is included in Hansen's theory, I will skip it. Does Hansen's proposal satisfy condition I? I have some doubts that all the parameters

posited are necessary for analysing sentences in which the colour is not predicated of an object, as in 'Imagine a red spot' (This example appears in (Recanati, 2010)) or 'Look, the rainbow is blue and green, etc.' It is not clear that the frame of reference variable is playing any role here. But let's leave this question aside (the analysis could probably be restricted to certain uses of colour predicates or made more complex)[35] . The problem I want to focus on is a different one. Due to the inclusion of a classificatory reading in which 'green' goes proxy for a correlated property, Hansen's analysis can be expected to explain a great variety of Travis cases. In all cases in which the variation cannot be explained as a variation in the part of the object that is supposed to be coloured, or in the observation conditions, observer or frame of reference, one can claim that 'green' refers to a different, although related, property. Now, the problem is that precisely because of this reading, Hansen's analysis does not respect condition II[36] . Which correlated property an instance of 'is green' expresses is not determined by the semantics of green—semantics just fixes that the predicate is used to express a different property, it does not contain any receipt as to how to identify that property[37] .

Let me quote how Hansen understands the project of Compositional Truth-Conditional Semantics to make the point clearer:

> The truth condition (or content, or proposition) of an occurrence of a sentence S in a context* C is determined by the semantic properties of the parts of S, their mode of combination, the context* C, and nothing else. (Hansen, 2011, p. 202)

By context* Hansen understands 'an n-tuple of features required to represent the contextually variable, but regular, contribution that context sensitive expressions like indexicals and demonstratives make to the truth condition of sentences.' (Hansen, 2011, p. 202).

[35] There are some examples involving objects covered by other objects, for example, tables covered by tablecloths. Here the predicate can apply to the table or to the tablecloth. Giberman's analysis fares better than Hansen's in this respect.

[36] It is not clear either that the determination of the value of the posited variables is purely semantic, but it is enough for my point to focus on the classificatory reading. This reading clearly calls for pragmatic interpretation.

[37] See Clapp (2012c) for a similar argument. Clapp argues that Kennedy and McNally's ambiguity view posits something equivalent to a lax discretionary indexical. The interpreter needs to consider what is reasonable to take the speaker to mean. Thus the approach obliterates the distinction between semantic interpretation and other forms of interpretation having to do with grasping speaker meaning.

We are not told how to apply this notion of context to the green leaves example. A reasonable interpretation is that this n-tuple of features includes the property for which 'is green' goes proxy (naturally green, in the green leaves example). However, if this is the proposal, then it is clear that it does not respect Semantic Propositionalism, since the semantics of 'is green' are not doing any work here in the determination of the property expressed. Whatever property is pragmatically determined goes into composition. The role of semantics would not be that of determining truth-conditions of S on an occasion of use but that of (1) constraining the extension of the expressions of S on an occasion of use and (2) composing those pragmatically obtained values. On the other hand, if the feature is not the property itself, then we need a further story about the features included in the context and how they determine which property is expressed on an occasion of use—a story we don't have yet.

Rothschild and Segal's analysis is subject to a similar worry, as Clapp has argued (Clapp, 2012a). The aspects of contexts that are relevant for the shifts in extension are not even specified by the character. It is a wholly pragmatic affair to determine the extension of the predicate on an occasion of use.

2.6 Conclusions

I have argued that neither minimalism nor indexicalism work as a defence of Semantic Propositionalism. The minimalist arguments offered so far don't manage to counter the pragmatist challenge. Some of them rely on there being context-insensitive ways of classifying things, which is precisely what Travis cases call into question. Others unwarrantedly rely on some privileged speakers or experts having access to the literal contents of colour predicates, assuming, further, that their verdicts will coincide with the minimalists preferred semantics. And yet others predict that, in quite ordinary situations, well-positioned competent speakers are ignorant of the literal truth-value of the sentences they use, thus casting doubts on the possibility of ever finding out the literal truth-value of the sentences used in Travis cases. On the other hand, the best indexicalist theories we have, that is, those able to deal with the corpus of examples that motivate the pragmatist challenge, have recourse to pragmatic interpretation. As a result, these are not theories in which semantics determines truth-evaluable content—pragmatics is

doing a crucial job. Although they might be compatible with both a principle of Semantic Compositionality and with a principle of Pragmatic Compositionality (see sec. 1), they are not a response to the pragmatist challenge, for the pragmatist is attacking Semantic Propositionalism, not (Semantic or Pragmatic) Compositionality.

At this point, it is important to consider why would one want to maintain Semantic Propositionalism[38] . A possible answer is that Semantic Propositionalism is on the basis of truth-conditional semantics and, as for instance Rothschild and Segal claim, 'Truth-conditional semantics is the major research project of linguistic semantics and the project and its prospects are a central concern in contemporary philosophy of language.' (Rothschild and Segal, 2009, p. 467).

Semantic theories are often conceived as interpretive theories, i.e., as theories that deliver interpretations for indefinitely many sentences in a language L[39] . These theories are supposed to model our ability to interpret speech. As an example, Larson and Segal

[38]One of the reasons Borg (2012) offers in support of minimalism has to do with the modularity of mind. Her idea is that the mind has a module for semantic interpretation. However, a pragmatist position as Travis' is perfectly compatible with this. On an occasionalist view, pragmatics do not play any role in the determination of semantic content, for semantic content is only content that sentences have qua types. By the same token, an occasionalist agrees with Borg that it is not the role of semantics to account for speech act content. The disagreement concerns whether semantics thus viewed can deliver something truth-evaluable. The point of disagreement is not whether semantics is free of pragmatic intrusion, it is whether semantics can deliver something truth-evaluable, or only an indeterminate content compatible with different truth-conditions.

[39]This is not the only aim of semantic theories. One of the things semantics is often taken to explain are logico-semantic relations between sentences. For example, one might be interested in how the truth of a sentence relates to the truth of other sentences, or about synonymy or contradiction. As Davies (2011) has argued, abandoning Semantic Propositionalism is not a problem for this aim, for we can distinguish substantial and logico-syntactic truth-conditions. In formal theories we assign truth-values to sentences and see how the values of other sentences in the theory depend of those assigned truth-values. These are logico-syntactic truth-values: they are either stipulated or derived from stipulated truth-values. By contrast, the substantial truth-value is the truth-value a given sentence has as assessed against the actual word. Semantic Propositionalism is about truth in this second sense. Similarly, one can take formal semantics to provide the tools to solve philosophical puzzles. The idea here is that by uncovering the complexities of a language one might be able to gain some insight into philosophical problems. As an example, consider the sentence 'John believes that the tallest spy is in Paraguay' (Dever, 2014). This sentence contains an ambiguity, it can be read de dicto or the re. A semantic theory might reveal such an ambiguity. As sentences become more complicate ('John believes that some linguist followed every philosopher'), semantic theories become more useful in philosophical discussion. Again, abandoning Semantic Propositionalism need not be a problem for this view, for we can have semantic theories working with logico-syntactic, not substantial, truth-conditions. This, together with a revision of the logical form of the problematic sentence, is enough to account for the possibility of having different readings.

write:

> A speaker's knowledge of meaning for a language L is knowledge of a deductive system (i.e., a system of axioms and production rules) proving theorems of the form of (T) that are interpretive for sentences of L. (Larson and Segal, 1995, p. 33).

(T) is, of course, the following schema:

> (T) S is true if and only if *p*.

These theories take our knowledge of meaning (together with our knowledge of composition rules) to be sufficient for us to interpret an utterance, where interpreting an utterance means knowing under what conditions the sentence uttered would be true. Renouncing to Semantic Propositionalism is a problem for this way of seeing the role of semantics, for it amounts to rejecting the idea that meaning is the route to extension—that we can identify knowledge of the meaning of 'is green' (together with the meaning of the other expressions and syntax) with something that would be sufficient for us to grasp the truth-conditions of Pia's utterances and, in particular, to grasp what counts as 'green' on the occasions described. If knowledge of meaning is not the key to interpretation (to truth-conditional content), then the project is doomed to failure—our knowledge of meaning is not knowledge of a deductive system proving theorems of the form of (T)[40] . We need to provide a different explanation of our ability to interpret speech.

This might sound as a strong reason for trying to secure Semantic Propositionalism. However, if we look at current defences of Semantic Propositionalism there are some doubts that the original project is still being pursued. I start with indexicalism. As I have argued, there are two indexicalist theories that can account for Travis cases, Hansen (2011) and Rothschild and Segal (2009). Are these a development of Larson and Segal's mentioned project? I think they are not. Both theories have already given up the idea that knowing the meaning of 'is green' is sufficient for a speaker to grasp the truth-conditions of Pia's utterances. As I have argued, both in Hansen's (2011) and in Rothschild and Segal's (2009) theories, knowledge of truth-conditions requires something more than knowledge of meaning and syntax—it requires pragmatic interpretation, the ability to grasp what counts as 'green' in a particular occasion of use, which

[40] On this point, the pragmatist position bears some similarities with (Chomsky, 2000). One way to take apart the identification of meaning with extension is to go for an internalist view of meaning (see (Pietroski, 2003)).

property the predicate expresses. And this is something that knowledge of meaning alone will not get us.

What about minimalism? It is also doubtful that minimalism is in the business of providing semantic theories as theories of interpretation, but for a different reason. In order to keep the kind of semantic theories they like (ones where predicates have a stable extension), minimalists need to explain away the actual interpretation many users of language make of the utterances described in Travis cases. However, because of this, it is not at all clear that their theories model *our* linguistic knowledge. If we took speakers' judgements as a test for correctness of semantic theories, minimalist theories will not be in a good position. Because of this, minimalists aim at dismissing the kind of judgements Travis cases trigger. However, in doing so, they loose all warrant that their semantic theories are interpretive theories that model our linguistic competence. If we dismiss ordinary speakers' judgements about truth-values, how can we be sure that minimalist theories are correct as theories that model the semantics of *our* languages?

Chapter 3

Are mental representations underdeterminacy-free?

According to some views (Carston, Fodor), natural language suffers from underdeterminacy, but thought doesn't. According to the underdeterminacy claim, sentence types underdetermine the truth-conditions of sentence tokens. In particular, the semantics of a predicate type seems to underdetermine the satisfaction conditions of its tokens. By contrast, mental representation-types are supposed to determine the truth-conditions of its tokens. If it works, this approach would motivate a certain version of truth-conditional pragmatics whose defining feature is its reliance on what I will call 'Pragmatic Propositionalism'. In this chapter I critically examine these mixed views. First, I argue that the arguments supporting the indispensability of including in one's theory mental representations that are free of the underdeterminacy exhibited by natural language are not sound. As a result, the possibility that mental representation-types are as underdetermined as natural language sentence-types has not been ruled out. Second, I argue that Carston's *ad hoc* concept-types are as underdetermined as word-types. After this, I argue that mental representations are also underdetermined in a second sense—mental representation-tokens only determine a partial function from possible worlds to truth-values. I finish with a discussion about the possibility that structured propositions (not mental representations) are underdeterminacy-free.

3.1 Underdeterminacy

As I have argued in the previous chapters, there are reasons to think that the truth-conditional content expressed by an utterance of a well-formed declarative sentence S is not determined by the semantics of S. Again, I will focus here on the underdeterminacy that can be traced to predicates. The contribution a large class of predicates make to the truth-conditions of an utterance seems to be affected by contextual factors that are not part of the predicate's meaning. Consider the following examples, discussed by Carston (2002)[1] :

(1) The kettle is black.
(2) Anne is happy.
(3) I want to meet some bachelors.

Imagine an old aluminium kettle that, after years of use, has turned black on most of its surface. As has been repeatedly noted, sentence (1) can be used to describe the original colour of the kettle, or the observable colour of most of its surface. The truth-conditions of an utterance of (1) will consequently vary, and two utterances of (1) used to describe the old aluminium kettle (in the same state) can have different truth-values. Similarly, 'happy' can be used to describe a range of positive emotions. Someone can be rightly described on an occasion as 'happy' because of being momentarily in a state of intense joy, whereas in other occasions a more stable positive feeling is required in order to count as 'happy'. As to the last sentence, Carston notes that 'bachelor' can be used to refer to unmarried man or in a more restrictive sense excluding men that have committed themselves to celibacy. Thus, if a woman tells her friend that she is going through a divorce and wants to meet some bachelors, 'bachelor' would most likely be restricted to unmarried heterosexual men. Moreover, who exactly counts as an unmarried man can be decided differently on different occasions. The extension of 'bachelor' on an occasion of use could also include men that are currently going through a divorce and exclude men that are involved in long-term relationships (men who have common-law partners but are not legally married), or exclude men married in a religious ceremony but not legally (or the other way round).

As a result, sentences (1)-(3) can express different truth-conditions in different con-

[1] The reason why I focus on these examples is that these are the ones Carston uses in order to illustrate her view.

texts, even when the referent of indexicals and definite descriptions is kept fixed. In this sense, their linguistic meaning can be said to underdetermine their truth-conditions on an occasion of use—the satisfaction conditions of some expressions in those sentences are sensitive to facts about the conversation in which they are used and, probably, to the interlocutors' interests and communicative intentions. If that is so, then we need truth-conditional pragmatics.

The question now is: What sorts of representations exhibit truth-conditional underdeterminacy? I think that, besides natural language sentences, underdeterminacy could be a feature of Mentalese sentences (mental representations). In particular, both sorts of representations might exhibit what I will call Type-Underdeterminacy[2] :

> *Type-Underdeterminacy:* A non-indexical structured representational item[3] S is type-underdetermined if and only if there are tokens of S that have distinct truth-values[4] .

That tokens of S have distinct truth-values means that the type does not determine a unique truth-value (given the state of the world). The tokens are true under different conditions[5] . The previous considerations about sentences (1)-(3) suggest that *linguistic* representational items, i.e., sentences, exhibit Type-Underdeterminacy. Once it is accepted that natural language underdetermines truth-conditions in this sense, it might be disputed whether other representational items are also type-underdetermined. We can take Type-Underdeterminacy to be a feature of representations in general, and generalize the underdeterminacy claim to other structured representational systems—saliently,

[2] In section 5 I will introduce a different notion of underdeterminacy, namely Token-Underdeterminacy.

[3] By 'structured representational item' I mean a representational item that is identified by its structure, such as a sentence. The definition is intended to apply both to sentences in natural language and to sentences in Mentalese.

[4] Type-Underdeterminacy is equivalent to occasion-sensitivity as I introduced the term in chapter 1. The reason why I am using the term 'Type-Underdeterminacy' here is that advocates of truth-conditional pragmatics, and in particular Carston, talk of 'underdeterminacy'. Saying that a sentence is type-underdetermined is also equivalent to saying that it fails to be truth-condition compositional (more on this later).

[5] This principle might sound trivial in case the state of the world changes from the first to the second tokening context. However, in the underdeterminacy scenarios that advocates of truth-conditional pragmatics discuss (such as Travis cases), the change in truth-value is not due to a change in the state of the world. Moreover, if the time at which the sentence is evaluated affects the truth-value of the tokens, then it must be incomplete with respect to time and, in this sense, underdetermined.

to mental representations. Or, we can take it to be restricted to language and try to find another representational system that is, in this sense, underdeterminacy-free. Again, the candidates are mental representations. Following this second option, the underdeterminacy detected in natural language can be explained as concerning the relation between natural language sentences and other representational items.

Fodor (2001) and Carston (2002) put forward an approach in which language suffers from Type-Underdeterminacy, but thought does not. Given that they endorse the Language of Thought (LOT) hypothesis, their claim is to be understood as saying that natural language sentence are underdetermined but LOT sentences, i.e., mental representations, are not. The underdeterminacy detected in sentences (1)-(3) concerns, according to this view, the relation between utterances and the mental representations expressed by them. What is underdetermined is what mental representation corresponds to an utterance of a natural language sentence S, S's meaning being compatible with, let's say, two different mental representations. Mental representations are seen as fully propositional. Pragmatics is supposed to bridge the gap between sentences and mental representations: our mind reading abilities allow us to infer which thought (i.e., which mental representation) a given utterance expresses. I will call approaches that fit this second option 'mixed views'.

A feature of the mixed view that is open to criticism has to do with its reliance on the existence of items (mental representations) that are identified by a structure and have fully determined truth-conditions independently of the context of use[6] . The first challenge for these approaches arises from the fact that we haven't be shown that other representational system behave differently from natural language when it comes to truth-conditions. Natural language sentences are representational items (i.e., items that we typically use for representing) that are identified by structure, i.e., syntax. If two tokens have the same syntactic form, then they are tokens of the same sentence. According to the Type-Underdeterminacy claim for natural language, there is a gap between these structural items and the truth-conditions they express on an occasion of use. In the case of natural language sentences, structure doesn't automatically get us truth-conditions. The problem is that, once we have seen that this is what happens with sentences, the assumption that there are other representational items identified by structure with context-independent truth-conditions needs to be justified. Why is it not the case that mental representations, like sentences of natural language, also express different truth-conditions at different tokening contexts? This line of reasoning

[6] Travis (2000) calls this the Janus-faced picture of thoughts.

can be used in order to motivate a generalization of Type-Underdeterminacy to mental representations[7].

The point I will address in this chapter concerns the relation between representations in the mental realm and truth-conditions. The plan is the following. In section 2, I will present the mixed view in more detail. In section 3, I will point a consequence of the mixed view—namely, ineffability. In section 4, I will assess what I call the indispensability arguments. Proponents of the mixed view argue that there are strong reasons for including in one's theory these structured representational items that are (allegedly) free of Type-Underdeterminacy. The main idea here is that they play a role that only representation that is free of Type-Underdeterminacy could play. In this sense, they are indispensable. If their arguments work, then we have reasons for positing fully determined mental representations. Against these views, I will argue that their arguments fail to establish that free of Underdeterminacy-Type representations are indispensable. As a consequence, alternative approaches are, at least, tenable. In section 5, I will argue, that there are reasons to doubt that Mentalese sentences are unlike natural language sentences regarding Type-Underdeterminacy. In particular, I will argue that Carston's description of the process of *ad hoc* concept creation suggests that Mentalese sentences are as type-underdetermined as natural language sentences. In section 6, I will introduce a second sense of 'underdeterminacy'—Token-Underdeterminacy—and argue that mental representations are also underdetermined in this second sense. In section 7 I will examine whether having recourse to propositions, instead of mental representations, could be a reasonable strategy for a proponent of a (different version of) the mixed view.

3.2 The mixed view

The mixed view can be summarized as follows:

> The variety of examples that contextualists have put forward involving intuitive variation in truth-conditions of non-indexical sentences across contexts show that linguistic meaning systematically underdetermines the

[7] This generalization has been pursued by Travis (2000) and Searle (1983).

> proposition expressed by a given sentence on an occasion of use. Identifying the proposition expressed always, or nearly always, goes via pragmatic interpretation (whatever exactly that amounts to).

What is wrong with this view? It is important to note that it shares an important feature with traditional approaches. Even if pragmatic interpretation is mandatory, the view relies on the assumption that there are classical propositions to be had. How should we understand this talk of propositions? Propositions can be simply conceived as truth-conditions. Talk about different propositions being expressed might be equivalent to talk about different truth-conditions being expressed. Let me rephrase the view with this notion of proposition in mind:

> The variety of examples that contextualists have put forward involving intuitive variation in truth-conditions of non-indexical sentences across contexts show that linguistic meaning systematically underdetermines the truth-conditions of a given sentence on an occasion of use. Identifying truth-conditions always, or nearly always, goes via pragmatic interpretation (whatever exactly that amounts to).

There is, nonetheless, another possible way of conceiving propositions: propositions can also be conceived as structured strings of items (such as properties, senses or concepts) with context-independent truth-conditions. In this second sense, it would not be adequate to replace the reference to propositions with a reference to truth-conditions. Even if propositions are conceived as the primary bearers of truth-conditions, they are supposed to be so in virtue of having a certain structure—for example, in virtue of attributing a property to an individual. I will be concerned with this second notion of 'proposition'.

The mixed view makes three claims. First, it claims that linguistic representations, i.e., sentences with their linguistic meanings, or sentences plus the contextual information determined by their linguistic meanings, type-underdetermine truth-conditions. Second, it claims that, besides sentences, there are other items that do not type-underdetermine their truth-conditions on an occasion of use. Third, it claims that which of this second kind of item is expressed by a use of a sentence is to be determined via pragmatic interpretation.

I thus take the defining feature of the mixed view to be its reliance on Pragmatic Propositionalism:

> *Pragmatic Propositionalism:* there are free of Type-Underdeterminacy propositions to be had. Which proposition a given utterance expresses is to be determined[8] pragmatically (for example, by having recourse to speaker meaning or the topic of the conversation) rather than semantically[9] .

As I have mentioned, I am focusing on views that work with structured representational items, not on views that understand 'proposition' as 'truth-conditional content' o 'set of possible worlds'. Hence, Pragmatic Propositionalism must be understood as follows:

> *Pragmatic Propositionalism (structured propositions):* there are free of Type-Underdeterminacy structured representational items to be had. Which structured representational items a given utterance expresses is to be determined pragmatically (for example, by having recourse to speaker meaning or the topic of the conversation) rather than semantically.

In Carston and Fodor's views, these contents that are free of Type-Underdeterminacy correspond to mental representations. In the last section of the paper I will consider the possibility that they correspond to structured propositions.

Carston's theory is explicitly a version of the mixed view. Carston (2002) writes:

> [T]he position I've been arguing for is that there are no eternal sentences in natural languages (that is, no sentences which encode a proposition or thought which is constant across all contexts), from which it follows that the linguistic underdeterminacy of the proposition expressed by an utterance is an essential feature of natural language (Carston, 2002, p. 42).

[8] In the metaphysical and epistemic sense.

[9] According to Borg, this is the defining feature of contextualism. She writes: 'According to the contextualist there is determinate, context-insensitive content to be had, their objection is just that the content provided via the lexico-syntactic constituents of the sentence alone isn't it. Yet min- imalists and occasionalists agree that, if you are swayed by the phenomenon thrown up by CSAs [Context-Shifting Arguments] at all, then this is a reason to think that no (or perhaps almost no) content offered in a context-independent manner will ever reach the standard of a complete proposition. For instance, the contextualist will want to claim that, though a sentence like 'the apple is red' expresses an incomplete or inappropriate content if we look just to the lexico-syntactic constituents of the sentence, there is a complete and appropriate proposition to be had (perhaps at the level of thought).The task then is to get from the former to the latter.' (Borg, 2012, p. 36)

> It is fully propositional conceptual representations, rather than sentences, or even utterances of sentences, that are the primary bearers of truth conditions (Carston, 2002, p. 60).

According to Carston, there are propositions to be had, they are just not encoded by natural language sentences. Rather, they correspond to mental representations. Strings of concepts are supposed to be free of Type-Underdeterminacy. The arguments that Carston presents in support of the underdeterminacy claim include the examples presented at the beginning of this paper. Concerning predication, she relies on Travis's examples[10] .

Given that linguistic meaning underdetermines the thought expressed by a sentence on an occasion of use, there must be some mechanism that enables the interpreter to grasp the thought that the speaker intends to communicate by means of a sentence, i.e., a mechanism that bridges the gap between sentences and propositional thoughts. In Carston's theory, the mechanism consists in the creation new concepts, slightly different from the concepts encoded in language. Concepts of this second kind are called *ad hoc* concepts. The idea is that the interpreter uses his ability for pragmatic interpretation in order to recover the thought that the speaker intends to communicate by using as input the encoded content of a sentence and the available contextual information. Often the recovery of the intended thought is achieved via the adjustment of the encoded concepts. In those cases, a new concept is 'constructed on-line (on the fly) in response to specific expectations of relevance raised in specific contexts' (Carston, 2002, p. 322).

Carston follows relevance theory in her understanding of what a concept is. Atomic concepts consist on three kinds of information: logical content (a set of inference rules capturing analytic implications of the concept), encyclopaedic knowledge (scientific information, general knowledge about the object, personal observations) and lexical properties (phonological and syntactic properties). Complex concepts are structured strings of atomic concepts. Language codifies concepts such as CAT[11] , with certain logical content (if something is a cat, then it is an animal), encyclopaedic knowledge (cats are domestic animals, visual images of cats), and lexical properties. As I said, in a conversation, in the process of utterance interpretation, interlocutors construct *ad hoc* concepts by adjusting the information of the lexically encoded concepts to the specifics

[10] Besides rejecting eternal predication, Carston also rejects eternal reference. However, here I will only consider the underdeterminacy that can be traced to predication.

[11] I use capital letters for encoded concepts. *Ad hoc* concepts are marked with an asterisk.

of the context. As is common, Carston distinguishes two pragmatic processes of adjustment: narrowing and broadening. In cases of narrowing, the concept is made more specific. Let us imagine a use of (3) by a woman who is chatting with a friend about her desire to meet a man, get married and have children (Carston, 2002, p. 326). The encoded concept BACHELOR makes reference to non-married men. The encyclopaedic entry might contain information about different types of bachelors: irresponsible and forever-young-and-free, capable of long-term commitment, etc. However, in this conversation the concept expressed by the word 'bachelor' is more specific: the speaker wants to meet some men eligible for marriage (capable of long-term commitment, heterosexual). Given what the hearer knows about the speaker (marital interest), during the process of interpretation a new *ad hoc* concept will be constructed excluding in the encyclopaedic entry features standardly associated with bachelors as irresponsible and forever-young-and-free. The account for broadening is symmetrical. Consider the sentence 'France is hexagonal'. France is not a geometrical hexagon. However, loosely speaking, its shape can be considered hexagonal. What is going on here is that whereas the concept of HEXAGON includes only strict hexagons, the *ad hoc* HEXAGON* includes shapes that deviate to some degree.

Let me pause on two problematic aspects of Carston's proposal. The first problematic aspect has to do with the relation between encoded concepts and *ad hoc* concepts. It is not at all clear whether lexically encoded concepts as HEXAGON are being conceived here as determining an extension, or whether it is only *ad hoc* concepts as HEXAGON* that do. In her (2002), Carston holds both that there is no eternal predication (one of the reasons why natural languages are underdetermined), and that lexically encoded concepts have extensions. For example, she writes, about narrowing, that 'the extension of the concept pragmatically constructed is a subset of the extension of the lexical concept from which it has been derived.' (Carston, 2002, p. 325). Now, this presupposes that the lexical concept has an extension. However, Carston's radical underdeterminacy claim is incompatible with lexically encoded concepts having extensions: if the lexically encoded concept BACHELOR has an extension, then 'is a bachelor' should have a constant extension. And if so, then there is eternal predication. Nonetheless, simply dropping the assumption that lexically encoded concepts determine extensions is also problematic, for encoded and *ad hoc* concepts have the same structure. As long as encoded and *ad hoc* concepts are described in analogous terms it is mysterious why they should behave differently. In order to avoid this, one could argue that words do

not encode concepts, but instead that they point towards conceptual spaces[12] . This option has an unpalatable consequence: if there are no concepts corresponding to our words, then there are no concepts corresponding to the unmodified 'black', 'bachelor', 'happy', etc.

The second problematic aspect has to do with the supposed atomic character of *ad hoc* concepts. *Ad hoc* concepts, just like lexically encoded concepts, are supposed to be atomic[13] . However, it is not clear what is meant by 'atomic' here, given that they are created by adding or subtracting information from already existing atomic concepts. For instance, if we go to Carston's description of the process of creation of the concept BACHELOR* (see quote on pp. 91-92), it is not at all clear that it should be considered atomic instead of the result of composing the concepts BACHELOR, ELIGIBLE and MARRIAGE—thus giving rise to the complex ELIGIBLE FOR MARRIAGE BACHELOR. I will argue (sections 3 and 4) that both options are problematic. If *ad hoc* concepts are atomic, then Carston cannot avail herself of the productivity argument (section 3.1)—something she can do if *ad hoc* concepts are complex. However, if they are complex, as the description of the creation of BACHELOR* suggests, i.e., if they are compositions of lexically encodable concepts, then one should expect that they be as underdetermined as combinations of words are (section 4.1).

Fodor (2001) also holds a version of the mixed view. He takes compositionality to be non-negotiable and argues that, between language and thought, whichever is compositional is the one that has content in the first place. He further takes language to be not compositional. Language being non-compositional seems to mean that linguistic meaning does not determine the (truth-conditional) content of complex expressions (of declarative sentences)—i.e., that language is type-underdetermined. By contrast, the contents of the simple constituents of thought, together with a mode of composition, are supposed to determine the (truth-conditional) content of complex thoughts.

I think that Recanati's contextualism could also be read as a version of Pragmatic Propositionalism. According to Recanati (2004; 2010), it is often the case, when a sentence is used, that its linguistic meaning is adjusted so as to fit the context. He writes:

[12] Carston explores the possibility that the meaning of substantive words (nouns, verbs, and adjectives) is procedural, without fully endorsing it, in her (2016).

[13] 'This term [*ad hoc*] is used to refer to concepts that are constructed pragmatically by a hearer in the process of utterance comprehension. The idea is that speakers can use a lexically encoded concept to communicate a distinct non-lexicalized *(atomic)* concept, which resembles the encoded one in that it shares elements of its logical and encyclopaedic entries, and that hearers can pragmatically infer the intended concept on the basis of the encoded one.' (Carston, 2002, p. 322). Emphasis added.

> When someone talks of 'wearing rabbit', the literal meaning of the mass term 'rabbit' (namely *rabbit stuff*) is accessed, but it has to compete with other candidates for semantic value. The more specific representation *rabbit fur* is also activated since it is associatively connected to the representations encoded by both 'rabbit' and 'wear'. As a result of this multiple activation, it is possible for the representation *rabbit fur* to be more active, in this context, than the less specific representation *rabbit stuff* which is linguistically encoded. (Recanati, 2004, p. 24)

Recanti also refers to properties and senses:

> Similarly, the expression 'ham sandwich' in 'The ham sandwich left without paying' arguably denotes, through transfer, the derived property HAM-SANDWICH-ORDERER rather than the linguistically encoded property HAM-SANDWICH. (Recanati, 2004, p. 26)

> [Modulation,] a family of primary pragmatic processes that make it possible to adjust the meaning of words and phrases in response to conversational needs, by endowing them with contextual senses distinct from their literal meanings. (2010)

Depending on how we understand this talk of 'representations', 'senses' and 'properties', the approach will or will not be an instance of the mixed view or not. As it is described, the outcome of modulation could be a classical structured proposition (the second quote is about properties) or a mental representation conceived as fully propositional. As an alternative, the modulated value of a predicate could be an extension or intension. If so, what the utterance expresses would not be a representational structure with context-independent truth-conditions, but simply a set of truth-conditions.

3.3 Effability

Both Fodor and Carston hold a representational theory of mental content according to which mental representations are sentences in LOT. Whereas natural language sentence types are taken to suffer from Type-Underdeterminacy, Mentalese sentence types are

supposed to be free of it[14] . As a consequence, natural language sentences are not apt to encode Mentalese sentences. In this sense, thoughts are ineffable—although we can often express them, we cannot find a sentence whose linguistic meaning corresponds to them.

This ineffability need not be a problem for communication. When we engage in a conversation and interpret an utterance we typically have access to information beyond the linguistic meaning of the sentence uttered. We might use this information in order to infer the thought that the speaker intends to communicate. In spite of this, it is important to see how radical the implied ineffability is.

First thing to note is that the Type-Underdeterminacy of natural language cannot be overcome by coining new words. Let us go back to the example involving the word 'happy'. On an occasion of use, the sentence 'Anne is happy' is used to describe Anne as, roughly, having a mild sense of acceptance of life. Suppose that we coin a new word, 'quappy', express this concept and add it to our vocabulary. As soon as we start using 'quappy' in different contexts it will most likely behave like any other word in natural language and, in turn, underdetermine its satisfaction conditions on an occasion of use. Imagine that some uses 'quappy' to describe Mary's feelings. Now we need to decide how similar to Anne's Mary's feelings have to be for the description 'Mary is quappy' to be true. But nothing prevents us from solving this question differently in different occasions, just like what counts as 'happy' varies across occasions. If so, 'quappy' will be type-underdetermined.

Second thing to note is that if conscious thought occurs in natural language, as introspection suggests[15] , then conscious thought also exhibits Type-Underdeterminacy. This means that (assuming that introspection is a reliable source here) conscious thought is as type-underdetermined as natural language. This does not imply that conscious thought lacks truth-conditional content. Just as we do with speech, we can distinguish conscious thought-types from episodes of thinking. Whereas the former, just like sentences, might not encode truth-conditional contents, the second can express them. In particular, this need not be a problem for Carston's account, for she can argue that episodes of conscious thought, as well as utterances of natural language, express mental representations (Mentalese sentences) and inherit from them their truth-conditions.

However, the Type-Underdeterminacy of conscious thought could be a problem

[14] Carston (2002) admits some context-dependence in thought, namely the presence of indexicals. However, she doesn't admit context-sensitive predicates at the level of thought.

[15] See (Carruthers, 1996).

for views that rely on availability as Recanati's, as Martínez-Manrique and Vicente (2004) argue. Here is Recanati's Availability Constraint: 'What is said must be intuitively accessible to the conversational participants (unless something goes wrong and they do not count as 'normal interpreters')' (Recanati, 2004, p. 20). If conscious thought occurs in natural language, then it is not clear in what sense what is said is accessible to normal speakers, given that what is said is supposed to be truth-conditional and natural language sentences are not[16] .

To sum up, the problem with ineffability is that it has the unpalatable consequence that those mental representations that are free of Type-Underdeterminacy can neither be encoded in a natural language sentence nor consciously entertained. Strictly speaking, we cannot encode those representations in natural language. We can express our thoughts via natural language, i.e., communicate them by using a sentence, but this will involve *ad hoc* concepts. Moreover, it seems that we cannot consciously entertain them. If conscious thought at least sometimes occurs in natural language, as introspection suggests, then conscious thought might also fail to encode determinate mental representations-types . What are the grounds, then, for assuming that there is a level of representations that are similar to natural language sentences yet not type-underdetermined?

3.4 The indispensability arguments

In this section, I will review and reject two arguments that have been offered in support of the claim that mental representations (Mentalese sentences) cannot exhibit Type-Underdeterminacy (Fodor (2001) and (2003), Carston (2002)). Fodor presents the arguments as concerning compositionality and does not distinguish them. However, in order to appreciate the dialectics of the discussion it is better to assess them separately.

There are different principles of compositionality. If we focus on natural language we can distinguish, at least, compositionality of meaning from truth-conditional compositionality[17] :

[16] I think that if what is said is understood as truth-conditions, the problem can be avoided, for we need not conceive 'being intuitively accessible' as equivalent to being consciously entertained.

[17] See (Searle, 1980) and (Clapp, 2012b) for similar distinctions. Clapp notes that this distinction un-

> *Meaning compositionality:* The meaning of a well-formed declarative sentence S is determined by the meaning of the expressions in S and the syntactic structure of S.
>
> *Truth-conditional compositionality:* The truth-conditions of a well-formed declarative sentence S are determined by the semantics (or the meaning) of the expressions in S and the syntactic structure of S.

It is this second principle that is at stake here[18] . The idea is that the truth-conditional content of a representation is exhausted by the semantic content of its simple constituents and their arrangement. If natural language is type-underdetermined, then it is not truth-conditional compositional[19] . However, it can still be meaning compositional. If mental representations are also type-underdetermined, then they are not truth-conditional compositional either, which means that their truth-conditional content is not exhausted by the concepts that form them and their arrangement. Their truth-conditions would depend on something else, as Travis sometimes puts it.

Why do mental representations have to be truth-conditionally compositional? Fodor (2001) addresses the question whether it is language or thought that has content in first instance. He claims that compositionality is non-negotiable and takes it that between language or thought the one which is compositional is the one which has content in first instance. Given the linguistic evidence (including the underdeterminacy arguments), it seems that language is not compositional. Because of this, Fodor assumes that thought is. However, one must ask why compositionality is supposed to be non-negotiable. As it is usual, Fodor mentions productivity and systematicity: 'Nobody knows exactly what compositionality demands, but everybody knows why its demands have to be satisfied. Here too the arguments are familiar; and, in my view, they're decisive. Both human thought and human language are, invariably, productive and systematic; and the only way that they could be is by being compositional.'[20] (Fodor, 2001, p. 6). I will focus on productivity. So the first argument is that mental representations have to

dermines the systematicity and productivity arguments. He focuses on systematicity.

[18] Fodor (2001) does not state any principle of compositionality. However, it is in this sense of 'compositional' that language fails to be compositional.

[19] Following Carston (2002), I have framed the discussion in terms of underdeterminacy. However, instead of talking about sentences of natural language and Mentalese being Type-Underdetermined one could talk of them failing Truth-conditional Compositionality.

[20] It is already odd that Fodor is here taking language to be compositional, when he explicitly rejects it. I will not try to solve this apparent contradiction.

be compositional because they are productive, and compositionality is the best (or the only) explanation we have for this.

This argument, as I will show, is off-target. Given Fodor's notion of productivity, meaning compositionality is sufficient for a system of representations to be productive. Moreover, if it is creativity we are interested in, or our ability to think new thoughts, then the best explanation given the scenarios of underdeterminacy involves the creation of *ad hoc* concepts, as Carston defends. But then productivity is not the key to our ability to think new thoughts—the creation of new concepts is. It follows that a system that is not truth-conditional compositional can still be productive, and that some forms of creativity are unrelated to productivity. So Fodor hasn't in fact provided any reason in support of the claim that mental representations must be free of Type-Underdeterminacy.

There is, however, a second argument. It has to do with the individuation of content. The idea is that if mental representations were not truth-condition compositional, then thought would be ambiguous or equivocal: the same representation could express different truth-conditions. The outcome would be that we wouldn't be able to tell some thoughts apart. I will argue that the conditional is false. Mental representations (types) could be equivocal, yet the tokening context could resolve the equivocation.

3.4.1 Productivity

This is Fodor's notion of productivity: 'Productivity is the property that a system of representations has if it includes infinitely many syntactically and semantically distinct symbols.' (Fodor, 2001, p. 6).

Natural languages are productive in this sense: there is an infinite number of well-formed, meaningful sentences. Given that the number of simple expressions is finite, compositionality is regarded as explaining the productivity of language: there can be an infinite number of meaningful complex expressions (sentences) because their semantics are determined by the semantics of simple expressions plus their syntax.

The relation between productivity and compositionality is often used to explain our ability to understand new sentences. As has been repeatedly noted, we are able to understand utterances of sentences we have never heard before. The best explanation

for this ability is that we understand these new sentences because we know the meaning of the simple expressions that form them and the syntactic rules. A finite mind can thus be reconciled with the capacity to interpret an infinite number of sentences.

Despite this argument, it has been questioned that knowledge of meaning and syntax suffices for working out the truth-conditions of an utterance, where this questions the idea that language is truth-conditionally compositional. Searle (1980) argues that knowing the meaning of 'cut', 'the', and 'sun' doesn't automatically enable us to understand an utterance of 'Cut the sun!'—we might fail to see what are the satisfaction conditions of this order, what exactly we are supposed to do. We lack some background that enables us to see what action would fit the order. This suggests that language is not truth-conditional compositional, but only meaning compositional. There certainly is something we understand when we first hear 'Cut the sun!', but we don't automatically grasp satisfaction-conditions.

What about thought? As Fodor (2003) notes, we are able to entertain new thoughts–potentially, an infinite number of them. However, being finite creatures, we only possess a finite number of simple concepts. Mental representations being productive would explain our infinite ability. Again, the reason why we are able to entertain an infinite number of thoughts could be that simple concepts can be arranged so as to form an infinite number of complex mental representations. Now, it is important to note two things. First, as Fodor defines productivity, something equivalent to meaning compositionality would be sufficient for mental representations to be productive. Being able to entertain an infinite number of mental representations only requires that simple concepts can be arranged so as to form an infinite number of complex concepts. However, Meaning Compositionality for mental representations is compatible with mental representations exhibiting Type-Underdeterminacy (i.e., with them failing to be truth-condition compositional).

Second, it is one thing to have the capacity to form or entertain indefinitely many mental representations, it is another to have the capacity to entertain indefinitely many truth-conditional contents. One could have the latter without having the former. Imagine a group of people who only possess two concepts, HUNGRY and THIRSTY, and no rule of composition (so no complex mental representation). Are these people only capable of entertaining two truth-conditional contents? No. Imagine that one of them tokens the concepts HUNGRY at 10am. And then, he tokens it again at 5pm. Even if the mental representation does not include the time, it is possible that the tokening context adds it and that the truth-conditional contents he entertains are different—the

first is true if and only if he is hungry at 10am, whereas the second is true if and only if he is hungry at 5pm. So he can entertain an indefinite number of truth-conditional content by tokening only one simple concept. As a consequence, our ability to think new thoughts does not inescapably go via compositionality.

What about Carston's view? Carston's concern is directly related to the underdeterminacy scenarios. In her explanation of theses cases, speaker and hearer token a concept that they create on-line, as a response to the specifics of the occasion. Now, if that on-line created concept is, as Carston labels it, atomic, then she is not in a position to use the productivity argument is support of her view. Recall that productivity is supposed to reconcile our ability to think an infinite number of thoughts with our having a finite mind. The traditional answer is that we can create an infinite number of complex concepts. However, Carston's explanation is a different one: in the underdeterminacy scenarios we think new thoughts because we create new simple concepts. Creativity, then, has to do with an ability to create indefinitely many concepts that fit the indefinitely many situations we encounter.

As a conclusion, Fodor hasn't offered any argument to the effect that mental representations are truth-conditionally compositional, and not merely meaning compositional—Meaning Compositionality being compatible with Type-Underdeterminacy. On the other hand, if *ad hoc* concepts are really atomic, ad not a combination of pre-existing concepts[21] , then Carston cannot use the productivity argument.

3.4.2 Equivocation

It has been argued (Fodor (2001) and (2003), Carston (2002) and (2008a)) that ambiguity and equivocation cannot occur at the mental level, since it is thoughts that disambiguate sentences. According to this line of reasoning, it makes no sense to take thoughts to be ambiguous themselves: if they were, there would be nothing that could disambiguate them. The argument can be reconstructed as a *reductio*. Some English

[21] As I noted in sec. 2, Carston's description of the process of creation of *ad hoc* concepts raises some doubts that they are not a combination of pre-existent representations. If they are, in this sense, complex, then she can avail herself of the productivity argument. However, a different worry would arise, as I argue in sec. 5.

words, as for example 'bank', have distinct meanings (financial institution, side of the river). Suppose that the same happens at the mental level, that is, suppose that English speakers only have one concept for BANK. Now, if that is so, then thoughts about financial institutions and thoughts about the side of a river are indistinguishable. But, clearly, thoughts about financial institutions and thoughts about the side of a river are easily distinguishable. So it is not the case that English speakers have one ambiguous concept BANK. Rather, they must have two different concepts, corresponding to the two different meanings of the word 'bank'. Moreover, these two concepts are the key to resolving the ambiguity with the English word 'bank'.

If this argument works for ambiguity, then it also works for other forms of equivocation such as the ones involved in (1)-(3). Again, if speakers only have one concept corresponding to the different senses of 'black', 'bachelor' and 'happy', then thoughts about a kettle being superficially black would be indistinguishable from thoughts about a kettle being originally black; thoughts about bachelors capable of long-term commitment would be indistinguishable from thoughts about not legally married men; and we could not discriminate thoughts about the different degrees and kinds of positive emotions that 'happy' seems to cover. The idea, in short, is that whereas language can, and often does, equivocate, thought cannot. Thought resolves equivocation. As Fodor puts it:

> [W]hereas the content of a sentence may be inexplicit with respect to the content of the thought it expresses, a thought can't be inexplicit with respect to its own content; there can't be more—or less—to a thought than there is to its content because a thought just is its content. ((Fodor, 2001, p. 14); quoted in (Carston, 2008a, p. 339)).

Fodor and Carston endorse a representational theory of mental content according to which 'having a thought with a particular content P involves the occurrence (the mental 'tokening', as it is often put) of a sentence of the language of thought (Mentalese) that means that P.' (Carston, 2002, p. 74). In this framework, Mentalese sentences (and in particular, their types) disambiguate or resolve the equivocations of natural language sentences. Whenever there are two distinguishable thoughts, there must be two different Mentalese sentence-types being tokened.

Now, as Clapp (2012b) has argued, the thesis that a thought just is its content does not entail the thesis that mental representations, conceived as Mentalese sentences (and specifically, sentence-types), have context-independent truth-conditions. 'Thought' can mean thought-content or thought-vehicle (see also (Recanati, 2007)). In the content

sense, and assuming content is conceived as truth-conditions, it is a conceptual truth that a thought just is its truth-conditional content. However, Mentalese sentences are thought-vehicles. Thought-vehicles are not truth-conditional content themselves—at most, they determine truth-conditional content, or are the bearers of truth-conditional content. Thus seen, the argument from equivocation relies on the premise that the only things that can disambiguate between two uses of an ambiguous or equivocal sentence—the only thing that can distinguish their truth-conditions—are thought-vehicles, and more specifically, types of thought-vehicles. Hence, the argument concludes that, whenever uses of a sentence involving the word 'bank' ('black', 'bachelor' and 'happy') differ in truth-conditions, these uses must be expressions of different Mentalese sentence-types.

Making the premise explicit and applied to equivocation, the argument is the following:

1. Two uses of an equivocal sentence can (and often do) differ in truth-conditions.
2. The only thing that can resolve the equivocation (=account for the difference in truth-conditions) is a mental representation-type which is conceived as non-equivocal.

C. Therefore, there must be non-equivocal mental representations.

However, premise 2 is false, because mental representation-types can be equivocal, yet the tokening context can resolve the equivocation. To show this, I will turn to various examples in which the same mental representation-type gives rise to different truth-conditions in different contexts of tokening. Let me start with an argument by Recanati (2007) concerning demonstratives and progressively move to general terms.

Recanati (2007, pp. 33-34) asks whether it is possible to find thought-vehicles with different truth-conditions at different contexts. His affirmative answer is supported by the following example. Imagine a subject who entertains the thought corresponding to 'This man is happy' while looking at a certain man—Bob. His thought is true if and only if Bob is happy. Now, had the context been different, the truth-conditions of his thought-vehicle could have been different. Suppose that it is Bill, not Bob, the man that the subject is looking at. In this case, the truth-conditions of his thought-vehicle would depend on Bill's properties, not Bob's. It seems that the same thought-vehicle can have different truth-conditional content. If the fact that the visual information that the

subject has is different is considered a problem, we can imagine that Bob and Bill are visually indistinguishable for the subject (they are twins).

This example shows that the tokening context can affect the truth-conditional content of a mental representation-type. Let's consider another example. Suppose that some mental representations, those concerning the present, are time neutral. Take the thought corresponding to 'Diana is happy'. To say that this mental representation is time neutral is to say that it contains no element encoding the time at which Diana is supposed to be happy. Would tokens of this mental representation equivocate as to when Diana is supposed to be happy? Not necessarily. The time can be provided by the tokening context. A tokening of the mental representation corresponding to Diana is happy at time t will be true if and only if Diana is happy at t, whereas a tokening of the mental representation corresponding to Diana is happy at time t* will be true if and only if Diana is happy at t*. Not everything that affects truth-conditional content needs to be encoded in the mental representation[22] .

Can the tokening context resolve the equivocation of predicates and general terms? A potential problem comes from the fact that, whereas time and place are automatically given by the tokening context, the same is not true of the satisfaction conditions of a predicate or a general term. In the previous example, time can be said to only take one value—something like now, let's say. When we move to the cases of underdeterminacy discussed at the beginning of this paper things get more complicated. The predicate 'is black' can get different satisfaction conditions—roughly, black in some specific relevant part, originally black, black all over, black in such-and-such observation conditions, etc. The act of tokening the mental representation corresponding to 'The kettle is black', of itself, is not going to decide what counts as 'black'. I think, however, that there is something else.

Typically, utterances take place in the context of a conversation. Following Lewis (1979), we can think about the context as the conversational score. The conversational score can be seen a specification of the state of the conversation. What is interesting now is that the interpretation of a given utterance can depend on the conversational score. One of the examples Lewis discusses involves the verbs 'coming' and 'going'. These expressions require a point of reference. Now, this point of reference can be part of the conversational score, for instance because of being given by a previous utterance:

[22] Perry (1986) argues that something similar happen with the location needed to get the truth-value of an utterance of 'It's raining'. Although the location is not, as he puts it, articulated in the sentence, it is necessary in order to get a truth-value. His proposal is that it is provided by the context.

> One way to fix the point of reference at the beginning of a narrative, or to shift it later, is by means of a sentence that describes the direction of some movement both with respect to the point of reference and in some other way. "The beggars are coming to town" requires for its acceptability, and perhaps even for its truth, that the point of reference be in town. Else the beggars' townwards movement is not properly called "coming". This sentence can be used to fix or to shift the point of reference. When it is said, straightaway the point of reference is in town where it is required to be. Thereafter, unless something is done to shift it elsewhere, coming is movement toward town and going is movement away. If we are told that when the soldiers came the beggars went, we know who ended up in town and who did not. (Lewis, 1979, p. 351).

Something similar could be going on in a conversation in which 'That kettle is black' is uttered. Imagine a pair of friends who have just moved together. In their new flat, there is no kettle and so one of them says 'We need a kettle, let's go buy one'. On their way to the shop, they talk about what kind of kettle they will buy. They say things like 'I don't like aluminium kettles, they are so ugly', and 'Let's buy something modern and elegant'. When they arrive to the shop and one of them says 'Look, that kettle is black' the previous conversation has made it clear that the colour predicate applies to the apparent surface of the kettle. The dialogue is enough for us to see what kind of kettles count as black (nor burnt ones, for example, not ones that look black on black and white photographs but not in normal conditions). Had the conversation been different, the satisfaction-conditions of the predicate could have been different.

Thought episodes also take place in context. Instead of a conversational score, we have a mental score. Whether they are chatting or they are independently thinking about the kettle they need, it is plausible to think that the two friends consider the need to buy a kettle, think about the kinds of kettles they like, etc. Suppose, now, that the mental representation they token at the shop includes the lexically encoded concept BLACK and not an *ad hoc* concept. Is their thought distinguishable from one in which only originally black things count as black? It is, for the mental score determines that what is at stake is the aspect of the kettle. The same mental representation type, tokened in two different thought contexts, can express different truth-conditions.

In conclusion, the indispensability arguments fail to secure the claim that mental representations cannot be type-underdetermined. Alternative approaches are tenable.

3.5 Persistent underdeterminacy

In the previous section I have considered and rejected two arguments in support of the claim that mental representations do not suffer from Type-Underdeterminacy. I will now argue that mental representations involving *ad hoc* concepts should be expected to be type-underdetermined, as natural language sentences.

The reason has to do with the source of Type-Underdeterminacy. Let us go back to linguistic Type-Underdeterminacy. The linguistic meaning of sentences (1)-(3) type-underdetermines the truth-conditions expressed by its tokens. It might seem that the reason why these sentences suffer from Type-Underdeterminacy lies in the fact that they do not make everything explicit. For example, (1) doesn't explicitly tell us where, in what spots, the kettle is supposed to be black; (2) doesn't specify the sense of 'happy' in which Anne is happy; and (3) doesn't make explicit what kind of bachelor the speaker wants to meet or who counts as a bachelor. However, as soon as we try to find other sentences that explicitly state where the kettle is supposed to be 'black', etc., we realize that, again, the new sentences do not make everything explicit either[23] . We can replace (1)-(3) by more complete versions:

(1*) The kettle is superficially black.
(2*) Anne has a stable feeling of happiness.
(3*) I want to meet some bachelors interested in long-term relations with women.

These sentences have solved some doubts, but not others. (1*) doesn't specify how much of the surface needs to be black for the utterance to be true, or what exactly counts as the surface (what about black plastic glued to the original surface?). (2*) doesn't specify how long a feeling needs to last in order for it to count as 'stable'. And the same goes for (3*), which doesn't give any clue as to what is the threshold for a relation to count as a 'long-term' one. So they are, in this sense, equivalent to (1)-(3): they also exhibit Type-Underdeterminacy.

The question now is whether it is possible to come up with a more complex sentence-type that leaves no doubts open. We, normal speakers, might be unable, for our cognitive capacities are limited. The question, however, is whether, were our capacities to be improved, we could come up with a sentence that encodes the content that ut-

[23] See (Searle, 1978, 1980), (Travis, 1997), (Cappelen and Lepore, 2005) and Recanati (Recanati, 2010) for similar arguments.

terances of (1)-(3) convey. There is a reason why we should give a negative answer. If we go to (1*)-(3*), we will see that the reason why these sentences also exhibit Type-Underdeterminacy is that the new expressions in the sentence can be understood differently on different occasions. For example, the threshold for something to count as 'stable' can vary across contexts, depending on what is at stake. To be sure, a speaker could use (2*) and solve some doubts concerning a use of (2) ('What do you mean 'Anne is happy'? Has she won the lottery?'). But this doesn't mean that he has thereby found a sentence that is free of Type-Underdeterminacy. The same happens with 'superficially'. For example, now it is open how much of the surface needs to be black in order for an object to be 'superficially black', or what counts as the object's surface. The added material comes with different possible interpretations, with different satisfaction conditions in different contexts. Complexifying a sentence is not the route to overcome Type-Underdeterminacy. Let us suppose that the mental representation-types corresponding to sentences (1)-(3) are composed of the concepts BLACK, HAPPY, BACHELORS, etc. Since adding linguistic expressions is not a way to eliminate the Type-Underdeterminacy, adding concepts to mental representation like SURFACE, STABLE or LONG-TERM will not be a way of getting mental representations that are free of Type-Underdeterminacy. Combinations of concepts-types corresponding to words are as type-underdetermined as the combinations of words themselves are. The problem is that it hasn't been shown that there are ways of conceiving the mental representation-types corresponding to utterances of (1)-(3) other than combining the concepts we already have—those corresponding to 'bachelor', 'black', 'happy', 'surface', 'long-term', 'stable'. As they are described, *ad hoc* concepts are no exception.

As I said, *ad hoc* concepts are supposed to be non-linguistically encoded atomic concepts. They don't correspond to words, and they are not combinations of concepts. So they might seem to escape this problem. However, they are created by activating some information belonging to the encyclopaedic entry of linguistically encoded concepts, and this information sounds very similar to the kind of information that we can encode in natural language and has different satisfaction conditions in different occasions. Concerning (2), Carston writes:

> Suppose, as above, that the context is one in which the addressee knows that the speaker, who is a woman, wants to get married and have children. Having accessed the lexical concept BACHELOR, which makes available its associated logical and encyclopaedic information, he uses a subset of this information to construct a more specific concept BACHELOR*, which is relevant (that is, gives

> rise to cognitive effects) in the context. The encyclopaedic entry might well contain information about certain sorts of bachelor, the irresponsible, fun-loving, forever-young-and-free sort, the elderly, solitary, misogynous sort, and those who are youngish, heterosexual, and capable of long-term commitment, i.e. eligible for marriage. Given the hearer's alertness to the speaker's marital interest, it is probable that information about this third sort of bachelor will be more highly activated than that about either of the others, so it will be accessed first, together with the logical entry, and used to construct the *ad hoc* concept. Provided this gives rise to a satisfactory range of cognitive effects, it is retained as the intended interpretation. In different contexts, other narrowings might be effected, yielding concepts which denote different subsets of the category of unmarried adult males. (Carston, 2002, p. 326)

In order to create the *ad hoc* concept, some information included in the encyclopaedic entry of BACHELOR is activated. Now, this information doesn't seem to be essentially different from the information we can encode in natural language. As I mentioned, 'being capable of long-term commitment' can be interpreted in different ways in different occasions, and the same happens with 'solitary', 'irresponsible', and so on. This information, per se, is no different from the information contained in 'bachelor'. Just as the latter can be understood in different ways, so can the former. Even if *ad hoc* concepts are supposed to be atomic, they are created by selecting features, so, in the end, they are not very different from a combination of concepts. And combinations of concepts encodable in natural language (like long-term commitment, solitary, irresponsible) are still type-underdeterminate.

The problem lies on the assumption that some structured representational items beyond language are free of Type-Underdeterminacy. We have seen that, in language, given a general expression F and a referential expression a, it is not automatically decided, by the semantics of F alone, whether a satisfies F. In general, some doubts concerning how to apply F must be solved. Concepts, as common nouns or adjectives, also have a general content. So it is plausible to think that the mental context is doing some work, just like the conversational context does. This is specially pressing if, as Fodor and Carston do, we conceive mental representations as sentences of a language[24] .

[24] Wittgenstein ([1953] (2009) 139-141) considers whether understanding a word (in his example, 'cube') can be explained as having a picture before one's mind (an image of a cube) and asks the question whether it is correct to apply the word 'cube' to a prism. He argues that the picture, the mental image, can be both made to fit and not to fit this particular use of the word 'cube', depending on the method of projection one applies. Now, if the method of projection is made part of the image that comes before one's mind,

At this point, a proponent of the mixed view might have recourse to an externalist theory of mental content . In order to counter the previous argument, he could claim that the content of the tokened *ad hoc* concept-type (a concept-type that is created on-line) is determined by the world itself. Imagine Peter, a friend of Anne who, while chatting with her, tokens the mental representation that could be expressed by saying 'Anne is happy'. According to this theory, the content of his representation will depend on how Anne happens to be. If, at the time of Peter's tokening, she experiences a stable positive emotion, then that will be the content of Peter's thought[25] . Thus, an externalist could claim that mental representations have determinate content, even though these contents might not be fully transparent to us and despite the fact that we cannot encode them in natural language.

I think that this externalist position faces two important worries. First, there are cases where the world itself does not help resolve the equivocation. Belleri (2014b) imagines such a case. I adapt from her example. The predicate 'is black' can be used to describe, among others, the original aspect of an object or its apparent surface. In this second sense, a burnt or painted kettle can be rightly described as 'black'. Now, we can imagine a kettle that is both originally black and painted black, or burnt, or very dirty. Imagine, further, someone thinking, of that kettle, what could be expressed by the sentence 'The kettle is black'. Is his mental representation about the original colour or about the apparent surface? Here, the world doesn't help decide. If the content of the *ad hoc* BLACK* is determined by the world itself, then it doesn't resolve the equivocation between the possible interpretations of 'black'. However, in many cases, the context will resolve the equivocation—for example, in cases where the person who tokens the mental representation is checking whether his flatmate has properly washed the kettle after using it.

More importantly, in the underdeterminacy scenarios the relevant understanding of the predicate doesn't seem to be determined by the world. Rather, it depends on what is at stake—on what is relevant for the purposes of the conversation or the thought

then the same problem arises again: how am I to know how to apply the image of the method of projection? Even if one abandons the idea that understanding a word has to do with having a mental image, the meaning of a general word seems to involve some kind of descriptive content. The method of projection problem arises for this descriptive content as well.

[25] It is not at all clear that this kind of externalism will even be applicable to all cases. Although it might sound appealing for cases of singular representation, it is not easy to see how could the world resolve the equivocation involved in (3), given that a mental representation corresponding to this sentence will not be about any particular individual.

episode. Imagine someone who wants to buy a modern, elegant kettle. Looking at a kettle, he tokens a mental representation that could be expressed by saying 'The kettle is black'. Let us imagine now that the relevant kettle is only black on its interior—a part of the kettle the subject cannot perceive—and that he was misperceiving the object because, let's say, the kettle is placed behind a dark glass. Should we say that his mental representation is such that the concept he tokens is one that is true of black-in-the-interior objects? I think this would be an odd description of the situation. Rather, given his purposes, it would be more reasonable to take his thought to be true if and only if the relevant kettle is superficially black. Typically, it is mental score, rather than the world itself, that resolves equivocation. But if it is conceptual material that is meant to fix the ensuing underdeterminacy, then the problems described in this section arise.

3.6 Partiality

In this section I will introduce a second notion of underdeterminacy and argue that mental representations are also underdetermined in this second sense[26] .

So far underdeterminacy has been understood as concerning types. However, it can also be understood as concerning tokens of structured representations:

> *Token-Underdeterminacy:* A token of a structured representational item S is token-underdetermined if and only if for some possible states of affairs its truth-value is indeterminate (i.e., if and only if it determines a partial function from possible worlds to truth-values).

Token-Underdeterminacy is stronger than Type-Underdeterminacy. Suppose that a sentence-type S has truth-value T at state of affairs SA. Presumably, if that is so, then tokens of S will all be T at SA. The semantic properties of the type are inherited by the token. Hence, if the truth-value of a token of S is indeterminate, the truth-value of the type S must be indeterminate as well. Otherwise, the truth-value of a token of

[26] In the previous sections I have addressed the problem of underdeterminacy as it arises in the literature concerning language and thought (mainly, in Carston's work). This section addresses a separate problem concerning a different notion of underdeterminacy.

S would not be indeterminate. By contrast, a representational system could be type-underdetermined without being token-underdetermined. This will be the case whenever the tokens determine truth-values for all possible states of affairs.

Token-Underdeterminacy can be motivated by appealing to what Waismann called the 'open texture' of our words. Here is a case Wittgenstein discusses ([1953] (Wittgenstein, 2009) 80). Imagine that I invite a friend over and say "There is a chair over there". When he approaches it, it disappears, but seconds later it appears again. Is my utterance true in the scenario described? The conversation might have fixed a relevant understanding of the word 'chair' (for example, one that excludes doll chairs and baby chairs) but still fail to determine whether the thing that appears and disappears counts as a 'chair'.

Token-Underdeterminacy becomes interesting when we realize that something similar to the case described by Wittgenstein can happen in more mundane scenarios. Imagine that instead of a chair that disappears I have some boxes that I usually use as chairs, or some pieces of modern furniture. There can be conversations about chairs that would leave it open whether these objects count as 'chairs'—conversations taking place at buildings where there are no such objects. Nonetheless, in other conversations it will be determinate whether these objects count as 'chairs' (saliently, in conversations at my place). Given a state of affairs (an apartment with some boxes and pieces of modern furniture), it can be the case that a token T_1 of 'There's a chair over there' is true at this state of affairs, whereas a second token T_2 is false and and the truth-value of a third token T_3 is indeterminate (where T_1, T_2 and T_3 occur in different conversations but the state of the relevant objects and the value of indexicals remain constant). Different tokens of the same sentence can leave the truth-value of different state of affairs indeterminate.

As I mentioned in the previous section, natural language sentences do not make everything explicit. In this sense, they can be thought of as leaving some questions open[27] such as: where does a kettle need to be black in order to count as 'black'? Or, for how long does one have to have a positive emotion in order to count as 'happy'? Usually, the context of use resolves these questions. For example, if everything goes well, previous conversation will fix whether 'black' is supposed to apply to the whole surface or a specific part of it, or whether 'happy' refers to a rather stable emotion or not. Depending on what is at stake, the question will be answered in one way or another. However, there can be questions that are left unanswered in a conversation. For example, all that is said in a conversation might leave unanswered questions about abnormal illumina-

[27] See (Travis, 1989) for this kind of approach to underdeterminacy.

tion conditions or about for what period of time the superficial colour is supposed to last (does a kettle that has been painted black with a painting that will only last for a month—because it progressively disappears as it is washed, let's say—count as 'black' in a conversation about superficially black kettles?). Whenever some questions remain unanswered there might be possible states of affairs for which the utterance does not determine a truth-value. As I will argue, the same goes for mental representations.

Here is an example. Imagine someone asking for a red pen to mark some essays. He utters the sentence 'I need a red pen. I've some essays to mark'. Now, 'red pen' can be used to refer to pens with red ink or to superficially red pens, etc. To what does it refer in this case? The second sentence solves this doubt. Because of it, it might seem plausible to assume that the mental representation that the speaker was tokening contained information about the ink. However, there are other things that the sentence 'I need a red pen' does not make explicit. For example, it does not settle whether the ink must look red on recycled paper, or on yellow paper. Imagine that there are no recycled papers available in the building where the conversation takes place and that, because of this, the speaker never prints the essays that his students send him in recycled paper. His mental representation (token) will very likely simply leave that doubt open. As a result, it will not determine a truth-value for any possible state of affairs. In what follows I develop this point in some more detail.

Here is a well-known example by Travis, often used as an argument for the Type-Underdeterminacy of natural language (it is similar to the black kettle example Carston mentions):

> Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, 'That's better. The leaves are green now.' She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those.' But now Pia speaks falsehood. (Travis, 1997, p. 111)

The example discloses some features that might matter to the application conditions of the colour predicate, such as the origin of the colour and the part of the object that exhibits it. It might seem plausible to say, about this example, that the mental representation that Pia tokens is (at least in the first occasion) one that specifies that the leaves are superficially, or visibly, green. Now, there are other examples disclosing other possible doubts concerning the application of the colour term. Let us call this kind of example a Travis case. There are other possible Travis cases. Here is one (based on some remarks by Travis (1997)):

> Jean's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, he paints them. He has no green paint, and he enjoys experimenting, so he uses blue and yellow paint. He paints small yellow and blue dots. The paints don't get mixed. However, at a certain distance (0.5m, let's say), the leaves look green. Returning, he reports, 'That's better. The leaves are green now.' He speaks truth. A photographer friend then phones, seeking green leaves for photography of a single leaf, taken at a close distance. 'The leaves (on my tree) are green,' Jean says. 'You can have those.' But now Jean speaks falsehood.

This example discloses something else that the sentence 'The leaves are green' doesn't make explicit: the distance at which the object is observed might matter. If Pia's mental representation is free of Token-Underdeterminacy, it must determine whether Jean's leaves would satisfy her use of 'are green'. However, this might very well be indeterminate. Pia need not have considered the possibility that at a certain distance objects might look different (as typically happens with printed colours). In Jean's scenario, this, however, is relevant, and this doubt can be considered as solved by the context. But this does not imply that the doubt is solved in every context in which the predicate 'is green' is used. Let us go back to 'bachelor'. The meaning of this word gives rise to the following questions: is someone married by some religious ritual but not legally a bachelor? Is someone in the middle of a divorce a bachelor? Is someone not married but who has a law-partner a bachelor? Is a man who married two women simultaneously in a ceremony taking place in his home country but that now lives in some European country (where polygamy is not permitted) a bachelor? These questions will be answered differently in different cases, and in some cases, some of them will not be answered at all. Imagine the following scenarios. Case 1: Anne is interested on the relation between marital status and the amount of taxes paid. She suspects that being married has not entailed any economic advantage for her. Thinking that it would be interesting to ask for personal experiences, she tells a friend: 'I want to meet some bachelors'. In this setting, it is plausible to take Anne's mental representation to token a narrower concept than BACHELOR, roughly equivalent to non-legally-married-men. However, it is not plausible to take her mental representation to solve, for example, the question about men going through a process of divorce or about men who got married in countries with marital laws radically different from the ones in the country were she lives, even though some of these men happen to live in the same country as she does. These questions simply didn't occur to her. Case 2: Tom wants to meet someone special and begin a long-term relationship, but most of the men he knows are married or have a part-

ner. He says to a friend: 'I want to meet some bachelors'. Again, it is plausible to take his mental representation to solve some questions, but not all of them. For example, his mental representation will very likely exclude men who are married by some ritual regardless of whether they are legally married, and it will also exclude men married in other countries even though they don't count as legally married in the country where he lives. The point is that the concept he tokens need not decide whether a man in the middle of a divorce satisfies his use of 'bachelor', or what happens with men who got married in order to get a European nationality, or with men that got divorced but are still involved in a relationship with their ex-husbands. As a conclusion, mental representations can be Token-Underdetermined. It seems safe to hold that the truth-value of a mental representation-token at a state of affairs that is not relevant for the tokening context can simply be indeterminate [28] .

3.7 What about propositions?

In this section I will consider a different version of what I have called 'Pragmatic Propositionalism'.I have argued against the claim that mental representations are free of Type-Underdeterminacy. However, having recourse to mental representations is not the only possible implementation of Pragmatic Propositionalism. Here is a different version of the mixed view:

> The linguistic meaning of sentence type-underdetermines the truth-conditions of the tokens, and the linguistic meaning of predicate type-underdetermines the satisfaction conditions of the tokens. This is so because the linguistic meaning of a sentence is compatible with (literally) expressing a variety of (structured) propositions and the linguistic meaning of a predicate is compatible with expressing a variety of properties. On an occasion of use, a predicate P expresses a property F that, being a property, establishes a partition among objects (perhaps of a given domain): it divides all objects into two sets, the set of those objects that have the property and the set of

[28] Belleri (2014b) defends a similar view according to which the notion of determinacy for the contents of our thoughts is context-relative.

those objects that lack it. Thus, the predicate token gets satisfaction conditions.

Perry and Williamson suggest this view in the following passages:

> I'm now looking at the Sprite can on the table, and I just said, out loud, "That is empty." [...]Did I mean "contains no Sprite" or "contains nothing at all"? The word "empty" can be used to get at a number of properties, depending on what one intends to say with it, and it can also be used less definitely. (?, p. 197)

> 'Mine', with its usual meaning, does not always refer to the property of being mine; why should 'round', with its usual meaning always refer to the property of being round? With this correction, Travis' argument can be reconstructed as leading to the less radical conclusion that meaning alone ubiquitously underdetermines what is said. (Williamson, p.1998, p. 10 n, quoted in (Travis, 2008a)).

The idea Perry and Williamson are putting forward is that, when we apply the predicate 'is empty' ('is round') to some object, we are not always attributing the same property to it. In some contexts we will attribute the property Empty_1 (Round_1), in others Empty_2 (Round_2), etc. Travis' view, by contrast, is that in all these cases we describe the object as round (no quotation marks), where being round is an occasion-sensitive property, that is, a property that an object might have in some occasion and lack in others, without undergoing any change. If the term 'property' is defined as a partition among objects, or a function from objects to truth-values, or something that determines extension/intension, then Travis is wrong by definition. However, if a more substantial notion of property is assumed, then Pragmatic Propositionalism might be problematic.

Travis (1989, pp. 23-24) and Borg (2012) argue against this approach. Travis notes that it can be objected to occasion-sensitivity that 'is green' or 'weighs 79 kilos' refer not to a property but to a family of properties. Let us suppose that this is the case and that, even if we cannot talk about the property that 'is green' or 'weighs 79 kilos' refer to, for there are a variety of them, we can talk about the property that the words 'is green' or 'weighs 79 kilos' refer to on a particular occasion, and assure a level of occasion-insensitive content. Travis argues that the assumption that our words, in use, express an occasion-insensitive property in problematic because of the same reasons why the assumption that 'is green' expresses an occasion-insensitive property is problematic. The problem is that in deciding that 'weighs 79 kilos' applies to Hugo in such and such conditions, or that 'There's milk in the fridge' is true of a certain fridge in such and such

conditions, we have solved some doubts but not all of them. So suppose, for the sake of the argument, that there is a property that 'is green' and 'weighs 79 kilos' expressed in the scenario described by Travis. Call them P and Q. Now we can ask, would 'There's P in the fridge' be true if there was (only) a slice of cheese? Or synthetic milk? Or a bottle with coagulated milk? Would 'Hugo Qs' be true if the earth's gravitational force were to be halved overnight? We cannot answer these questions, for they raise doubts we didn't solve when we settled the issue about the bridge or about cleaning the fridge. So, Travis concludes, it turns out that these properties, in turn, leave some questions open—which, for him, means that they are occasion-sensitive after all[29] .

Relying on Travis' argument, Borg claims that the proponent of the mixed view is under pressure to specify the occasion-insensitive contents that, according to him, are expressed by our utterances. The claim is that these contents are recovered via pragmatic interpretation, not that they cannot be had. However, it is reasonable to think that for any specification of a given understanding that can be given, or for any sense we are able to somehow select, new Travis cases can be created and, as Borg says, 'occasionalism at one remove is still occasionalism'.

I think that those who endorse Pragmatic Propositionalism (structured propositions) can counter this argument, for they can reject the identification of the properties Travis specifies (what he calls P and Q) with the contents they postulate. As soon as we name or specify properties, we are dealing with linguistic meaning, and these contents admittedly type-underdetermine the property expressed. Thus, a proponent of the mixed view might counter the Travis-Borg argument by carefully distinguishing metaphysical determination from what we can call 'epistemology of content'. Thus, it can be argued that utterances express structured proposition that are free of Type-Underdeterminacy but that these are not fully transparent to us—we don't know the answer to the questions Travis asks, for propositions are not fully transparent, but this doesn't mean that there are no answers.

However, if mental representations are not free of Type-Underdeterminacy, then the idea that utterances express propositions that are free of Type-Underdeterminacy looses its attractive—this sort of propositions doesn't seem to be playing any role in our cognitive lives, so why should we use them in our linguistic theories? It is not clear what the role of those properties would be, for they do not correspond to our mental representations. Hence, even if one can have a notion of 'property' such that properties

[29] I think that Travis is here aiming at something similar to Token-Underdeterminacy. This suggests that occasion-sensitivity would be, for him, a combination of Type- and Token-Underdeterminacy.

are free of Type-Underdeetrminacy, it is not at all clear that we should have recourse to this kind of properties when studying natural language.

3.8 Conclusions

I have argued against the claim that mental representations are underdeterminacy-free. I have distinguished two notions of underdeterminacy—Type-Underdeterminacy and Token-Underdeterminacy. On the one hand, I have considered and rejected two arguments in support of the claim that mental representations must be free of Type-Underdeterminacy: the productivity argument and the argument from equivocation. Against them I have argued, first, that the productivity argument only establishes that natural language and mental representations must be meaning-compositional, and this is compatible with Type-Underdeterminacy. Moreover, I have noted that views that have recourse to online concept creation cannot avail themselves of the productivity argument. Second, I have argued that the argument from equivocation relies on a false premise—namely, the premise that mental representation-types are non-equivocal and, as such, they are suitable to resolve the equivocations involved in an utterance of a natural language sentence. The tokening context can also resolve the equivocation. As a result, proponents of the mixed view have failed to provide good reasons in support of their view. After this, I have argued that *ad hoc* concepts should not be expected to behave differently than lexically encoded concepts. Hence, mental representations, even when they involve *ad hoc*, non-lexically encoded concepts should be expected to be Type-Underdetermined.

On the other hand, I have argued that there are reasons to take mental representations to be token-underdetermined. I have done so by presenting some cases where it seems reasonable to conclude that the truth-value of the mental representation-token is indeterminate at some possible states of affairs. The resulting picture—that mental representations are token-underdetermined—is admittedly a radical one. In particular, it is more radical than approaches that take mental representations to be type-underdetermined. The thesis that mental representations are type-underdetermined motivates approaches according to which the situation where a representation is tokened plays a crucial role in resolving equivocations and, consequently, in determining

the truth-conditional content of the token . This is compatible with the tokens being free of Token-Underdeterminacy. By contrast, the partiality argument motivates an approach according to which tokens exhibit this second kind of underdeterminacy Although at the tokening context the questions concerning application conditions that are relevant for whatever is at stake will typically be solved, other questions that are not relevant can remain unanswered.

I have further argued that, if mental representations are Type-Underdetermined, as I think they are, having recourse to structured propositions that are free of Type-Underdeterminacy doesn't seem like a good option.

My aim in this chapter was to cast doubt on the assumption that there is another system (saliently, Mentalese) that, unlike natural language, is free of Type-Underdeterminacy and that can be used in order to confer truth-conditional content to our utterances. If my arguments are sounds, there are reasons not to understand truth-conditional pragmatics as the view that natural language sentences type-underdetermine the (structured) proposition or mental representation expressed in a use of sentence, with those (structured) propositions or mental representations being underdeterminacy-free. Instead, in what follows I will put forward a view of utterance content that can be classed as relativism or situationalism.

Chapter 4

Utterance content

The aim of this chapter is to put forward a notion of utterance content that is in accordance with the phenomenon of occasion-sensitivity. I will adopt Recanati's relativist framework and model utterance content as an Austinian proposition. My notion of utterance content has three features. First, Austinian propositions are modelled as <lekton, activity> pairs, thus substituting Recanati's use of situations with activities. Second, this framework leaves room for Token-Underdeterminacy. The reason is that activities constrain the application conditions of words but do not determine a truth-value for any possible states of affairs. Third, in order to account for intercontextual content-sharing, and in particular for speech reports, my framework admits of Austinian propositions with different granularity. I will finish by arguing that this notion of content escapes the non-objectivity threat introduced in chap. 1.

4.1 Where are we?

In the previous chapters I have countered several arguments intended to deactivate the pragmatist challenge raised by the phenomenon of occasion-sensitivity. In chap. 2 I have addressed and rejected several arguments to the effect that the sentences used in Travis cases are (per se) truth-evaluable. Moreover, I have argued that the theories that posit some kind of context-sensitivity in colour predicates either are insufficient to deal with the possibilities of variation that these predicates exhibit or have recourse to prag-

matic interpretation. In chap. 3 I have argued that the arguments used for supporting the claim that mental representations are not subject to the kind of underdeterminacy natural language is subject to are not sound. As a consequence, I think that we should abandon both Semantic and Pragmatic Propositionalism. Instead of relying on there being mental representations with occasion-insensitive truth-conditions or having recourse to occasion-insensitive structured propositions or mental representations[1] , I will present here a new way of conceiving truth-conditional content inspired on Recanati's relativism.

We can think of the role of context as not being that of identifying an enriched proposition. When we engage in a conversation and interpret an utterance the information we have includes not only the sentence uttered (with its linguistic meaning), but also the surroundings of the utterance. For instance, we typically have information about what the conversation is about, which objects are salient, etc. We can model the content of an utterance as a pair including the conventional meaning of the sentence uttered plus those surroundings. My suggestion is that we need a dual notion of content, capturing both the meaning of the sentence (for example, that it describes a certain object as having a certain property) and the activity in which it is used.

On this view, truth-evaluable content is content seen as embedded in an activity. Let us use Wittgenstein's remarks on seeing-as and the duck-rabbit picture as an illustration[2] . In the duck-rabbit picture, we have some lines on a piece of paper. We can see those lines as a picture of a duck or as a picture of a rabbit. It wouldn't make sense to ask whether it *really* is a picture of a duck or a picture of a rabbit. The picture remains the same whether we see it as a duck or as a rabbit. This, however, doesn't prevent us from seeing the picture differently. We perceive different aspects, let's say. There is a difference in the visual experience. Something similar happens with linguistic meaning. Take an occasion-sensitive sentence as 'The leaves are green'. Different occasions of use, or different language games, are like different ways of seeing that sentence (it's meaning). We can see a colour description as describing the natural or the visible colour. When we embed it in an activity, the occasion triggers one way or another of seeing it. Travis' examples, with their deployment of a setting, create the conditions for a linguistic Gestalt switch to happen. Linguistic meaning can be perceived in different ways. Given a context, and assuming the context is cooperative enough, one of those ways is automatically

[1] In chapter 3 I mentioned the option of conceiving propositions directly as truth-conditions (or sets of possible worlds) without discussing it. I will come back to this in sec. 3.

[2] I merely intend this as a rough analogy. In my discussion, 'seeing-as' is metaphorical.

available to us.

Following the idea that the truth-value shift in Travis cases is due to the sentence being embedded in different activities, we can see occasionalism as a form of situationalism (following (Corazza, 2007)). Thought and talk is situated in the sense that it typically occurs at some time and place, in the context of a conversation and extralinguistic practices. Instead of taking the role of these surroundings to be that of identifying a fully determinate proposition, situationalism sees them as affecting the truth-value of the sentence or the mental representation[3] . Utterances typically occur in conversations where people have goals and purposes and are engaged in extralinguistic activities. These surroundings provide the key to interpretation–they give us the measure of what counts as 'green', for example. Corazza (2007) distinguishes contextualism and situationalism on the basis of their approaches to context-dependence. Contextualists claim that our utterances typically express enriched, or contextually adjusted, propositions. By contrast, situationalism has it that the proposition expressed by an utterance of a non-indexical sentence S ('There's beer in the fridge') is semantically determined (it is the proposition *that there's beer in the fridge*) but admit that the truth-value of this sentence is relative to the situation of use[4] .

I think that the lesson of Travis cases is that sentences are true or false only when embedded in an activity. That what we need is an activity is suggested by the examples themselves. In the examples we have people painting leaves, cleaning fridges, crossing bridges, etc. These activities bring in some constraints as to how to use words, yet they are extralinguistic. For example, the reason why we take Pia's first utterance to be true is that she is painting leaves. In this activity, what matters is how things look—how we want them to look, how they look as a result of our action, etc. By contrast, in Pia's second utterance the relevant activity is scientific research—and science is about natural properties. The activity itself comes with certain constraints. For example, that only natural properties matter.

My aim is not to present a formal framework that can deal with Travis cases but instead to put forward a notion of utterance content that is faithful to occasion-sensitivity. The question I will address: If language is occasion-sensitive, how should we think of the content of our utterances? I think that, given that the activity plays a crucial role in determining, for example, what counts as 'green', a good option is to model utterance content, following Recanati's Austinian propositions, as including a lekton and an ac-

[3] Travis (2000) speaks of representations as being situated.

[4] Situationalism can be seen as a form of relativism.

tivity. In the remaining sections of this chapter I develop his idea in more detail. I will argue that, in order to account for same-saying we need Austinian propositions with different granularity. Thus, I will endorse a form of multipropositionalism.

I will finish this section by summing up what I take to be the main insights of the pragmatist challenge regarding the notion of utterance content. Occasion-sensitivity motivates three claims. First, it motivates a relativist claim. Utterances of the same (non-indexical) sentence can have different (literal) truth-values at different occasions. Sentential truth is relative to the occasion of use. Thus, truth-value is determined by three elements: linguistic meaning, state of the world, and the specifics of the occasion[5] . Following MacFarlane (2007), this can be seen as concerning propositions and properties, and not only sentences—propositions can be seen as having different truth-values at different circumstances of evaluation, where circumstances include a 'count-as' parameter. Thus, the same property can be said to be expressed by all uses of 'green', etc. This property might determine different extensions at the same possible world, for the extension depends on the value of the 'count-as' parameter.

Second, it can be seen as motivating a contextualist claim, for one can hold that the two utterances depicted in a Travis case, in a sense, don't say the same. If I describe a certain object as being 'green' in two different occasions, I might not be saying the same thing twice (leaving time aside). For example, it might be the case that what I say in the first occasion is true if and only if the object is superficially green, whereas in the second the object must be originally green in order to count as 'green'. If that is so, I didn't describe things as being the same way[6] . To capture this, we need a notion of utterance content according to which what is said by the utterance transcends the conventional meaning of the sentence used.

Third, occasion-sensitivity motivates a context principle according to which the role of context is not mere identification of the mental representation (or structured proposition) expressed by an utterance[7] . Searle claims that we are incapable of understanding

[5] Travis is particularly clear about this in his (1996). Also Searle, when he talks about truth-conditions being relative to a background.

[6] This assumption is in place throughout (Travis, 2008a).

[7] The notion of context that I aim at is similar to Wittgensteinian language games. According to Conant (1998), '[Wittgenstein] seeks to generalize Frege's context-principle so that it applies not only to words (and their role within the context of a significant proposition) but to sentences (and their role within the context of circumstances of significant use, or – as Wittgenstein prefers to call them – language-games).' (Conant, 1998, p. 233) And later: 'Its point is that 'for a large class' of occasions of speaking there isn't anything which can properly count as asking the question 'What do the words [which have been spoken] mean?' apart from a simultaneous consideration of questions such as 'When was it

some well-formed sentences when we lack an appropriate background. He offers as an example the order 'Cut the sun'. Although we know the literal meanings of the expressions composing the sentence, we are unable to see what would count as obeying the order. For one thing, we cut different things in different ways, so there is no object-independent way of cutting that we could automatically apply to the sun. But more importantly, we lack some story that prepares us to the interpretation of the sentence (that tells us, for example, what we are doing and what is the aim of cutting the sun). We can't even decide among those potential ways of cutting without it. However, once we provide a story, that is, some context, understanding becomes easy.

In order to show how important the surroundings are, let me sketch two stories in which to embed the sentence 'Cut the sun!'.

> First scenario: a science fiction novel.
>
> There is an interstellar war going on, and one of the armies involved has developed a laser that cuts planets into two. Going for a more destructive use of the laser, the commander orders 'Cut the sun!'

> Second scenario: a non-literal use.
>
> Some children are cutting figures painted on a paper: a house, a dog, a child, a butterfly, a sun, etc. The teacher says to one of them: 'You are doing great, you already have the house and the dog. Now, cut the sun'.

That understanding requires an appropriate background is shown in that we can imagine the sentence as being used in a science-fiction novel or in a non-literal way, and easily understand it. One might reply that in some cases we understand questions or are able to obey orders without a previous story. I think that there are two reasons why sometimes we seem to understand an utterance in absence of a context. The first reason is that we understand something—namely, the linguistic meaning of the sentence used. The second reason is that we might have recourse to what I will call default contexts[8] . Suppose that 99% of cases in which we use the words 'France' and 'hexagonal' together the conversation is about countries resembling things (loose talk describing countries). Since we are equipped with this implicit knowledge, we can use it to interpret an out-of-the-blue utterance of 'France is hexagonal'.

---

said?', 'Where?', 'By whom?', 'To whom?', etc.' (Conant, 1998, p. 239)

[8] This idea will be explored in chap. 5.

In what follows I present a notion of content that is faithful to main in insights of occasion-sensitivity. This notion can be said to be dual, since it includes not only linguistic meaning but also an activity. In section 2 I present Recanati's discussion about Austinian propositions and put forward a new version of this notion. In section 3 I argue that language, as mental representations, can be token-underdetermined and that Austinian propositions can account for that. Section 4 is about how same-saying imposes an additional desideratum on the notion of utterance content. In order to meet this requirement, I adopt, in section 5, multipropositionalism. In section 6 I address the non-objectivity threat.

4.2 Towards a dual notion of content

Recanati (2006b; 2007), following Perry (1986), provides the means for defining a dual notion of content that captures the three claims motivated by occasion-sensitivity. Recanati's Moderate Relativism admits of two different kinds of content: the lekton and the Austinian proposition. The lekton is defined as the explicit, articulated content. Its truth-value is relative to situations of evaluation. By contrast, the Austinian proposition is the complete truth-conditional content, including the lekton plus a situation of evaluation. It has absolute truth-conditions. These two kinds of content correspond to two different kinds of linguistic items. Whereas the lekton is the content of the sentence, the Austinian proposition is the content of an utterance of the sentence. Thus, utterance content is not identical to sentence content. Take as an example the temporal proposition expressed by 'Socrates is sitting' (assuming time is not articulated in the sentence). The explicit content of this sentence (lekton) has different truth-values at different times. Now, an utterance of 'Socrates is sitting' at t has absolute truth-conditions: it is true if and only if Socrates is sitting at t.

Besides lekton and Austinian proposition, Recanati distinguishes another level of meaning: meaning of the sentence type or kaplanian character. His framework includes three levels of content[9] :

[9]Recanati uses the resulting framework to account not only for sentence/utterance content but also for the contents of thought. In what follows, I will focus on utterances.

1. The meaning of the sentence type (or kaplanian character): a function from contexts of utterance to lekta. This level is context-independent. Since I am dealing with predicates, not indexicals or demonstratives, I will leave this level aside.

2. The lekton: a relativized proposition (a function from situations to truth-values).

3. The Austinian proposition: the lekton and the contextually relevant situation. According to Recanati, it determines a classical proposition (in the sense that it has absolute truth-conditions).

The grounds for the distinction lekton-Austinian proposition is given by two principles:

> [Duality] To get a truth-value, we need a circumstance of evaluation as well as a content to evaluate. (As Austin puts it, 'It takes two to make a truth'.)
>
> [Distribution] The determinants of truth-value distribute over the two basic components truth-evaluation involves : content and circumstance. That is, a determinant of truth-value, e.g. a time, is either given as an ingredient of content or as an aspect of the circumstance of evaluation. (Recanati, 2007, pp. 33-34)

Duality is a well-accepted principle. It simply states the thesis that truth-value depends both on content and on how things are. Whether an utterance is true depends on what it means together with the state of the world. Distribution introduces a novel idea[10] : some of the elements that are needed to fix truth-value need not be given by content—they can be given by the circumstance. Recanati follows here Perry (1986). According to Perry, some of our utterances have truth-conditions only when some parameters not articulated in the uttered sentence itself are fixed. As an example, Perry uses the sentence 'It's raining'. In order for an utterance of this sentence to have a truth-value, not only a time must be determined, but also a place. For example, it might be raining now in Barcelona but not in Sevilla, so we need to determine the place in order to get a truth-value. Now, this place can be given as part of the situation the utterance concerns. In this case, Perry calls it an 'unarticulated constituent'. In many cases, the place will just be the place of utterance. In other cases, it might be a salient place, for example the one the conversation is about[11] . Using Perry's terminology, we can say that an

[10] This idea is the basis for relativist frameworks in general, not only Recanati's.

[11] In this framework, the situation of evaluation need not be identical to the situation of utterance, but it is nonetheless determined by it. This is why Recanati calls his view 'moderate relativism'. By contrast,

utterance of the sentence 'It's raining' concerns a place, even though it is not (explicitly) about it.

The way Recanati defines the Distribution principle suggests that some elements that affect truth-value can but are not necessarily given by the circumstance. The place where it is raining can be given by the circumstance, but it could also be articulated in the sentence. The utterer of 'It's raining' could have used the sentence 'It's raining in Barcelona' or 'It's raining here'. If this is the case, we still need a circumstance of evaluation in order to get a truth-value, but the role of the circumstance of evaluation has become uncontroversial. It is a consequence of Duality, not of Distribution. So, in absence of further argument, even if we cannot model 'It's raining' as a complete proposition, we can model 'It's raining here' as one.

However, Recanati gives a second argument extending the need for Austinian propositions beyond unarticulated constituents. The idea is that utterance truth-value does not necessarily coincide with lekton truth-value, even when the articulated content is enough to get a truth-value. Here is the example he uses[12] . I'm watching a poker game, and I say: 'Claire has a good hand'. As it happens, Claire is not among the players. I was mistaken about who the players are. Intuitively, my utterance is not true, for my utterance was about the poker game I am watching. However, suppose that, by coincidence, Claire is playing poker at a different place, and has indeed a good hand. There is a sense in which what I said is true. Moderate relativism can accommodate both intuitions. On the one hand, the sentence is true, for it says that Claire has a good hand and that is the case. On the other, my utterance of it is false, for it purported to characterise a given poker game (the one I was watching) and Claire was not among the players.

As a result, the two arguments provide two different reasons supporting the need for two levels of content. According to the first, we need Austinian propositions because lekta, or explicit content, sometimes lack a determinant of truth-value. If we want a content with absolute truth-conditions, we need, for some sentences, a content that includes the circumstance of evaluation (providing a time, a place, etc.). According to the second, we need the Austinian proposition, even when everything is articulated, because the situation at which the utterance must be evaluated depends on the context of utterance. In principle, we can always have a 'Claire has a good hand' case. The idea is that, even if the sentence is truth-evaluable, this doesn't imply that the truth-value of

MacFarlane (2014) reserves the term 'relativism' for views, like his own, in which the circumstance of assessment is not determined by the circumstance of use. On Recanati's terminology, MacFarlane's view is dubbed 'radical relativism'.

[12] Recanati takes this example from (Barwise and Etchemendy, 1987).

an utterance of it will always be the same—our judgements of truth-value are sensitive to the situation utterance purports to depict.

Occasion-sensitivity provides a third argument for the distinction lekton-Austinian proposition. It shows that truth-conditions depend on something that cannot be articulated in the sentence. Why? Because there are reasons to think that, for any sentence, a Travis case can be generated—for 'The leaves are painted green', 'The leaves are naturally green', and so on. Not all the elements on which truth-value depends can be articulated in natural language. Thus, the possibility of iterating Travis cases motivates a reading of the Distribution principle that extends the need for Austinian propositions beyond Perry-style unarticulated constituents. It is not just that some linguistic representations are not fully articulated. Rather, what we have now is the claim that no linguistic representation is fully articulated, where a fully articulated representation would be one that makes everything explicit and gets rid of occasion-sensitivity. As a result, occasion-sensitivity motivates a stronger reading of the Distribution principle:

> [Distribution*] The determinants of truth-value distribute over the two basic components truth-evaluation involves : content and circumstance. Not all determinants can be fully articulated in the sentence. Some must be provided by the circumstance.

The reason why not everything can be made explicit is that adding new linguistic material to a sentence brings in new indeterminacies (chapter 3). Although we can articulare some things, we cannot articulate everything.

Corazza and Dokic (2007; 2012) use the idea that utterances concern situations in order to explain some of the examples that have been discussed in the contextualist debate, including cases of incompleteness as 'Tipper is ready', where we need to add what Tipper is ready for in order to get a complete proposition, or more Travis-style example, as 'There's beer in the fridge' (an example analogous to Travis's 'There's milk in the refrigerator'). According to their account, called situated minimalism, the truth-conditions of an utterance are given by relative T-sentences such as 'An utterance u of "There's some beer in the fridge" is true iff there's some beer in the fridge in the situation of u.' (Corazza and Dokic, 2012, p. 187). Following their idea, we could model the content of the utterance as an Austinian proposition involving a lekton and a situation. However, it is not clear that the notion of situation will do the required job here. In order for this explanation to work, the notion of situation at stake must be one different from Barwise and Perry (1983), for the state of the refrigerator can remain

constant across the occasions described in the example. It seems that what is doing the work here is something else[13] .

The need for situations in semantic analysis can be traced back to Austin, who wrote:

> A statement is made and its making is a historic event, the utterance by a certain speaker or writer of certain words (a sentence) to an audience with reference to a historical situation. (Austin, 1950, p. 20)

This idea that utterances, or statements, refer to situations opened the door to a new notion of proposition, explored by Barwise and Etchemendy (1987). Whereas indexical-free Russellian propositions have the same truth-conditional content regardless of the context in which they are tokened, on the Austinian view truth is sensitive to the situation an utterance is about—as the 'Claire has a good hand' example shows. I think that Austin's remark is also the key to a proper understanding of occasion-sensitivity[14] . However, the notion of situation that is needed in order to account for occasion-sensitivity is a bit different from the one used in situation semantics.

In Barwise and Perry's seminal work on situation semantics (Barwise and Perry, 1983), reality is conceived as consisting of situations—parts of the world. The leading idea is that we always find ourselves in a situation: 'we see them, case them to come about, and have attitudes towards them' (Barwise and Perry, 1983, p. 7). Individuals, properties and locations are uniformities across real situations. These uniformities are the building blocks of abstract situations—the tools for the semantic theory. Now, in Travis cases, it seems that something else must be added to the situation, for individuals, properties and locations remain constant across the two occasions depicted.

As I said in the previous section, I think that occasion-sensitivity shows that, beyond including objects, places and times, situations must be conceived as involving linguis-

[13] Corazza and Dokic's explanation of the example mentioned is in fact so similar to Travis's that their notion of situation seems to be equivalent to Travis's notion of occasion. They write: 'When Jane hears John say that there's beer in the fridge, she answers that he's wrong, because she knows that the fridge is free of bottles of beer. What counts as the fact that there's beer in the fridge in John's situation is the presence in the fridge of at least one bottle of beer. On the other hand, when John hears Jane's "There's beer in the fridge" he has to admit that she's right. What counts as the fact that there's beer in the fridge in Jane's situation is the presence of beer stains that have not been properly removed. Jane evaluates John's utterance as false because there's nothing to drink in the fridge, and John evaluates Jane's utterance as true because the fridge is still not quite clean'(Corazza and Dokic, 2012, p. 190).

[14] Unsurprisingly, given Travis's acknowledgement that he found the core idea of occasion-sensitivity in Austin's work.

tic and extralinguistic activities—more in line with Wittgenstein's language games than with Barwise and Perry's situations[15] . Thus, we can take the salient variation in Pia's example to be due to the activity in which the act of utterance is embedded. In the first context, Pia is painting leaves so as to match the colour they have in spring (or for decorative purposes, let's say), whereas the botanist is doing scientific research[16] .

The distinction lekton-Austinian proposition can be easily adapted to Travis cases. I will model Austinian propositions as <lekton, activity> pairs. Let us go back to the green leaves example: It is winter, and Pia's tree is full of brown leaves. She is trying to decorate the garden and, thinking that the tree looks very ugly and that green leaves are always beautiful, she decides to paint the leaves green. After doing it, she says: 'That's better. The leaves are green now'. What she says is true. Later, a botanist phones, seeking green leaves for some scientific research on the properties of green plants. Again, Pia says (this time addressing the botanist): 'The leaves (on my tree) are green. You can have them'. Now, what she says if false. So the semantics of 'The leaves are green' is compatible with both saying something true and something false of the same leaves.

On the one hand, there is the lekton, that is, the sentence meaning. The lekton can be thought of as a structured proposition. In the simplest case, it is composed of a property (the colour green) predicated of an object (certain leaves). The lekton has different truth-values at different activities. On the other hand, there is the Austinian proposition, that is, the content of the utterance. The Austinian proposition is the lekton as used and evaluated at a specific activity. As a result, in Travis cases we have two utterances that share the lekton but are not expressions of the same Austinian proposition; they concern different activities. Given that truth-value depends on some features of the activity not articulated in the lekton, these utterances can have different truth-values. In particular the shared lekton of Pia's utterance is, simplifying a bit the counterpart of the description 'the leaves', the proposition Fa, where F refers to the property of being green and a refers to certain leaves. This lekton, as used to talk about Pia's leaves, is true

[15] One can speak of Wittgensteinian propositions if the difference seems substantial.

[16] In Travis cases, the occasion is usually considered to be the occasion of use because it coincides with the occasion of evaluation. However, there could be cases in which the interlocutors are engaged in a conversation about something they are not doing. For example, they could be working in a lab, on some research on chlorophyll, but talking about decoration or photography. In this case, what counts as 'green' in their conversation would have to do with the activity of decorating leaves or taking pictures, not with scientific research. In both cases, the relevant activity seems to be fixed by the circumstances of use. By contrast, MacFarlane advocates a relativistic framework in which the value of the parameters needed to get a truth-value are determined, not by the context of use, but by the context of assessment. I will not address the issue whether something similar could be the case with the predicates involved in Travis cases.

at some occasions and false at others. It is false at occasions where the relevant activity is one in which what matters are the natural properties of the leaves, and true at occasions is which what matters is the visible aspect, let's say. The content of the first utterance can be modelled as <Fa, $activity_1$>, whereas the content of the second utterance would be <Fa, $activity_2$>. As a result, the first utterance is true if and only if Fa is true at $activity_1$ (or: if and only if the leaves count as 'green 'at $activity_1$). The second utterance is true if and only if Fa is true at $activity_2$. Since the leaves count as 'green' in $activity_1$ but not in $activity_2$, the utterances have different truth-values.

I am using the notion of activity here because it seems like a natural way to capture the idea that what counts as 'green' on an occasion depends on what is going on on that occasion. Alternatively, we could talk about the purposes of the conversation. I take purposes and activities here to be roughly equivalent: if the conversation concerns scientific research we can also say that its purpose is to do scientific research, and vice versa. In both descriptions, what makes it the case that painted leaves do not satisfy an utterance of 'The leaves are green' is that science is about natural properties[17] [18] .

I will finish with two remarks about activities. First, it can be doubted that all talk concerns an activity and that, even if that is so, the activity is what is relevant in the determination of truth-value. Let us consider first some Travis-style cases. As I argued at the beginning of this chapter, when we embed the sentence 'Cut the sun' into different stories, we understand it in different ways. Are activities the relevant aspect of the stories? I think they are. Take the example in which children are cutting some shapes drawn on paper. The reason why we interpret 'Cut the sun' as being non-literal and as asking the child to cut a certain figure is that we are seeing the sentence as part of an extralinguistic activity consisting in cutting shapes. Similar explanations can be given of other examples. In the science fiction novel example, the sentence is interpreted in relation to an activity that has been introduced: cutting planets into two with a laser.

Can we generalize this claim? I think it might sound a bit unnatural to talk about activities when it comes to more complex Travis cases (iterated, or higher order, Travis cases). For example, a good way to present Pia's example is to say 'Pia has painted the leaves green'. What activity would an utterance of this sentence belong to? I think we could say that it still concerns the activity of decorating the garden. Alternatively, we could talk about activities in a loose sense. In this loose sense, we can say that an utter-

[17] By contrast, it wouldn't be entirely accurate to take the second component of the Austinian proposition to be the purposes of uttering the sentence. These can be very unspecific. For example, in Pia's second utterance they could be something like helping a friend to get what he is looking for.

[18] Kölbel (2008) mentions purposes.

ance of 'Pia has painted the leaves green' concerns to the activity of describing things as we perceive them in normal illumination conditions and at the distance at which we usually interact with them, as opposed to the activity of describing things as they look when observed through a microscope. It is in this loose sense of 'activity' that I take activities to be part of Austinian propositions[19] .

Second, activities are not further representations. Travis cases can be created for specifications or descriptions of activities, but not for activities themselves. I see activities as practices we engage in and are familiar with. It is beyond my goal here to provide a substantial notion of 'activity'. Instead, let me simply note that the notion is inspired by Wittgenstein remarks on language games, practices and customs, and on Searle's Background. Loosely following Wittgenstein, I think of activities as practices in which things are typically done in certain ways, practices we have been trained into and that we have probably developed in response to our particular biological and psychological traits[20] . The notion is also related to Searle's Background. The Background encompasses non-representational assumptions, skills, capacities, practices and habits (Searle, 1983). Although I understand activities as being external, it could be said, in a Searlian spirit, that when one is acquainted with an activity one thereby acquires a set of non-representational assumptions and habits[21] .

[19] As I mentioned, the activity can be given by default. If we very often use colour predicates in order to describe how things look in such-and-such conditions, then, in absence of further context, we will use that as a default interpretation of sentences such as 'Pia has painted the leaves green'.

[20] We usually consider an object as 'painted' even when small parts of its surface are not covered with paint. Perhaps had our visual system been more precise, we would have different standards as to what counts as 'painted'.

[21] My proposal is also in the spirit of Gauker's claim that 'Typically, conversations have goals'. Gauker notes that: 'One sort of basic goal that people might have in conversation is finding something: finding prey, finding water, finding a good place to sleep. Other goals have to do with the management of society: settling a territorial dispute, arranging a marriage, deciding how to punish a misbehaving child. I do not assume that all such goals serve survival. Our goals may include finding a beautiful vista, or learning how to play a flute.' (Gauker, 2003, p. 49). According to Gauker, an action in pursuit of a goal can be in accord or not in accord with a body of information. Actions include linguistic actions as assertions. On my view, Guaker's goals are activities.

4.3 Partiality

In chapter 3 I argued that there are reasons to take mental representations to be token-underdetermined. Let me recall the distinction between Type-Underdeterminacy and Token-Underdeterminacy:

> *Type-Underdeterminacy:* A non-indexical structured representational item S is type-underdetermined if and only if there are tokens of S that have distinct truth-values.

> *Token-Underdeterminacy:* A token of a structured representational item S is token-underdetermined if and only if for some possible states of affairs its truth-value is indeterminate (i.e., if and only if it determines a partial function from possible worlds to truth-values).

I think that the same examples that motivated Token-Underdeterminacy for mental representations suggest that this principle also holds for utterances. One of the examples I considered was inspired on the green leaves example. Let's imagine, again, that Pia and her friend are painting some leaves with green paint. After finishing, Pia says 'The leaves are green'. Given the setting, her use of 'green' is to be understood as, let's say, 'superficially green', or perhaps simply 'painted green'. However, the setting might be insufficient to determine, for any possible leaf l, whether l counts as 'green' on Pia's use of this expression. Again, imagine a leaf that has been painted with small blue and yellow dots and that, at a certain distance, looks green. It is possible that nothing in the conversational setting determines whether an utterance of 'It's green' would be true or false of that leaf. And, plausibly, the same will happen with leaves that have been only painted on one side, or, going for a more strange example, with leaves that have been painted with some paint that only looks green under certain illumination conditions (and is otherwise invisible).

In chap. 3 I suggested that we can think of the indeterminacies involved in Travis cases as questions that are not answered by the conventional meaning but by the context instead. When it comes to 'is green', these questions include the following: what part of the object must be green? On what observation conditions? Etc. Activities can be seen as providing answers to these questions. For example, in the activity 'decorating the garden', the part of the colour that is relevant is the surface. It is part of the activity that it is about the visible parts of objects. By contrast, in the activity, 'doing botany',

the answer will be answered differently. The object must contain green pigments in its interior, let's say. Some activities might leave some questions unanswered. Plausibly, the activity Pia and her friend are performing leaves it unanswered whether an object that looks green at 0,5m but yellow and blue at 0,25m counts as 'green'.

As a result, Austinian propositions determine partial functions from states of affairs (or possible worlds) to truth-values. For some states of affairs, the truth-value of the Austinian proposition might not be determined. The Austinian proposition expressed by Pia's utterance is true at states of affairs where the contextually determined leaf has been covered with green paint, false at states of affairs where the contextually determined leaf has been covered with blue paint or is brown and has not been painted, and its truth-value is not determined at states of affairs where the contextually determined leaf is painted with blue and yellow dots.

There is a reason why my proposal can be seen as more traditional than a pragmatist approach as Travis's: it allows for modelling the content of an utterance u as a function from situations (where these include activities) to truth-values. However, I think that the linguistic meaning of a sentence should not be identified with a function from situations to truth-values. Activities can be created and abandoned, and they can even be created online, whereas linguistic meaning can remain fix throughout the process. Truth-conditions are the product of embedding sentences into those activities. But sentences are embedded with their linguistic meanings.

4.4 Specificity and content-sharing

In sections 2 and 3 I have introduced the basics of my approach to utterance content. On this view, utterance content is modelled as a <lekton, activity> pair and can be underdetermined, in the sense that it doesn't determine a truth-value for any possible state of affairs. In this section, I will argue that, in order to account for intercontextual content sharing, the approach needs to be improved.

Cappelen and Lepore (2006) argue that the concept of what is said faces a tension. On the one hand, contextualists have it that what is said is closely tied to the context of use. On the other, we often share contents across contexts—we report what others say,

have inter-contextual discussions, etc[22] .

According to contextualists, what is said in an utterance is the result of con-textual adjustment. The point that Cappelen and Lepore find problematic is that this adjustment seems to be, according to contextualists, highly specific. The reason is that every small detail of a conversation might be relevant to how one understands words. Two people who have been trained in different painting traditions can understand differently an utterance of 'The leaf is green' when they are painting leaves—there can be a leaf that counts as 'painted green' for one of them but not the other. If con-tent needs to be highly specific, then activities, as they figure in Austinian propositions, should be very fine-grained.

However, it should be possible to share content across contexts. Modeling activities too fine-grained makes it difficult. In particular, it makes most speech re-ports strictly speaking false. Imagine that I report Pia's utterance by saying: 'Pia said that the leaves are green. She is painting them'. In most circumstances, this report is enough for us to grasp what Pia said. We have a broad understanding of what painting consists in. Although there are different ways of painting things we know that, roughly speaking, painting consisting in applying paint to the surface so as to hide the original colour. However, if every small detail of the context matters, we should model Pia's original utterance as involving a more fine-grained activity than the one conveyed by 'She has painted them'.

As a consequence we need to find a middle course between contents that are as context-bound as Travis cases seem to call for but that cannot be shared across contexts, and contents that can be shared across contexts but that are too broad to respect the intuition that any detail might matter.

Corazza and Dokic's situated minimalism seems to find that middle course. According to them, propositional truth in cases similar to the ones I've been considering should be relativized to the situation of use. This explains the variation in truth-value and the intuition that any detail of the context might matter to what is said. But on the other hand, Corazza and Dokic are minimalists when it comes to what an utterance says, and claim that 'Two people using the same alleged underdetermined sentence can be characterized, pace contextualism, as having said the same thing even if they are not

[22]Contextualists can explain content sharing on the basis of the similarity between the contents expressed in different contexts (Bezuidenhout, 1997). My explanation of shared content is similar to Cappelen and Lepore's speech act pluralism in that they reject monism, i.e., the view that only one proposition is expressed by an utterance, and so do I (although I will speak of utterance content being modelled with different fineness rather than of utterances expressing a multiplicity of propositions).

co-situated' (2012, p. 196).

However, same-saying at least sometimes requires something else than sameness of linguistic meaning. Some philosophers have argued that whether two utterances share content is something that needs to be decided by taking into account not only linguistic meaning (lekton) but also some features of the context of use[23]. Wieland (2010a) makes this point by focusing on indirect reports. As she argues, we rarely report what someone has said by uttering the exact same words. Instead, some additional linguistic material is usually needed. Imagine the following situation. We have a barrel full of apples. Some of the apples are affected by some fungus and we need to discard them. The fungus makes the interior of the apples red. Anne cuts an apple and says 'The apple is red'. Now imagine that the apple is left on the table. When Nelly arrives home she says that she's hungry and, since the lights are out and cannot properly see the apple asks if that's a red apple. A report such as 'Anne said that the apple is red' would be incorrect here: we need to add something else in order to capture the content of Anne's utterance. The moral of these examples is that sharing a lekton is not always enough for same-saying.

The problem with Corazza and Dokic's account of same-saying is that a homophonic speech report, as 'Jane said that there's a beer in the fridge' is false whenever the report context is relevantly different from the original context. And it is so because the sentence 'There's a beer in the fridge' does not say the same in both occasions.

My solution to the tension will involve Austinian propositions with different granularities. Before that, let me note that conceiving Austinian propositions as in-cluding activity types, instead of particular situations, already allows us to explain some cases of same-saying without the need to identify same-saying with sameness of lekton. Imagine a teacher teaching geography to some children. In order to make it easier for them to identify countries in a map, he tells them: 'Look, Italy is the one that looks like a boot. And France is hexagonal. You see?'. Her utterance concerns the activity of comparing countries with objects, let's say. As it happens, this teacher uses the same example every year, as do many other teachers. She says the same year after year. If we model her utterance as including the activity type 'comparing countries with objects', we can account for that.

Instead of particular situations, we can think of Austinian propositions as including situation-types. Thus, two utterances of the same sentence that refer to two different particular situations can express the same Austinian propositions, because of the par-

[23]See Wieland (2010a,b), Recanati (2006a) and Travis (2006a). Recanati discusses an example by Leslie.

ticular situations being of the same situation type. The content of an utterance, then, will be conceived as a pair <lekton, activity-type>. Two different activity-tokens can correspond to the same activity-type.

Despite this, I admit that a content that is closely tied to the context of use is difficult to share across contexts. If we model the content of Pia's first utterance as including a very fine-grained activity, such as 'painting the leaves in the way Pia likes' or 'painting the leaves according to such-and-such tradition', then it is going to be difficult to have indirect reports that capture the content of the original utterance. But usually, less specific contents are good enough. For most purposes, it is enough to grasp that when Pia said 'The leaves are green' (first utterance) she only cared about the superficial aspect or, roughly, about the leaves being covered by paint (as opposed to them being naturally green).

In this sense, understanding can be said to come in degrees. When two interlocutors are co-situated and aware of each other's intentions as well as the topic of the conversation, and so on, then their understanding of each other's utterances can be very deep. By contrast, when they are no co-situated, or have only broadly grasped what the conversation is about, their understanding of the utterance will be more superficial and they will only roughly grasp what the speaker means–for instance, they might grasp that Pia's utterance has to do with painting leaves, but not whether it is important that the whole surface is covered with paint.

My proposal goes as follows. We can classify activities more or less fine-grained. My suggestion is that we allow for different ways of classifying activities, ranging from very rough to very precise classifications in which any detail matters. If we look at the tension with this in mind, the threat will turn out not to be that serious. In the limiting case, the Austinian proposition can include the particular occasion–an activity so fine-grained that it will admittedly be difficult to share across contexts.

Given that we can class particular activities into activity-types with different fineness of grain, we can represent the content of an utterance as being more or less fine-grained, depending on the aim of the classification–ranging from very rough to very precise classifications in which any detail matters. Thus, we can have more or less fine-grained propositions. For example, we could have a range of propositions that could model the content of Pia and John's utterance:

P1: <The leaves are green, Decoration>

P2: <The leaves are green, Painting leaves>

P3: <The leaves are green, Painting leaves in a way that the original colour is no

longer observable>

P4: <The leaves are green, Painting leaves in a way that the original colour is no longer observable and excluding pointillism>

And the limiting case:

P5: <The leaves are green, Particular occasion>

We can class a given utterance as expressing one proposition or another, depending on the purposes of the classing. If we want to model the content of Pia's utterance as being closely tied to the context of use, we can use P4 or P5. If we want to model the content that is shared across contexts (for example, in a speech report), we can use P1 or P2.

Propositions can be understood here, following Perry (2001) and Korta and Perry (2007; 2011), as abstract objects that are used for classificatory purposes:

> [T]he reflexive-referential theory sees propositions as abstract objects that are used to classify events of certain types (cognitive states and utterances, paradigmatically) by conditions of truth (or other relevant forms of success)—used explicitly by theorists such as ourselves, and implicitly in the practice of those who have mastered the propositional attitudes and similar constructions. We do not see propositions as denizens of a third realm to which some quasi-causal relation relates us, but as devices by which we can classify events along different dimensions of similarity and difference. Different propositions can be used to classify the same act, relative to different frameworks for associating success-conditions of various sorts. (Korta and Perry, 2007, pp. 176-177).

Thus, the range of Austinian propositions I have described can be seen as a tool for classifying linguistic events. Sometimes, very fine-grained propositions are needed. These fine-grained propositions allow us to capture the insight that, as Travis cases suggest, any detail of the utterance context might matter to truth-conditions. On the other hand, less fine-grained propositions can be used in classifying linguistic events that are similar enough to count as same-saying.

To sum up, I have modelled utterance content as a <lekton, activity-type> pair and advocated a form of multipropositionalism according to which we can represent the content of an utterance with different granularities, depending on whether we are interested on same-saying or very detailed truth-conditions. The notion presented captures the three claims motivated by occasion-sensitivity. First, the truth-value of the lekton is relative to activities. Second, utterance content is truth-conditional. And third, the role of the context is properly acknowledged.

I will finish by noting that the approach presented here can be of some interest for the minimalist as well. According to the minimalist, non-indexical sentences are truth-evaluable. Semantics determines classical propositions. However, minimalists typically agree that, when it comes to communication, these semantically determined propositions are insufficient. Thus, Cappelen and Lepore (2005) opt for speech act pluralism. Much of the contextualist debate concerns the notion of content relevant for communication and the interest on the truth-evaluability of semantically determined propositions is arguably of limited interest when it comes to explaining communication. In a review of Borg's work Carston writes:

> [Borg] discusses and puts to rest (hopefully forever) the idea that a semantic theory is required to deliver for any sentence s the content p , where for a speaker S who utters s, 'S said that p' is a correct indirect report. As she argues, both intuitive pre-theoretic notions of 'what is said' and the technical Gricean notion of 'what is said' are a matter of communication (or speaker meaning), hence involve consideration of speaker intentions, and so fall outside the domain of a formal semantic theory, while any more restrictive notion of ' what is said ' turns out either to be dependent on the very semantic judgements it is meant to constrain or to be redundant in determining semantic content since the real work is being done by syntactic features of the sentence uttered. As suggested above, once we recognise that what the contextualist is focussing on is this pragmatic notion of what is said (that is, the truth-conditional content of an utterance or speech act) as distinct from a formal 'pure' level of linguistic semantics, much of the fire goes out of the minimalist/contextualist debate. (Carston, 2008b, p. 361).

The notion of Austinian proposition I have put forward can be of interest in accounting for utterance content regardless of what one thinks of minimal propositions. A minimalist can have it that lekta are truth-evaluable and yet acknowledge that the truth of an utterance (not of the sentence) is relative to the relevant activity. The reason why a minimalist might find Austinian propositions useful is that of the 'Claire has a good hand' example: one thing is the truth-value of the sentence (in case it has one), another is the truth-value of the utterance. The latter arguably depends on what is at stake.

4.5 Occasion-sensitivity in the tradition of situation semantics

In chapter 1 I presented a potential problem for occasion-sensitivity. The problem is the following. If meaning is occasion-sensitive, then it seems that meaning is not objective, for it does not extend of itself. The problem could be overcome if we thought of predicates as expressing properties more specific than what is encoded in natural language. However, in chapter 3 I have casted doubts on the existence of occasion-insensitive mental representations and the role of occasion-insensitive structured propositions in our linguistic and cognitive lives. As a result, there can be doubts that occasion-sensitivity is compatible with the objectivity of content, i.e., the independence of the truth-value of a content from our judgements. Why? Because occasion-sensitive representational items do not extend of themselves and so an occasionalist might seem to end up claiming that whether a given representational item is true as used in a given occasion partly depends on how we treat it—for example, on whether we take it to be reasonably taken as true.

In his (2010b) Travis attributes to Wittgenstein an account that is similar to the one I have advocated in this chapter. In this account, the truth of an utterance is contingent upon the language game in which the utterance takes place. Language games are defined by introduction and elimination rules. He offers as an example a film-developing game. Uttering 'The room is dark' is a move in this game. By its elimination rule, when a player utters this sentence, the other player can remove the film from the camera and place it in the developer. This fixes the conditions under which it is correct to use the sentence (introduction rule)—it is correct if the room is dark enough not to spoil film, let's say. Travis writes:

> My calling the room dark raises a question of truth: is it dark on the understanding on which I spoke of that? This question still leaves room for the parochial to work. There are things to be decided as to the understanding of that on which I did so speak. Now, the idea is, say what you like as to what that understanding is, and there is still the same sort of room for the parochial to work. There is always room for fresh questions of truth to arise. Find new circumstances of which my words might be true or not, and there are, correspondingly, new language games for me to have been playing or not, games which provide for such circumstances in any of many ways (or none). With that, there is more of the usual work for the parochial to do. The parochial, like rust, never sleeps. On this

> new picture, as opposed to Frege's, it never cancels out absolutely.
>
> Invisible, intangible truth bearers thus lose the special status Frege accords them. The parochial takes on an ineliminable role in representing as so. Such a view must be seen as a picture of what representing is as such—though it is as yet unclear how it could be that. (Travis, 2010b, p. 179).

According to Travis, the role of the parochial never cancels out. However, it is difficult to see what is meant by that here. If the role of the parochial is non-eliminable, does this mean that the truth-conditions of an utterance somehow depend on us—on our judgements or reactions? Here is where the problem would appear. Travis is well aware of the threat:

> And does it have that objectivity which is the hallmark of real thinking (taking) something to be so? Does its would-be truth turn simply on how things are, independent of our so taking them? Or does whether things are as we represent them turn, in part, at least, on our responses to them—what we are inclined, or designed, to feel or say, confronted with the world? Does the parochial contribute merely to how we represent things to be? Or must its attempts to do that ineluctably taint when things would count as thus represented? The phenomena of occasion-sensitivity can, in fact, make it difficult to see the answer here. (Travis, 2010a, p. 16)

I think that a situationalist framework, together with a careful distinction of two different roles for the parochial to play, provides the means to solve the problem. I will start with the role of the parochial. Travis calls 'the parochial' our particular way of thinking, including here our psychological traits that need not be shared by other kinds of thinkers. It could be playing two different roles. First, it can determine what content a given utterance expresses. Take demonstratives as an example. Gauker (Gauker, 2008) has argued that the referent of a demonstrative is determined by an all-things-considered judgement taking into account, among others, salience, pointing, relevance, and previous reference. If Gauker is right, then the parochial is plausibly doing some work here, for it plausibly depends on our specific psychological traits how we calibrate the importance of the things that go into the all-things-considered judgements. Similarly, the truth-conditional content of an utterance (what counts as 'green', for example), could be the case of an all-things-considered judgement, or depend on what is relevant—for thinkers like us. If so, the parochial plays a role in the identification of the

content of our utterances[24] . Second, the parochial could be playing a role in determining whether the content expressed by an utterance is true or false. This would be the case if, for example, the truth-value of an utterance depended on whether the interlocutors of the conversation found it reasonable to take it as true. In this approach, whether 'The room is dark' is true or false depends on the participants in the conversation where the sentence is used taking it to be reasonable that the room is described as 'dark'. In general, the parochial will be playing this second role whenever truth-value depends on how we happen to react. This is what might preclude objectivity.

I think that adopting a situationalist framework can solve the threat. In the view I have presented, utterance content is conceived as an Austinian proposition. The parochial will very likely play a role in determining what Austinian proposition a given utterance expresses. What activity a given utterance concerns might be up to us—for example because of being the outcome of an all-things-considered judgement. However, it does not, on top of that, play a role in determining whether the Austinian proposition expressed is true or false. That will depend only on how things are, the lekton and the activity in place. Using Wright's terminology, within an activity meaning can be said to extend of itself to unconsidered cases. The activity contributes further criteria that constrain the application conditions of the lekton, independently of our reactions.

Let us go back to Travis' view on Fregean truth-bearers from chapter 1:

> If circumstances of a stating matter to what was stated along the lines indicated here—lines on which what was stated is fixed by what one then had a right to expect of things being as stated (of a cat's having mange, say) — then there is in principle no end of opportunities for circumstances of a stating to matter to what was stated. There is no point at which circumstances choose for us some truth-evaluable item which is itself immune in principle to admitting of different further understandings — no point at which, through appeal to circumstance, we arrive at the sort of invisible, intangible truth-bearer (what Frege called a 'thought') which, Frege held, was the only thing that could really make a determinate question of truth arise. (Travis, 2008a, p. 6)

Frege's notion of a thought is one according to which thoughts are complete. This invites a picture where propositions have the same truth-conditions whenever they are

[24]Moreover, it might be the case that which words we coin depend on our psychological setup and our needs. This, however, does not call the objectivity of content into question. Once it is settled what the content of an utterance is, the role of the parochial cancels out.

tokened. Travis cast doubts on the accuracy of this picture by showing that predicating the property GREEN of certain leaves does not always result on the same truth-conditions (keeping the time fixed). The truth-conditions of the resulting structured proposition seem to be contingent upon what is at stake at the tokening context.

Now, admitting that structured propositions are incomplete in the sense that something else is needed to get non-relativized truth-conditional content does not commit one to a picture where our judgements about what is reasonable play an important role, for the tokening context can provide an activity that, together with the structured proposition, determines truth-conditions. This option, in the tradition of situation semantics, respects the main insights of Travis cases, yet does not run the risk of precluding objectivity.

4.6 Conclusions

In this chapter I have put forward a notion of utterance content in accord with the lessons from Travis cases. Following Recanati, I have distinguished sentence meaning from utterance content and modelled the latter as a <lekton, activity> pair. This notion of content escapes two potential problems.

The first problem, common to other forms of truth-conditional pragmatics, is that of content sharing. On the one hand, the content of our utterances is closely tied to the specifics of the context of use. But on the other hand, we sometimes share contents across contexts. For example, we report what others have said. I have claimed that Austinian propositions as <lekton, activity-type> pairs, combined with multipropositionalism, can accommodate both desiderata. First, activity-types can be shared across occasions, which explains some cases of same-saying. Second, we can have fine-grained and coarse-grained activities. Austinian propositions with coarse-grained activities can be shared across contents and sharing it is often enough for the correctness of a speech report.

The second problem was introduced in chap. 1. There I argued that not having semantic items whose meaning extends of itself precludes objectivity. In this chapter I have introduced a non-standard notion of content that, I have argued, escapes the problem. The reason is that within an activity meaning can be said to extend of itself.

Chapter 5

Grasping truth-conditions

Travis cases motivate a rejection of certain semantic theories[1]. In this chapter I will argue that abandoning all sorts of systematic theories of truth-conditions can be problematic. One of the main goals of semantic theories is to explain linguistic understanding. The reason why theories that deliver truth-conditions are useful to this aim is that understanding an utterance of a sentence is usually identified with grasping the conditions under which it would be true. Advocates of truth-conditional pragmatics, including advocates of occasion-sensitivity, do not dispute this view about understanding. However, if they were to reject any form of systematic semantics they would run the risk of turning our ability to interpret speech (grasping truth-conditions) into something mysterious. As a result, an advocate of occasion-sensitivity, as well as any other advocate of truth-conditional pragmatics, needs to show that his proposal is compatible with a plausible explanation of our ability to interpret speech. In order to do so, I will argue that, although occasion-sensitivity is incompatible with Davidson's T-schema, the notion of utterance content introduced in chapter 4 is compatible with there being systematic connections between sentences and activities, and truth-conditions. Moreover, following Predelli, I will show that it is compatible with double-indexed theories. I will also argue that occasion-sensitivity can be made compatible with theories that specify truth-conditions by having recourse to default understandings.

[1]By 'semantic theory' I mean here a theory of truth-conditions.

5.1 Linguistic competence and systematic theories of truth-conditions

Let me start by paraphrasing Davidson[2] : Pia utters the words 'The leaves are green' and under certain conditions we know that she said something that is true (false) of certain painted leaves. Having identified her utterance as intentional and linguistic, we are able to go on to interpret her words: we can say what her words, on that occasion, meant. What could we know that would enable us to do this?

Advocates of truth-conditional pragmatics think that Davidson's truth-theory won't work. They claim that the meaning of a non-indexical declarative sentence S underdetermines the truth-conditions of its tokens. In contrast to what is often assumed by semanticists, they argue that the truth-conditions of our uses of non-indexical sentences depend on the specifics of the contexts of use and are not determined (exhausted) by linguistic meaning. Moreover, they sometimes take it that truth-conditions cannot be fully systematized in a formal theory. Bezuidenhout writes:

> The alternative view holds that meaning underdetermines truth-conditions. What is expressed by the utterance of a sentence in a context goes beyond what is encoded in the sentence itself. *Truth-conditional content depends on an indefinite number of unstated background assumptions, not all of which can be made explicit.* A change in background assumptions can change truth-conditions, even bracketing disambiguation and reference assignment. That is, even after disambiguating any ambiguous words in a sentence and assigning semantic values to any indexical expressions in the sentence, truth-conditions may vary with variations in the background. (Bezuidenhout, 2002, p. 105) (Emphasis added.)

Following Searle, Bezuidenhout takes it that we cannot identify and make explicit all the background assumptions on which truth-conditions depend. Similarly, but on the Chomskian tradition, Pietroski seems to be sceptical that we can have theories of truth-conditions. He writes:

[2] Davidson writes: 'Kurt utters the words "Es regnet" and under the right conditions we know that he said that it is raining. Having identified his utterance as intentional and linguistic, we are able to go on to interpret his words: we can say what his words, on that occasion, meant. What could we know that would enable us to do this? How could we come to know it? (Davidson, 1973, p. 313)' A Davidsonian semantic theory, or a theory of interpretation, is intended to be what we could know that would enable us to interpret speech.

> The fact that (an utterance of) a sentence has a certain truth-condition is typically an *interaction effect* whose determinants include (i) intrinsic properties of the sentence that we can isolate and theorize about, and (ii) *a host of facts less amenable to theorizing, like facts about how "reasonable" speakers would use the sentence.* (Pietroski, 2003) (Emphasis added.)

Travis is sceptical that there can be what he calls generative theories of understandings or of (truth-conditional) contents:

> [S]uch a theory must have at its disposal an adequate range of representational features. Suppose we can find a pair of statements which differ in the understanding they bear. Then the theory must be able to attribute to each a feature that it has and the other lacks. Now, to suppose that there *is* some correct theory of the envisioned sort is to suppose at least this: for any statement, S, there is some specifiable set of representational features such that no statement with those features could differ from S in when it would be true [...]. Such a set of features would identify *precisely* one understanding there is for a statement to bear–one certain content that might be the content of that statement. Let us call such set of features, and the representational structure they identify, a disambiguation of S. [...] Then a theory of understandings of the sort described–as it were, a generative theory of contents–would generate disambiguation for any statement (perhaps within some range). [...]
>
> It is thus also clear how there could fail to be any correct theory of the envisioned sort. The idea would be this: take any statement and ascribe to it any set of specifiable representational features you like; then two or more statements might all share those features, yet differ in what they said, and hence in when they would be true. (Travis, 2000, pp. 34-36)

Semantic theories typically aim at systematically linking sentences, or, better, syntactic structures, with truth-conditions. This project works on the assumption that there is a set of representational features (a disambiguation) such that if two statements (utterances) express the same disambiguation, then they express the same truth-conditional content. The problem, for Travis, is that given that predicates, and for him also properties, have different satisfaction conditions on different occasions, one syntactic structure can express different contents (or bear different understandings). So generating disambiguations is not a way of disambiguating truth-conditional contents.

Theoreticians interested in natural language work with theories that match sentences with truth-conditions, using as input the meaning of the simple expressions and

the rules of composition. The reason is twofold. First, the aim of much of semantic theorizing is to explain linguistic understanding. Semantic theories are often conceived as modelling our linguistic knowledge[3] . As an example, let me quote Larson and Segal:

> We can see semantics as a theory of the knowledge that underlies our ability to make semantic judgements. Semantic theory addresses one part of our linguistic knowledge: *knowledge of meaning.* (Larson and Segal, 1995, p. 10)

Larson and Segal follow here Chomsky, who takes linguistics to be a branch of psychology dedicated to the study of an aspect of the mind—linguistic competence, i.e., the speaker's knowledge of his language[4] :

> Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of the language in actual performance. (Chomsky, 1965, p. 3)

> A grammar of a language purports to be a description of the ideal speaker-hearer's intrinsic competence. (Chomsky, 1965, p. 4)

One important feature that a theory of linguistic competence needs to explain is productivity, i.e., our ability to understand sentences we have never heard before. Because of this, it is usually thought that only a compositional theory can model our linguistic competence.

Second, understanding a sentence, or an utterance of a sentence, has to do with knowing under which conditions that sentence (or utterance) would be true. This

[3] Partee (1979) distinguishes two views on semantics. The first, the Chomskian view, takes the central goal of linguistics to be that of providing an account of the knowledge that a speaker has when he knows a language. The second, ascribed to Montague, takes semantics to be a branch of mathematics. These views are compatible. Interestingly, authors in the Montagovian tradition as Lewis (1970) endorse a principle of compositionality. The reasons typically offered in support of such a principle have to do with the productivity of natural language. This shows that even authors working on the second view might want semantic theories that can be used to model linguistic competence.

[4] Chomsky does not identify meaning with reference and therefore does not take the aim of semantics (as a theory of meaning) to be that of giving a theory of truth-conditions. This is a further step not taken by all of those who take the project of linguistics to be that of explaining linguistic competence.

identification of linguistic understanding with knowledge of truth-conditions has been standard in philosophy[5] .

With these two desiderata (productivity, truth-conditions) in mind, it has been thought that, if we had a theory delivering the truth-conditions of the well-formed declarative sentences of a language on the basis of the meaning of its simple expressions and the rules of composition, then that theory would be an important step towards an account of our linguistic knowledge—it would be a candidate to model our ability to interpret speech. This is the project that Davidson pursued (Davidson, 1967, 1984).

In the Davidsonian account, semantic theories are theories of interpretation, in the sense that they are theories that an interpreter could use in order to interpret utterances of a language. The theory generates, for every potential sentence of a language, a theorem that specifies its truth-conditions. Given productivity, the theory is supposed to generate an infinite number of theorems on the basis of a finite number of axioms. It can be conceived as a set of axioms establishing the meaning of simple expressions, together with a set of rules of composition from which theorems of the form of the T-schema can be derived:

(T) 'p' is true if and only p.

For example, "Snow is white' is true if and only if snow is white'. On the right-hand side of the biconditional the sentence is used. Thus, the semantic theory tells us, in the metalanguage, what the sentences of a language say. It is conceived as delivering interpretations, truth-conditional contents. In this sense, it is interpretive–it could be used in order to interpret an unknown language.

In the case of sentences containing indexicals and demonstratives things get more complicated. As a example, the truth-conditions of 'She is lazy' could be given by a conditional T-schema:

(1) If x is referred to by *she* in the course of an utterance of 'She is lazy', then that utterance is true just in case lazy(x)[6] .

Following Davidson, semantic theories have been seen as an attempt to model our linguistic competence. These theories have a set of axioms from which truth-conditions

[5] An early statement of this idea can be found in Wittgenstein's Tractatus (4.024): 'To understand a proposition means to know what is the case if it is true.' (Wittgenstein, 1994) [1922]. See Wiggins (1997) for a historical overview of the relation between meaning and truth-conditions.

[6] I follow here (Higginbotham, 1988).

for all well-formed sentences can be derived[7] . When given as input a sentence, usually together with some contextual information, the theory is supposed to deliver an interpretation.

With this in mind it is easy to see that Travis' rejection of semantic theories might be problematic. Truth-conditional semanticists are (allegedly) able to provide, at least to some extent, theories that explain, or even model, our capacity interpret speech. Advocates of truth-conditional pragmatics as Travis claim that truth-conditions are not determined by a set of features that we can specify and systematize, but rather that they depend on the specifics of the context, perhaps in unsystematic ways. If they are right, Davidson's project is ill-conceived. But then, if not in virtue of recognising a set of representational features, how is it that we are able to grasp the truth-conditions of an utterance?

Bezuidenhout is clear about the problem:

> If we assume that this understanding is not magical, there must be a systematic account to give of what is understood (the truth-conditional content of utterances) and how it is we are able to know this (what our semantic knowledge consists in and how this knowledge is used in context to understand what is said). (Bezuidenhout, 2002, p. 106)

Advocates of truth-conditional pragmatics are under pressure to show that understanding is not magical, and this amounts to showing that the truth-conditions of our utterances can be not driven by linguistic meaning, yet systematically obtained by using linguistic meaning and contextual information.

A question analogous to the one Davidson asked is motivated by the occasion-sensitivity scenarios. Pia utters the words 'The leaves are green' and under the right conditions we know that she has said something that is true (false) of certain painted leaves. What could we know that would enable us to do this, i.e., what could we know that would enable us to understand some given words in certain ways on certain occasions, to grasp the truth-conditional content of that utterance?

Giving up systematic theories of interpretation might mean giving up the project of explaining how we go about interpreting speech. And this is problematic, for claiming that it is impossible to have a systematic account of the process of interpretation might amount to claiming that we have a mysterious ability to classify some utterances as true

[7] As an example, Larson and Segal write: 'A speaker's knowledge of meaning for a language L is knowledge of a deductive system (i.e., a system of axioms and production rules) proving theorems of the form of (T) that are interpretive for sentences of L' (Larson and Segal, 1995, p. 33).

and some others as false, where whether it is one or the other depends on a very limited linguistic meaning together with some features of the context we cannot identify. The threat, in short, is that views as that of Travis', with its rejection of generative theories of content, turn interpretation into a mysterious process. Or, in other words, that they end up claiming that we understand utterances in different ways, and that we cannot explain why. But surely, if it is the case that we sometimes interpret utterances of the same sentence differently, then there must be something in virtue of which we do so. Even if there is no mechanism that tells us in advance, independently of the context, the truth-conditions of an utterance, there surely are some relevant similarities and differences across contexts that we recognize and use in order to interpret utterances. And if so, we might have systematic accounts establishing how truth-conditions depend on specific information we abstract, perhaps in unsystematic ways, from the context.

In what follows I will explore whether occasion-sensitivity is compatible with a systematic account of truth-conditional content. In section 2 I present what occasion-sensitivity is incompatible with. In section 3 I will argue, following Predelli, that the phenomenon deployed in Travis cases is compatible with double-indexed theories. However, Predelli's account might be subject to a criticism by Lepore against semantic theories in the Montagovian tradition–namely, that these theories are not interpretive. I will argue that occasion-sensitivity is incompatible with having semantic theories that, as Davidson's, aim at specifying truth-conditions. Nonetheless, if it is understood along the lines of the proposal presented in chapter 4, it is compatible with there being systematic connections between sorts of situations (or activities) and truth-conditions. These systematic connections can be used to show that occasion-sensitivity does not turn understanding into something magical. Still, in section 5 I will argue that one can have theories that specify truth-conditions if one accepts that these take the terms they involve to be understood by default.

5.2 Two problematic assumptions

Travis cases provide evidence against Semantic Propositionalism and semantic determination (chapter 2). According to Semantic Propositionalism, the semantics of a well-formed declarative sentence S determines a truth-evaluable content (a proposition, a

truth-condition). Let us restrict this principle to non-indexical sentences:

> *Semantic Propositionalism (non-indexical sentences):* The semantics (meaning) of a well-formed non-indexical declarative sentence S determines a truth-evaluable content (a proposition, a truth-condition).

According to this principle, some sentences (non-indexical declarative sentences, or most of them) are such that their meanings are sufficient to determine a truth-evaluable content[8] . Presumably, all the tokens of those sentences will express the same truth-evaluable content.

Semantic Propositionalism is present in two different parts of certain semantic theories. On the one hand, some theories attribute truth-evauable content to the well-formed declarative, non-indexical sentences of a language on the basis of the meanings of the simple expressions of the language and the rules of combination. So they must be assuming that sentences, and not speech acts, are bearers of truth-evaluable content. On the other hand, if at least some sentences express always the same truth-evaluable content, then it seems that we have the means to specify (in the sense of articulate, encode in natural language) what the truth-conditions of a sentence, or even sentence at a context, are. We just need to use a sentence that encodes them. Thus, when we use instances of the T-schema as ' 'Snow is white' is true if and only snow is white' we rely in that the right hand side of the biconditional will state the conditions under which the sentence 'Snow is white' is true.

Travis' criticism of Semantic Propositionalism poses a double problem. As it turns out, given that we can generate Travis cases for a large class of sentences, we have no reasons to think that there is a relevant class of sentences that express the same truth-evaluable content in any occasion of use (leaving mathematics aside). So views that state truth-conditions for sentences along the lines of Davidson's are problematic. Moreover, given that we have reasons to think that most (or perhaps all) sentences are occasion-sensitive, the possibility of stating the truth-conditions of a sentence, or a proposition,

[8] I will speak of truth-evaluable content, instead of truth-conditions, because in intensional semantics truth-conditions can be understood as truth-value distributions over possible worlds. Similarly, one can identify the truth-conditions of a sentence with a truth-value distribution over indices of evaluation including other parameters besides a possible world (standards of precision, a 'what is at stake' parameter, etc.). Although these theories, in a sense, attribute truth-conditions to sentences, they do not attribute truth-evaluable content to sentences, but rather to sentences at a context (where the context determines the relevant index).

is called into question, for any sentence that we use will be, as a sentence, occasion-sensitive and so its understanding will rely on inexplicit assumptions that can vary across contexts. As a result, we cannot make explicit all the assumptions on which truth-conditions depend by using natural language sentences, for those very same sentences will be understood against other implicit assumptions. This is particularly harmful for Davidsonian semantics. Travis writes:

> The driving force of [occasion-sensitivity] is this idea: the open sentences of language speak of ways for things to be which admit of understandings [...]. This blocks truth-conditional semantics. For suppose I say, 'The sentence "Sid grunts" is true iff Sid grunts'. Either I use that last 'grunts' on some particular understanding of being a grunter—one understanding among many—or I do not. If I do, then I assign the sentence a property it does not have. For it does not speak of being a grunter on any special understanding of this. But if I do not, then I fail to state any condition under which anything might be true. Being a grunter on no particular understanding of being one is just not a way for Sid to be. In brief, the choices here are falsehood or failure to say anything. What would be needed to block this result are ways for things to be, which one might speak of, and which do not admit of understandings. (Travis, 2006a, pp. 47-48)

The problem, then, is that in trying to state the truth-conditions of a given sentence—in purporting to specify the condition a sentence imposes on the world—either we are doing something inadequate, for we are attributing to the expression type a property of some tokens; or we simply fail to achieve the goal, for we are putting forward some words that, qua words, do not have a determinate content. That they do not have a determinate content is shown in that we can create a Travis case for them: on some understandings they are true, on some false, and so the type is neither true nor false.

In the previous quote, Travis targets Semantic Propositionalism for natural language sentences. Nonetheless, he extends his criticism to other representational forms. In chapter 1 I presented his more general criticism as targeting the following principle:

> *Truth-Conditional Compositionality for Structured Representations:* The truth-conditions of a token of a structured representation S are determined (exhausted) by features of S's type.

Now, assuming that Travis is right in his criticism, what are the prospects of having generative theories of truth-conditional content? I think we can still have theories of a certain kind. Travis cases only show that Truth-conditional Compositionality doesn't

hold for representational forms that we can specify. Moreover, when he talks about adding more and more representational features he seems to have in mind representational features similar to time, place, and perhaps standards of precision. However, this is compatible with <representational form, situation> pairs being truth-conditional compositional. The question now is whether we can have generative theories of truth-conditional content that take as input propositions (or lekta) and situations (or activities). In what follows I will argue that, although there are some restrictions to the possibility of specifying these theories, there can be systematic connections between lekta and activities and truth-conditional content and that some semantic theories (double-indexed theories) can deal with occasion-sensitivity.

The second assumption that Travis targets is Semantic Determination[9] . Semantic Determination is in place in theories that take the referent of context-sensitive expressions to be determined (not constrained) by meaning. It states that context-sensitivity in this sense is driven by semantics. If an expression is context-sensitive, then it is part of its semantics that it is. And it is its semantics that determines how the context matters. Thus, if Semantic Determination holds, then the semantics of any context-sensitive expression E encodes some kind of rule such that, given a context, it automatically picks up the value of E in that context. I think that Semantic Determination is implicit in Kaplan's definition of character: 'The character of an expression is set by linguistic conventions and, in turn, determines the content of the expression in every context' (Kaplan, 1989, p. 505). As I pointed out in chap. 2, the indexical 'I' provides a good model. The semantic content of 'I' (its character) is a rule specifying that the value of this expression on an occasion of use is the utterer. This character, given a context, automatically picks up a referent. However, we need not assume that everything that is contributed by the context is determined by meaning. In particular, the situation an utterance concerns need not be determined by the meaning of the utterance. It can simply be the situation where the conversation takes place.

In this sense, occasion-sensitivity is compatible with having a theory that matches features we abstract from particular contexts with truth-conditions. As a toy example, let us imagine a theory making use of topic of the conversation. Such a theory could have theorems as the following: 'If the conversation is about painting leaves, then the sentence 'The leaves are green' is true if and only if the salient leaves are painted green'. This theory doesn't tell us how to determine that a particular conversation is about painting leaves. It can even be agnostic about the existence of a mechanism that can

[9] Semantic Determination was introduced in chapter 2.

establish it. In order to guess what the conversation is about, we might need to pay attention to what people is doing, to previous speech, to background knowledge about how human beings behave in certain circumstances… and other things we can't even list in advance. However, once the topic of conversation has been established, the theory delivers the truth-conditions of the utterance. In this sense, the unsystematic determination of the topic is prior to the application of the theory. The proper input to the theory is a context-type, a possible topic of conversation.

This view is not uncommon. As an example, Schoubye and Stokke (2016) take the enriched proposition a given utterance expresses to be determined by the linguistic meaning of the sentence uttered (for them, sentences express minimal propositions) together with a Question Under Discussion. Contexts, on their view, contain Questions Under Discussion. The goal of the conversation is to answer them. Schoubye and Stokke give a semantic theory that takes as input minimal propositions and Questions Under Discussion and yields as output enriched propositions. It is not part of the theory to explain how is the Question Under Discussion determined (in the metaphysical and epistemic sense).

Although I think that it would be too ambitious to expect that the theory explains how is the topic of conversation–or in the case of Austinian propositions, the relevant situation or activity–determined, I think that any feature that the theory makes use of must be such that it is plausible to attribute it to the interpreter. In this respect, situations and activities are, I think, in a better position than Questions Under Discussion. Given that most conversations do not start with a question, Schoubye and Stokke take Questions Under Discussion to often be implicit. The problem with this is that it is not always easy to identify them. For example, if we go back to the green leaves example, it is easier to identify the activity that Pia's first utterance concerns (painting leaves or, in Predelli's reconstruction, taking pictures) than to identify the question that the utterance is supposed to address. We can imagine that the conversation goes as follows. Pia tells her friend that she doesn't like her maple's brown leaves. She takes some green paint, and paints them. After that, she says 'That's better. The leaves are green now'. What Question Under Discussion is she addressing? It would be odd to reconstruct the conversation as addressing the question 'What colour are the leaves painted?' or to take Pia's friend to be (implicitly) using this question in order to interpret Pia's utterance. But it doesn't seem that odd to describe both Pia and her friend to be aware that the utterance is to be interpreted in relation to the activity of painting the leaves. Explicit questions avoid this identification problem. In their explanation of the process

of completion of the sentence 'Steel is strong enough' Schoubye and Stokke imagine a conversation starting with an utterance that introduces an explicit question: 'The space shuttle must be able to carry 35 tons of cargo, endure extreme temperatures, and be capable of withstanding severe cyclonic dust storms. So, what material for the shuttle is sufficiently strong?' However, a potential problem here is that explicit questions will often be in need of pragmatic adjustment themselves (35 tons of cargo, with what standards of precision? What counts as a 'severe' dust storm, etc.?).

5.3 Natural language semantics and interpretation

In this section, I will argue, following Predelli, that occasion-sensitivity is compatible with natural language semantics in the Montagovian tradition. An advocate of occasion-sensitivity can use this in order to explain our ability to grasp the truth-conditional content of our utterances. Hence, occasion-sensitivity is not incompatible with any kind of systematicity.

Double-indexed theories are suited to deal with the context-sensitivity of natural language. Let me take as an example Lewis' framework. As Lewis (1970: 23) notes, the truth-value of a sentence (its extension) depends not only on its meaning but also on, among other things, the time and place of utterance, the speaker, the surrounding discourse, the possible world, etc. Lewis calls a package of such parameters on which extension depends 'an index'. More standardly, and because of linguistic and philosophical reasons, it is common to work with a notion of truth relativized to contexts and indices. This is usually called double-indexing[10] . Following the idea that sentential truth is relative to indices, the semantics of the sentence 'France is hexagonal' can be thought of as a function from contexts and indices that include standards of precision to truth-values. Semantics is thus made compatible with Austin's well-known example.

Now, these semantic frameworks typically deliver a notion of truth that is not relativized to indices. The reason is that it is the non-relativized notion that is of pragmatic relevance—as MacFarlane (2014, p. 53-55) argues, when we speak we typically try to

[10]See as an example Kaplan (1989) and Lewis (1980). In the kaplanian framework, the role of the context is to determine the proposition expressed .

say what is true, thereby conveying some information that might be useful for the addressee. The notion we need in order to account for this practice is truth-at-a-context, not truth-at-an-index. Lewis writes: 'a sentence s is true at context c iff s is true at c at the index of the context c' (1980, p. 88). Thus, a semantic theory might have it that, for example, the sentence 'France is hexagonal' at context c is true if and only if this sentence is true at the index determined by c. Given that different contexts of use will very likely determine different indices, not all the tokens will express the same truth-conditional content. Two tokens of the same non-indexical sentence can have different truth-values–this can be the case whenever they are evaluated against different indices.

Hence, Predelli (2004; 2005b; 2005a) argues that the shiftiness exhibited in Travis cases is compatible with natural language semantics. He uses a different terminology. According to Predelli, the formal system takes clause-index pairs as input and yields functions from points of evaluation to truth-values (intensions) as output. In his terminology, a clause is a syntactic construct and an index is a set of parameters (place, time, speaker, world). A clause-index pair can have different truth-values at different points of evaluation. Importantly, points of evaluation are not wordly conditions. One same wordly condition (some leaves being painted green) can correspond to two different points of evaluation. So, we can have two utterances u and v of the same sentence ('The leaves are green') with the same intensional profile, but two occasion of use that correspond to two different points of evaluation. Predelli writes:

> Returning to Pia's case, then, the intuitive requirement it puts forth is that u, her utterance of 'The leaves are green' during the discussion with the photographer, be evaluated as true at particular points, those reflecting our assessment of the leaves as green. In the more austere jargon of the interpretive system, what needs to be obtained is a result of truth at points k such that, in k, the object denoted by $[\text{the leaves}]_{\text{NP}}$ partakes in the extension of $[\text{is green}]_{\text{NP}}$. Moreover, as far as our intuitions go, it is also desired that v, her utterance of 'The leaves are green' during the exchange with the botanist, turns out false at points of a different type, those corresponding to our understanding of the leaves as non-green. To put it otherwise: falsehood must be obtained whenever points k' are taken into consideration, such that the *denotatum* of $[\text{the leaves}]_{\text{NP}}$ is not a member of the value of $[\text{is green}]_{\text{NP}}$ at k'. This much, of course, is perfectly consistent with the results straightforwardly provided by traditional systems. For in their approach it is by no means surprising that an indexical-free clause (paired with any index i), when evaluated with respect to a certain point, ends up being associated with a truth-value distinct from that assigned to that very clause–index pair with respect to

another point. (Predelli, 2005a, pp. 47-48)

In a similar vein, MacFarlane's proposal is to add a count-as parameter to the circumstance of evaluation (MacFarlane, 2007, 2009). One same proposition can have different truth-values when assessed against circumstances where the count-as parameter takes different values.

Double-indexed theories do not assume Semantic Propositionalism. Although they attribute intensions to sentences, they do not attribute truth-evaluable content to sentences, but to sentences-at-a-context. Moreover, they do not specify truth-conditions in the same way that Davidsonian theories do. Following these ideas, one can agree with Travis that propositions, or representational forms, fail Truth-Conditional Compositionality, for the truth-conditions of their tokens depend on something else, yet do not renounce to any sort of semantic theory. What is interesting here is that this 'something else' need not be a parameter that is added to the representational form. Rather, it can be the situation against which the proposition is evaluated. Given that, Travis' proliferation argument cannot be run again, and occasion-sensitivity is shown to be compatible with certain sort of semantic theory (by which I mean a theory of truth-conditions).

Now, are these proposals useful as an account of our ability to interpret speech? I think that, if conceived as theories of interpretation, Predelli and MacFarlane's accounts are subject to a worry that Lepore raised for model-theoretic semantics.

Lepore (1983) claims that theories that do not specify truth-conditions are insufficient as theories of understanding. His argument against model-theoretic semantics goes as follows. In Montagovian model-theoretic semantics, the truth-conditions of the Finish sentence 'Barbara sekoilee' are provided by the following theorem:

> (E1) 'Barbara sekoilee' is true in an interpretation A at a world w and a time t (in A) if and only if the extension picked out by the intension of 'Barbara' in A at w and t is a member of the extension picked out by the intension of 'sekoilee' in A at w and t.

But, Lepore notes, one could know this theorem and fail to discern what a Finish speaker says when he utters 'Barbara sekoilee'. All he would know is that 'sekoilee' is true of something named 'Barbara', but this is not enough to know what the Finish speaker would assert for a similar argument . The theory could only serve the purpose of delivering the interpretation of the sentence if it included disquotational schema as:

> (E2) The extension of 'Barbara'=Barbara.

(E3) (x) (x satisfies 'skoilee' if and only if x is confused).

Predelli writes:

> [G]iven a worldly condition w, the utterance u is true with respect to w iff j(f(w))=truth, that is, iff the intension j, when applied to a point of evaluation suitably corresponding to w, renders a result of truth. (Predelli, 2004, pp. 2122)

Again, one could know this theorem and fail to discern the content of Pia's utterances. Adding a 'count-as' parameter or distinguishing wordly conditions from points of evaluation is a way to make Travis cases compatible with natural language semantics, but it can be doubted that the resulting theory is sufficient for the hearer to interpret an utterance of 'The leaves are green', for it does not specify what counts as 'green' at the point of evaluation of the context. This is something that the interpreter needs to know in advance. So, if the theory is to be something that an interpreter could use in order to interpret a use of the sentence 'The leaves are green' it would have to include something equivalent to (E3), like (E4):

(E4) (x) (x satisfies 'is green' at C if and only if s).

Where 's' specifies what counts as 'green'. And this is something that the semantic theory will have trouble to do, because, plausibly, the only way of specifying what counts as 'green' will be to say that x is painted green, or that x is naturally green, etc., and these terms are themselves occasion-sensitive.

I think that Predelli and MacFarlane's semantic theories are insufficient as accounts of our ability to interpret speech[11] in the sense that they are not adding anything to Travis' claim that whether the sentence 'The leaves are green' is true of certain leaves on an certain occasion of use depends on whether the leaves count as 'green' on that occasion–they leave it unexplained how is the hearer supposed to grasp what counts as 'green' on that occasion.

However, we can distinguish here an ambitious and a modest goal. The ambitious goal (Davidson and Lepore's goal) would be to provide a theory that could be used to interpret speech. If Lepore is right, this theory will be one that specifies truth-conditions. The modest goal would be to identify the features that an interpreter could use in interpreting an utterance and how these features relate to truth-conditions. The modest goal

[11] I do not intend this as a criticism of Predelli or MacFarlane's theories. To my knowledge, explaining our ability to interpret speech was not their aim.

could then be seen as a (small) step in modelling our ability to grasp truth-conditions. And, more importantly for my purposes here, it could be used to show that occasion-sensitivity is compatible with a plausible account of interpretation.

An account similar to Predelli's can be used to show that occasion-sensitivity is compatible with there being systematic connections between sorts of situations and truth-conditions, and this is all we need in order to show that occasion-sensitivity does not turn interpretation into something magical. Rejecting Semantic Propositionalism casts doubt on the possibility of specifying truth-conditions, but not on the possibility of explaining our ability to interpret speech.

In order to show this, let me distinguish systematicity from specifiability. Truth-conditional content is systematic when it systematically depends on a set of features–same set of features, same truth-conditions. Now, it could be the case that truth-conditions are systematic, yet that we cannot have systematic theories of truth-conditions. For example, the set of features on which the truth-conditions of an utterance depend might be too complex for us to identify. More interestingly, there can be limits on what our theories can specify, by which I mean what we can encode or articulate in a natural language sentence. When Bezuidenhout writes that 'Truth-conditional content depends on an indefinite number of unstated background assumptions, not all of which can be made explicit' she is pointing to a problem detected by Searle. Searle notes that the truth-conditions of our utterances depend on a number of assumptions. We can make some of those assumptions explicit. The problem is that in making them explicit we bring in further assumptions (see chapter 3). As a result, we cannot specify all of those assumptions—any specification brings in new assumptions.

However, this is compatible with there being systematic connections between some features of the context or the background and truth-conditions. Let me use as an example Perry's Z-landers (Perry, 1986). Z-landers do not represent Z-land, the place where they live. The have no word or concept for it, and they do not travel. Because of this, if a Z-lander semanticist is to state the truth-conditions of an utterance of 'It's raining' he would most likely have something like the following theorem:

> (Z1) An utterance u of 'It's raining' at time t is true if and only if it is raining at t.

Z1 does not specify a place. However, when a Z-lander says 'It's raining' the content of his utterance concerns Z-land–it is true if and only if it's raining in Z-land, regardless of whether or nor it rains at other places. We could specify the truth-conditional content of his utterance with Z2:

(Z2) An utterance u of 'It's raining' at time t is true if and only if it is raining at t in Z-land.

The Z-lander sematicist, lacking a concept for Z-land, might be unable to specify the complete truth-conditional content of 'It's raining' (Z2). Despite this, there are systematic connections between certain features of the context (the place of utterance, Z-land), and the truth-conditions of the Z-landers utterances. Similarly, we cannot articulate a sentence that gets rid of all the implicit, context-dependent assumptions on which, if Searle and Travis are right, truth-conditions depend. But this does not entail that there are no systematic connections. We cannot specify an activity in a way that does not bring in further assumptions, but this does not mean that there are no systematic connections between activities and truth-conditional content.

The problem with the limits on what we can specify is not that it makes understanding mysterious, for it does not entail that there are no systematic connections, but that it casts serious doubts on the possibility of having theories specifying the features on which truth-conditions depend or the truth-condition of a sentence (or utterance).

My answer to the question about what could we know that would enable us to interpret Pia's utterance involves three elements. First, we know the literal meaning of the sentence uttered (that 'green' refers to a particular colour). Second, we are acquainted with different activities in which we use words. Third, there are systematic connections between activities with criteria of applicability for words. This approach has some similarities with Travis's. Here is Travis' view on interpretation (I quote at length):

> But then, how might we be able to do the things we in fact do in understanding words? Let us begin, yet again, with an example. [The example is: Pia and Max have received an invitation to a reception stating 'business attire'. Pia says to Max: 'Wear a tie'. The question concerns the interpretation of this utterance] Max is in his wardrobe, fingering his beloved bolo tie. Would it be wearing a tie on the understanding on which she spoke of that? Max might reason (more long-windedly than we ever need to) along these lines: In Pia's circle no one ever wears anything but a four-in-hand (and stodgy at that). Four-in-hands are all they would ever think of. They would be rather taken aback by a bolo. Pia would not want a thing like that. She would have meant: wear the kind of things her colleagues wear by way of a tie. So wearing a bolo would not be doing what she asked when she told me to wear a tie. So wearing a bolo would not be doing what she asked when she told me to wear a tie. It would not be wearing a tie. It would not be wearing a tie on the understating of doing that on which Pia spoke of it.

> Max's ability to perceive this particular fact as to what Pia said, and how he might conform to what she said, derives from two things: his worldliness; and his reasonableness (or, as one might say, sanity). His understanding of Pia's words extends just as far as those things do. His wordliness includes his knowledge of what that place and time are like, of the people and circles involved in the proceedings Pia's statement is part of. He knows their customs and habits, how things are done in those surroundings, what those people would expect, and what would surprise them. His reasonableness consists in his ability, so far as it reaches, to bring facts to bear in trains of thought where those facts are relevant, and to reason consecutively, and with some degree of imagination. So far as we know, there is no algorithm for reasonableness in this sense. But most of us have enough of it to see and appreciate the line of thought about bolo ties just rehearsed (given, of course, wordliness equal to Max's). If Max can do that, that is enough for him to perceive the particular fact which was in question as to the right way of understanding Pia's words. (Travis, 2000, pp. 209-210)

According to Travis, it is the combination of worldliness and reasonableness that enables us to interpret a given utterance. First, we need to have some knowledge about how things work—some acquaintance with receptions. Second, we, normal speakers have the ability to put this knowledge to work in the interpretation of a given utterance. Now, Travis claims that 'so far as we know, there is no algorithm for reasonableness'. The question now is: What limits would that impose?

In Travis' example, reasonableness is what enables Max to relate his knowledge about receptions and Pia's circles to the present question. It is what enables him to take the bolo tie not to fit the present purpose. As I see it, it has to do with identifying the activity that the utterance concerns and, once that is done, bringing knowledge about this activity in order to see what fits them. But if that is so, then there is room for some systematicity. Let me tell a slightly different story.

Max is acquainted with many different activities. Among them, going to informal dinners with friends, having a drink with work colleagues, going to receptions in his hometown and going to elegant receptions in Europe. Pia says: 'Wear a tie'. The meaning of 'tie' is insufficient to decide exactly which things would count as 'tie', in what cases he would be obeying the order. In order to interpret this utterance, he has to relate it to one of these activities he is acquainted with. But let's suppose that Max is aware of the surroundings of the conversation and comes to think that the utterance concerns the activity of going to elegant receptions in Europe. This activity comes with certain ways of doing things, including how to dress. Since Max is acquainted with this activ-

ity, his knowledge of it enables him to see that wearing a bolo would not be obeying the order.

What kind of systematicity does this story admit? It allows for systematically matching utterances and activities with interpretations of utterances. In other words, it is compatible with there being systematic connections between activities and application conditions for words.

I think that advocates of pragmatic approaches to truth-conditions should not abandon all forms of systematic accounts of truth-conditional content. Occasion-sensitivity motivates a rejection of Semantic Propositionalism and semantic determination. But abandoning these principles is compatible with the truth-conditions of our utterances systematically depending on the situation of utterance. An advocate of occasion-sensitivity can use this systematicity in order to explain our ability to interpret speech.

Using the idea that the satisfaction conditions of the predicates in Travis cases depend on the activity of the context, and following Predelli and MacFarlane, we can have the following schema:

> (S1) A sentence S at context C is true if and only if S is true in the activity A of C.

Or, if one prefers propositions:

> (S2) An utterance of S at context C is true if and only if the proposition expressed by S at C is true in the activity A of C.

I what follows I will make use of the notion of Austinian proposition introduced in chapter 4 (<lekton, activity> pair) in order to sketch an account of our capacity to grasp truth-conditional content compatible with occasion-sensitivity.

As I introduced in chapter 4, occasion-sensitivity calls for a form of non-indexical contextualism. What is needed in order to get a truth-value is an activity, or a practice, in which the sentence is used. Thus, the content of an utterance can be conceived as a <lekton, activity> pair, where the lekton, following Recanati, is the linguistic meaning of the sentence (plus the value of indexicals) and the activity is the type of practice the utterance concerns. One same sentence used relatively to different activities can have different truth-values. In particular, the activity determines the criteria of applicability of predicates (what counts as 'green', etc.). We can have schema of the kind of S1 and S2, such as S3:

> (S3) An utterance of a sentence S at context C is true if and only if the lekton expressed by S at C is true for the activity A of C.

This schema plausibly models part of our linguistic ability. When we are presented with the green leaves examples, we take the utterance as true (or false) because the state of the leaves is one that counts as 'green' in the activity described—the activity constrains the satisfaction conditions of 'green'. The schema captures two features that are needed to interpret speech: grasping the proposition expressed (what, following Recanati, I have called the lekton, or simply the meaning plus the referents of indexicals, demonstratives, etc.) and the activity of the context.

The knowledge that an interpreter needs in order to grasp the satisfaction conditions of 'is green' on an occasion of use includes the conventional meaning of this expression (that it refers to a certain colour, let's say) together with an acquaintance with the activity or practice in which it is being used. This activity constrains what counts as 'green', and being familiar with it enables us to see how the use of the expression is constrained. This is not linguistic knowledge, but rather knowledge about how are things classified for the purposes of an activity. Given this, it is plausible to think that what enables speakers to grasp truth-conditions is knowledge of the systematic connections between lekta and activities, and truth-conditions. Part of what we learn when we get acquainted with a new activity has to do with how to classify things within that activity, how it constrains the satisfaction conditions of words.

The question I have addressed here is: how do we go about interpreting utterances, given the shiftiness of natural language? The sketch I propose is the beginning to an answer. It uses information it is plausible to attribute to speakers, such as the recognition of an activity. It is systematic in that it assumes systematic relations between lekta and activities, and truth-conditions. In this way, it explains how is it that we are able to interpret utterances: the reason is that there are systematic relations between activities and satisfaction conditions, we are acquainted with many such activities, and, on an occasion of use, are aware of the relevant activity.

There are two things the account sketched is silent about. First, it does not explain in virtue of what we are able to recognise the activity an utterance concerns, or even what determines the relevant activity. Second, it does not tell what activities there are or could be.

As to the first question, I think that the information about the activity a conversation concerns can be included in the common ground of the conversation. Following Stalnaker (2002; 2014) we can think of the context of a conversation as a body of infor-

mation that is presumed to be shared among the interlocutor. The common ground includes the content of the utterances of the conversation and their presuppositions, but it can also include information not obtained through language—as the information we obtain through perception. Stalnaker writes that '[T]he information that a context set models includes all the information that is a resource for the interpretation of context-dependent expressions' (Stalnaker, 2014, p. 24). Recently, García-Carpintero has argued that the notion of common ground should include non-doxastic attitudes as Questions Under Discussion (García-Carpintero, 2015). Similarly, what I am suggesting here it that it should include information about what activity the conversation concerns. Thus, when faced with an utterance the interpreter will be in a position to use the information about which activity the utterance concerns, connect it with its knowledge of the activity, and interpret the utterance. Following Searle, we can think of his knowledge of the activity as partly based on the interpreter's background—a set of non-representational assumptions, skills, capacities, practices and habits (Searle, 1983).

Concerning the second, what activities we engage in are external to the systematic mechanism linking them with truth-conditions. Moreover, particular speakers might not be acquainted with all the activities in which a given word is used. In this sense, we probably have 'partial theories of interpretation'[12] .

5.4 Locally expressible truth-theoretic semantics

In the previous section I have argued that occasion-sensitivity is compatible with there being systematic connections between sorts of situations (or activities) and truth-conditional content. Hence, it does not necessarily turn interpretation into something magical. In this section I will explore to what extent it is compatible with theories that specify truth-conditions and argue that we can have instances of the T-schema on the condition that they rely on default understandings.

[12] This opens the door for an approach along the lines of Davidson (1986), where speakers constantly adjust their theories of interpretations. Davidson seems to hold that there is no such a thing as 'the' theory of interpretation for a given language, but provisional theories speakers use in conversations. Similarly, one can think of provisional theories matching specific activities that are created online (but related to pre-existing activities) and lekta with truth-conditions.

The aim of the previous section was modest. My goal there was to show that occasion-sensitivity is compatible with systematicity and that, as a consequence, it does not turn interpretation into something magical. Moreover, occasion-sensitivity can be integrated in natural language semantics. However, one might want to go for a more ambitious aim, such as providing a theory that an interpreter could use to interpret Pia's utterances.

In view of Lepore's argument, one might want a semantic theory that specifies what an utterance of a given sentence would say instead of merely delivering S3. As an example, let us consider a theory delivering conditional T-schema (I ignore the complexities introduced by the definite description and treat it as a proper name):

> (T1) If an utterance u of the sentence 'The leaves are green' concerns the activity of decorating the garden, then u is true if and only if the leaves are green for the purpose of decorating the garden.
>
> (T2) If an utterance u of the sentence 'The leaves are green' concerns the activity of doing a study on green leaf chemistry, then u is true if and only if the leaves are green for the purpose of doing a study on green leaf chemistry.

There are reasons to be suspicious about the prospects of such a theory. The problem lies in that we are using natural language as a metalanguage. The expressions that are used ('decorating the garden', etc.) can be understood in different ways. What does it mean 'for the purposes of decorating the garden'? Is a leaf that has been poorly covered with green paint count as 'green for the purposes of decorating the garden'? Because of this, there are reasons to doubt the possibility of stating the kind of theory that Davidson envisioned. Nonetheless, I will argue that T1 and T2 need not be incorrect.

I think that we have the resources to admit what I will call 'locally expressible theories'. The key is that advocates of truth-conditional pragmatics typically accept that sentences can be used to report what a given utterance says—even though they claim that any sentence is open to different understandings and that sentences do not encode truth-conditional content.

Let me draw an analogy with speech reports. Imagine (a modification of the green leaves example) that a botanist working on a study on chlorophyll says 'I need some green leaves'. A colleague can correctly report that to Pia, who happens to be painting the leaves on her tree, by saying 'He said that he needs green leaves for a study on chlorophyll'. The sentence 'He needs green leaves for a study on chlorophyll' does not fully

capture (in the sense of encoding) what the botanist said: as a sentence, it is open to new Travis cases, whereas the botanist uttered a truth-evaluable, definite content. For example, the sentence doesn't tell us if leaves containing artificially injected chlorophyll would fit the botanist's utterance. However, as a report offered to Pia, it can be seen as true (or perhaps as true enough). It relies on a default understanding of what counts as 'green leaves for a study on chlorophyll', but here that is not a problem. We know that, as a general rule, science studies natural properties, that only leaves that are naturally green have chlorophyll, etc. The interpretation of the report can rely on this common knowledge (not articulated in the sentence).

Default understandings are a variety of what Bach (1995) calls standardization. Suppose that most of the times we use colour predicates to describe the colour of someone's hair what we thus describe is how his hair looks in normal illumination conditions. In other words, the most common activity in which we use colour predicates as applied to hair is to describe how the hair of some person looks. We have, of course, other practices in which we use colour predicates as applied to hair: we talk about someone's original hair colour, about hair colours in poorly illuminated rooms, and so on. But if one practice is the most common, we can interpret by default utterances about hair colour as concerning that practice—that is, unless we have evidence to the contrary. The hearer can rely on precedent. In absence of a context, for example, we rely on that most of the times colour predicates are used to describe someone's hair they are used to describe the apparent (as opposed to the original) colour and use this as the default interpretation. So if I start a conversation by saying 'Have you met Paul, the red-haired man who works with Lisa?', my addressee will probably take Paul to have hair that looks red on normal illumination conditions—not necessarily originally red, etc[13] .

Moreover, in many cases a rough approximation is good enough. I ask you what Pia said. You tell me that she said that she has green leaves for decorating her garden. There are different ways of understanding these words, and the set of things that count as green in them need not fully coincide. However, I get a rough idea of what Pia said—one according to which 'green' has to do with how things look. This rough idea can be good enough.

Thus, even if we cannot have standard truth-theoretic semantics, we can, in the

[13]One might object that default understandings are the semantic values of words. It could be objected that if we, by default, interpret colour predicates in a certain way, then that is the semantic value of colour predicates. But default understandings are just the most common ones. We might have other understandings. And being (statistically) the most common interpretation is not equivalent to being the semantic value.

Davidsoinian spirit[14] , have theories that redescribe what an utterance says by having recourse to default understandings. We could have a theory delivering theorems as T1 and T2. This theory will deliver something that tells us what the utterance says, even if it does so by relying on default understandings of the terms involved. Strictly speaking, any theory will deliver a sentence that, as a sentence, is subject to Travis cases. Other parts of the theory might be subject to Travis cases as well. A number of implicit assumptions will always remain implicit. Yet, the theory might shed some light about how we go about interpreting speech given the shiftiness of natural language—we can use it to truly report what was said in a given utterance, to compare it to what would happen in other context-types, and to capture how could a speaker arrive at that particular interpretation. He knows that if the utterance concerns the activity of decorating the garden (and plausibly interlocutors know the activities that utterances concern), then it is true if and only if the leaves are green for the purposes of decorating the garden.

These theories will be only locally expressible[15] in the sense that we cannot state a theory that provides the truth-conditions of any potential sentence of the language without assuming that the metalanguage is understood by default. As Searle would put it, the right hand side of the biconditional comes with an indefinite number of implicit assumptions. The theory makes explicit how truth-conditions systematically depend on some features of the context. But, since it uses natural language, it cannot do so in a way that does not bring in new assumptions, or that does not leave some questions unanswered about the satisfaction conditions of the expressions it uses. No theory that uses natural language as a metalanguage, as Davidsonian theories do, can do without default understandings. So we can only give the truth-conditions of some parts of the language at a time, relying on other bits of the language being understood by default.

How do these locally expressible theories fare in view of Travis' criticism against the T-schema? As a quick reminder, the problem with 'The sentence 'The leaves are green' is true if and only if the leaves are green' is that either the right hand side of the biconditional does not presuppose an understanding of the words used, in which case it is not stating a truth-condition, or that it does, in which case it is wrongly attributing

[14] 'We interpret a bit of linguistic behavior when we say what a speaker's words mean on an occasion of use. The task may be seen as one of redescription. We know that the words 'Es schneit' have been uttered on a particular occasion and we want to redescribe this uttering as an act of saying that it is snowing. What do we need to know if we are to be in a position to redescribe speech in this way, that is, to interpret the utterances of a speaker?' (Davidson, 1974, p. 309)

[15] I take the expression local (vs. global) expressibility from (Recanati, 2002), a discussion about Searle's view on the expressibility of thought.

the sentence a property of a use of the sentence. I think the schema presented here escapes this criticism. The left hand side of the biconditional in T1 and T2 includes the activity. It is not matching sentences with presupposed understandings, but sentences and specified activities with default understandings. The words stating the activity on the left hand side bear the same understanding as the words stating the truth-conditions of the right. Thus, it is not wrongly attributing a sentence a property of some uses of the sentence and, relying on default understandings, it manages to state a truth-condition.

5.5 Conclusions

The aim of this chapter was to show that occasion-sensitivity does not turn understanding into something mysterious. Natural language semantics (theories about the truth-conditions of sentences-at-a-context) can be used to model our ability to interpret speech. However, Travis and other advocates of truth-conditional pragmatics have called certain approaches into question. In particular, Travis holds that there can be no generative theory of understandings.

Because of their insistence on the difficulty, or even impossibility, of systematizing the role of context, advocates of truth-conditional pragmatics run the risk of turning our ability to grasp the truth-conditions of our utterances into something magical. However, making use of the notion of Austinian proposition introduced in chapter 4, I have argued that occasion-sensitivity is compatible with there being systematic connections between lekta and activities, and truth-conditions. Following Predelli, I have also argued that occasion-sensitivity is compatible with double-indexed theories. Moreover, I have argued that occasion-sensitivity is compatible with theories that specify truth-conditions in case these theories are seen as relying on default understandings.

Conclusions

I have explored and defended the hypothesis that natural language is occasion-sensitive. I have done so by focusing on a number of examples involving different sorts of predicates, with special attention to colour predicates. To hold that language is occasion-sensitive is to hold that it exhibits semantic underdeterminacy. I have distinguished two notions of underdeterminacy: Type and Token-Underdeterminacy. A non-indexical sentence S is type-underdetermined if and only if there are tokens of it that have distinct truth-values. A token of a sentence S is token-underdetermined if and only if for some possible state of affairs its truth-value is indeterminate. I have argued that natural language is underdetermined in both senses and that there are no reasons to take mental representations to be any different in these respects. Because of this, I have endorsed a form of truth-conditional pragmatics in which the content of an utterance is not identified with a mental representation or structured proposition with intrinsic truth-conditions but, instead, with an Austinian proposition including a lekton and an activity. This conception of utterance content can be used in accounting for our ability to grasp truth-conditional content, on the basis of the systematic connections between lekta and activities, and truth-conditions.

In chapter 1 I have reconstructed and examined Travis' targets and arguments. I have taken the outcome of his critique against truth-conditional semantics to be the claim that no specifiable content (no content that can be articulated in natural language) has intrinsic truth-conditions. The justification for a such a claim is provided by Travis cases and by the proliferation argument. Moreover, Travis generalizes his underdeterminacy claim to non-linguistic representations, as structured propositions, thus targeting a general principle of Truth-conditional Compositionality for structured representational forms. The basis for such a generalization is that, given that in natural language representational structures are compatible with a variety of truth-conditions, there are no reasons to think that other representational systems will behave differently. Nonetheless, I have argued that some of Travis' positive remarks suggest a picture in

which the objectivity of content is undermined. I have concluded that an advocate of occasion-sensitivity needs to provide a non-standard notion of content.

The goal of chapters 2 and 3 was to defend occasion-sensitivity against several attempts to deactivate it (minimalism, indexicalism) or to restrict it to natural language (Fodor's arguments and Carston's version of truth-conditional pragmatics). In chapter 2 I have argued that advocates of minimal propositions find themselves in an unstable position. In order to secure minimal propositions, they need to dismiss common reactions and intuitions triggered by Travis cases. But then we loose all warrant that we will ever find out what the literal truth-value of our utterances is (whether, for example, the meaning of the predicate 'is green' is literally true or literally false of painted leaves). On the other hand, I have argued that indexicalism and similar proposals do not manage to secure the claim that semantics (i.e., the properties of the expression type) determines (and not merely constrains) truth-evaluable content.

This opens doors for truth-conditional pragmatics. In chapter 3 I have argued against a certain version of truth-conditional pragmatics–namely, against the view that it is a pragmatic matter to identify which (non-underdetermined) mental representation a given utterance expresses. I have done so by arguing that the productivity argument fails to show that mental representations are Truth-condition Compositional and that the argument from equivocation relies on a false premise. Moreover, I have argued that *ad hoc* concepts should be expected to behave as lexically encoded concepts. In this chapter I have also introduced the notion of Token-Underdeterminacy and argued that both linguistic and mental representations are token-underdetermined (in addition to being type-underdetermined).

In chapter 4 I have introduced a situationalist approach to occasion-sensitivity. Partly following Corazza and Dokic, and Recanati, I have modelled utterance content as a pair including a lekton and an activity. This approach has several advantages. First, it is compatible with Type- and Token-Underdeterminacy. Second, on the assumption that we can have Austinian propositions with different granularity, it is compatible with both having contents that are closely tied to the context of use and less context-bound contents that can be shared across contexts. Third, it escapes the non-objectivity threat, for here it is the activity (rather than our judgements about which speech acts are reasonable) that determines, for example, what counts as 'green'.

In chapter 5 I have explored the possibility of having systematic theories of truth-conditions. The reason for wanting such theories is that they can be used in explaining our ability to interpret speech. The reason for doubting that we can have them is that

some advocates of truth-conditional pragmatics have raised doubts concerning the possibility of amening truth-conditions to systematic treatment. Against this, I have held that that there can be systematic connections between lekta and activities, and truth-conditions. I have argued that Travis and Searle's critique only cast doubt on the possibility of having a theory that states (specifies, articulates in natural language) these connections, not on the existence of such systematicity. This systematicity can be used to explain our ability to grasp truth-conditional content. Moreover, I have argued, following Predelli, that occasion-sensitivity is compatible with double-indexed theories in the Lewisian tradition. I have also argued that occasion-sensitivity can be made compatible with theories that specify truth-conditions on the assumptions that these specifications bear interpretations by default.

Bibliography

Austin, J. L. (1950). Truth, *Proceedings of the Aristotelian Society. Supplementary volume* 24(1): 111–128.

Austin, J. L. (1975). *How to Do Things with Words*, Clarendon Press.

Bach, K. (1994). Conversational Impliciture, *Mind & Language* (2): 124–162.

Bach, K. (1995). Standardization vs. Conventionalization, *Linguistics and Philosophy* (18): 677–686.

Bach, K. (2006). The excluded middle: Semantic minimalism without minimal propositions, *Philosophy and Phenomenological Research* 73(2): 435–442.

Barwise, J. and Etchemendy, J. (1987). *The liar: An essay on truth and circularity*, Oxford University Press.

Barwise, J. and Perry, J. (1983). *Situations and Attitudes*, MIT Press.

Belleri, D. (2014a). *Semantic under-determinacy and communication*, Springer.

Belleri, D. (2014b). You can say what you think: vindicating the effability of our thoughts, *Synthese* 191(18): 4431–4450.

Berg, J. (2002). Is semantics still possible?, *Journal of Pragmatics* 34(4): 349–359.

Bezuidenhout, A. (1997). The communication of de re thoughts, *Noûs* 31(2): 197–225.

Bezuidenhout, A. (2002). Truth-Conditional Pragmatics, *Philosophical Perspectives* 16: 105–134.

Borg, E. (2004). *Minimal Semantics*, Oxford University Press.

Borg, E. (2012). *Pursuing Meaning*, Oxford University Press.

Cappelen, H. and Lepore, E. (2005). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*, Blackwell.

Cappelen, H. and Lepore, E. (2006). Shared content, *in* E. Lepore and B. C. Smith (eds), *The Oxford Handbook of Philosophy of Language*, Oxford University Press, pp. 1020–1055.

Carnap, R. (1947). *Meaning and Necessity: A Study in Semantics and Modal Logic*, University of Chicago Press.

Carruthers, P. (1996). *Language, Thought, and Consciousness*, Cambridge University Press.

Carston, R. (1988). Implicature, Explicature, and Truth-Theoretic Semantics, *in* R. Kempson (ed.), *Mental representations: The interface between language and reality*, Cambridge University Press, pp. 155–181.

Carston, R. (2002). *Thoughts and utterances: The pragmatics of explicit communication*, John Wiley & Sons.

Carston, R. (2008a). Linguistic communication and the semantics/pragmatics distinction, *Synthese* 165(3): 321–345.

Carston, R. (2008b). Review Minimal Semantics by Emma Borg, *Mind & Language* 23(3): 359–367.

Carston, R. (2016). The heterogeneity of procedural meaning, *Lingua* (175-176): 154–166.

Chomsky, N. (1965). *Aspects of the Theory of Syntax*, MIT Press.

Chomsky, N. (2000). *New Horizons in the Study of Language and Mind*, Cambridge University Press.

Clapp, L. (2012a). Indexical Color Predicates: Truth Conditional Semantics vs. Truth Conditional Pragmatics, *Canadian Journal of Philosophy* 42(2): 71–100.

Clapp, L. (2012b). Is even thought compositional?, *Philosophical Studies* 157(2): 299–322.

Clapp, L. (2012c). Three Challenges for Indexicalism, *Mind & Language* 27(4): 435–465.

Conant, J. (1998). Wittgenstein on Meaning and Use, *Philosophical Investigations* 21(3): 222–250.

Corazza, E. (2007). Contextualism, minimalism, and situationalism, *Pragmatics & Cognition* 15(1): 115–137.

Corazza, E. and Dokic, J. (2007). Sense and Insensibility: Or Where Minimalism Meets Contextualism, *in* G. Preyer and G. Peter (eds), *Context-Sensitivity and Semantic Minimalism. New Essays on Semantics and Pragmatics*, Oxford University Press, pp. 169–193.

Corazza, E. and Dokic, J. (2012). Situated minimalism versus free enrichment, *Synthese* 184(2): 179–198.

Davidson, D. (1967). Truth and Meaning, *Synthese* 17(3): 304–323.

Davidson, D. (1973). Radical interpretation, *Dialectica* 27(3-4): 313–328.

Davidson, D. (1974). Belief and the basis of meaning, *Synthese* 27(3-4): 309–323.

Davidson, D. (1984). *Inquiries Into Truth And Interpretation*, Oxford University Press.

Davidson, D. (1986). A nice derangement of epitaphs, *in* E. Lepore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, Blackwell, pp. 433–446.

Davies, A. (2011). Occasion-Sensitivity: Selected Essays, by Charles Travis, *Disputatio* IV(31): 309–315.

Davies, A. (2014). Off-Target Responses to Occasion-Sensitivity, *Dialectica* 68(4): 499–523.

DeRose, K. (1992). Contextualism and Knowledge Attributions, *Philosophy and Phenomenological Research* 52(4): 913–929.

Dever, J. (2014). Formal semantics, *in* M. García-Carpintero and M. Kölbel (eds), *The Bloomsbury companion to the philosophy of language*, Bloomsbury Publishing, pp. 47–83.

Egan, A., Hawthorne, J. and Weatherson, B. (2005). Epistemic modals in context, *in* G. Preyer and G. Peter (eds), *Contextualism in Philosophy*, Oxford University Press, pp. 131–170.

Fodor, J. (2001). Language, Thought and Compositionality, *Mind & Language* 16(1): 1–15.

Fodor, J. A. (2003). *Hume variations*, Cambridge University Press.

Frege, G. (1956). The thought: A logical inquiry, *Mind* 65(259): 289–311.

Frege, G. (2010). On sense and reference, *in* D. Byrne and M. Kölbel (eds), *Arguing About Language*, Routledge, pp. 36–56.

García-Carpintero, M. (2015). Contexts as Shared Commitments, *Frontiers in Psychology* 6(December): 1–13.

Gauker, C. (2003). *Words Without Meaning*, MIT Press.

Gauker, C. (2008). Zero tolerance for pragmatics, *Synthese* 165(3): 359–371.

Giberman, D. (2016). Moving parts: A new indexical treatment of context-shifting predication, *Synthese* 193(1): 95–124.

Glüer, K. (2011). *Donald Davidson: A short introduction*, Oxford University Press.

Grice, P. (1975). Logic and conversation, *Syntax and semantics* 3: 41–58.

Grice, P. (1991). *Studies in the Way of Words*, Harvard University Press.

Hansen, N. (2011). Color adjectives and radical contextualism, *Linguistics and Philosophy* 34(3): 201–221.

Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar*, Blackwell.

Higginbotham, J. (1988). Contexts, models, and meanings: a note on the data of semantics, *in* R. Kempson (ed.), *Mental representations: The interface between language and reality*, Cambridge University Press, pp. 29–48.

Kaplan, D. (1989). Demonstratives: an essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals, *in* J. Almog, J. Perry and H. Wettstein (eds), *Themes From Kaplan*, Oxford University Press, pp. 481–563.

Kennedy, C. and McNally, L. (2010). Color, context, and compositionality, *Synthese* 174(1): 79–98.

Kölbel, M. (2008). Introduction: Motivations for relativism, *in* M. García-Carpintero and M. Kölbel (eds), *Relative Truth*, Oxford University Press, pp. 1–40.

Korta, K. and Perry, J. (2007). How to say things with words, *in* S. L. Tsohatzidis (ed.), *John Searle's Philosophy of Language: Force, Meaning, and Mind*, Cambridge University Press.

Korta, K. and Perry, J. (2011). *Critical pragmatics: An inquiry into reference and communication*, Cambridge University Press.

Larson, R. and Segal, G. (1995). *Knowledge of meaning*, MIT Press.

Lepore, E. (1983). What model theoretic semantics cannot do?, *Synthese* 54(2): 167–187.

Lewis, D. (1970). General semantics, *Synthese* 22(1-2): 18–67.

Lewis, D. (1979). Scorekeeping in a language game, *Journal of Philosophical Logic* 8(1): 339–359.

Lewis, D. (1980). Index, Context, and Content, *Philosophy and Grammar* (2): 79–100.

MacFarlane, J. (2007). Semantic Minimalism and Nonindexical Contextualism, *in* G. Preyer and G. Peter (eds), *Context-Sensitivity and Semantic Minimalism. New Essays on Semantics and Pragmatics*, Oxford University Press.

MacFarlane, J. (2009). Nonindexical contextualism, *Synthese* 166(2): 231–250.

MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*, Oxford University Press.

Martínez-Manrique, F. and Vicente, A. (2004). Overhearing a sentence: Recanati and the cognitive view of language, *Pragmatics & Cognition* 12(2): 219–251.

McDowell, J. (1984). Wittgenstein on following a rule, *Synthese* 58(3): 325–363.

Partee, B. (1979). Semantics—mathematics or psychology?, *in* R. Bäuerle, U. Egli and A. von Stechow (eds), *Semantics From Different Points of View*, Springer Verlag, pp. 1–14.

Perry, J. (1986). Thought Without Representation, *Proceedings of the Aristotelian Society* 60(1986): 137–151.

Perry, J. (2001). *Reference and reflexivity*, CSLI Stanford.

Pietroski, P. (2003). The character of natural language semantics, *in* A. Barber (ed.), *Epistemology of Language*, Oxford University Press, pp. 217–256.

Predelli, S. (2004). Semantic contextuality, *Journal of Pragmatics* 36(12): 2107–2123.

Predelli, S. (2005a). *Contexts: Meaning, Truth, and the Use of Language*, Clarendon Press.

Predelli, S. (2005b). Painted leaves, context, and semantic analysis, *Linguistics and Philosophy* 28(3): 351–374.

Recanati, F. (1989). The Pragmatics of What is Said, *Mind & Language* 4(4): 295–329.

Recanati, F. (2001). What Is Said, *Synthese* 128(1-2): 75–91.

Recanati, F. (2002). The limits of expressibility, *in* B. Smith (ed.), *John Searle*, Cambridge University Press, pp. 189–213.

Recanati, F. (2004). *Literal Meaning*, Cambridge University Press.

Recanati, F. (2006a). Crazy minimalism, *Mind and Language* 21(1): 21–30.

Recanati, F. (2006b). Moderate relativism, *in* M. García-Carpintero and M. Kölbel (eds), *Relative Truth*, Oxford University Press, pp. 41–62.

Recanati, F. (2007). *Perspectival Thought: A Plea for (Moderate) Relativism*, Clarendon Press.

Recanati, F. (2010). *Truth-Conditional Pragmatics*, Oxford University Press.

Recanati, F. (2017). From Meaning to Content, *in* D. Ball and B. Rabern (eds), *The Science of Meaning*, Oxford University Press.

Rothschild, D. and Segal, G. (2009). Indexical predicates, *Mind and Language* 24(4): 467–493.

Sainsbury, M. (2001). Two ways to smoke a cigarette, *Ratio* 14(4): 386–406.

Sainsbury, M. (2008). Fly swatting: Davidsonian truth theories and context, *in* M. C. Amoretti and N. Vassallo (eds), *Knowledge, Language, and Interpretation. On the Philosophy of Donald Davidson*, Ontos Verlag, pp. 33–48.

Schoubye, A. J. and Stokke, A. (2016). What is said?, *Noûs* 50(4): 759–793.

Searle, J. (1978). Literal Meaning, *Erkenntnis* 13: 207–224.

Searle, J. (1980). The background of meaning, *in* J. R. Searle, F. Kiefer and M. Bierwisch (eds), *Speech act theory and pragmatics*, Springer, pp. 221–232.

Searle, J. (1983). *Intentionality*, Oxford University Press.

Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition*, Blackwell.

Stalnaker, R. (2002). Common ground, *Linguistics and Philosophy* 25(5-6): 701–721.

Stalnaker, R. (2014). *Context*, Oxford University Press.

Stanley, J. (2000). Context and logical form, *Linguistics and Philosophy* 23(1985): 391–434.

Stanley, J. and Szabó, Z. G. (2000). On Quantifier Domain Restriction, *Mind & Language* 15(2&3): 219–261.

Szabó, Z. G. (2001). Adjectives in context, *in* R. M. Harrish and I. Kenesei (eds), *Perspectives on Semantics, Pragmatics, and Discourse*, John Benjamins, pp. 119–46.

Travis, C. (1981). *The True and the False: The Domain of the Pragmatic*, Benjamins.

Travis, C. (1989). *The Uses of Sense: Wittgenstein's Philosophy of Language*, Oxford University Press.

Travis, C. (1991). Annals of Analysis, *Mind* 100(2): 237–264.

Travis, C. (1996). Meaning' s Role in Truth, *Mind* 105(419): 451–466.

Travis, C. (1997). Pragmatics, *in* B. Hale and C. Wright (eds), *A Companion to the Philosophy of Language*, Blackwell, pp. 87–107.

Travis, C. (2000). *Unshadowed Thought: Representation in Thought and Language*, Harvard University Press.

Travis, C. (2006a). Insensitive semantics, *Mind and Language* 21(1): 39–49.

Travis, C. (2006b). *Thought's footing: a theme in Wittgenstein's Philosophical investigations*, Oxford University Press.

Travis, C. (2008a). *Occasion-Sensitivity: Selected Essays*, Oxford University Press.

Travis, C. (2008b). On what is strictly speaking true, *Occasion-Sensitivity. Selected Essays*, Oxford University Press, pp. 19–64.

Travis, C. (2010a). *Objectivity and the Parochial*, Oxford University Press.

Travis, C. (2010b). To represent as so, *Objectivity and the Parochial*, Oxford University Press.

Vicente, A. (2012). On Travis cases, *Linguistics and Philosophy* 35(1): 3–19.

Vicente, A. (2015). The green leaves and the expert: Polysemy and truth-conditional variability, *Lingua* 157: 54–65.

Waismann, F. (1951). Verifiability, *in* G. Ryle and A. Flew (eds), *Logic and Language (First Series): Essays*, Vol. 19, Blackwell, pp. 117–44.

Wieland, N. (2010a). Context Sensitivity and Indirect Reports, *Philosophy and Phenomenological Research* LXXXI(1): 40–48.

Wieland, N. (2010b). Minimal propositions and real world utterances, *Philosophical Studies* 148(3): 401–412.

Wiggins, D. (1997). Meaning and truth conditions: from Frege's grand design to Davidson's, *in* B. Hale and C. Wright (eds), *A Companion to the Philosophy of Language*, Blackwell.

Wittgenstein, L. (1994). *Tractatus Logico-Philosophicus (Trans. Pears and McGuinness)*, Routledge.

Wittgenstein, L. (2009). *Philosophical Investigations, 4th Edition (Trans. Hacker and Schulte)*, Wiley-Blackwell.

Wright, C. (1984). Kripke's account of the argument against private language, *The Journal of Philosophy* 81(12): 759–778.

Wright, C. (2001a). *Rails to Infinity: Essays on Themes From Wittgenstein's Philosophical Investigations*, Harvard University Press.

Wright, C. (2001b). Rule-following, objectivity and the theory of meaning, *in* C. Wright (ed.), *Rails to Infinity: Essays on Themes From Wittgenstein's Philosophical Investigations*, Harvard University Press.