

On Travis cases

Agustin Vicente

Linguistics and Philosophy

ISSN 0165-0157

Volume 35

Number 1

Linguist and Philos (2012) 35:3-19

DOI 10.1007/s10988-012-9109-x



Your article is protected by copyright and all rights are held exclusively by Springer Science+Business Media B.V.. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your work, please use the accepted author's version for posting to your own website or your institution's repository. You may further deposit the accepted author's version on a funder's repository at a funder's request, provided it is not made publicly available until 12 months after publication.

On Travis cases

Agustin Vicente

Published online: 8 May 2012
© Springer Science+Business Media B.V. 2012

Abstract Charles Travis has been forcefully arguing that meaning does not determine truth-conditions for more than two decades now. To this end, he has devised ingenious examples whereby different utterances of the same prima facie non-ambiguous and non-indexical expression type have different truth-conditions depending on the occasion on which they are delivered. However, Travis does not argue that *meaning* varies with circumstances; only that truth-conditions do. He assumes that meaning is a stable feature of both words and sentences. After surveying some of the explanations that semanticists and pragmaticians have produced in order to account for Travis cases, I propose a view which differs substantially from all of them. I argue that the variability in the truth-conditions that an utterance type can have is due to meaning facts alone. To support my argument, I suggest that we think about the meanings of words (in particular, the meanings of nouns) as rich conceptual structures; so rich that the way in which a property concept applies to an object concept is not determined.

Keywords Meaning · Truth-conditions · Occasion-sensitivity · Conceptual semantics · Semantic knowledge · World knowledge

1 Introduction

Charles Travis has been forcefully arguing that meaning does not determine truth-conditions for more than two decades now. To this end, he has devised ingenious

A. Vicente (✉)
Department of Logic and Philosophy of Science, University of the Basque Country, UPV/EHU,
Avda. Tolosa, 20080 San Sebastián, Spain
e-mail: agustin_vicente@ehu.es

A. Vicente
Ikerbasque: Basque Foundation for Science, Alameda Urquijo, 36-5, Plaza Bizkaia, 48011 Bilbao,
Spain

examples where different utterances of the same prima facie non-ambiguous and (in the given context) indexically non-problematic expression type have different truth-conditions depending on contextual factors, or as he would have it, depending on the occasion on which they are uttered.¹ This has led him to argue that truth-conditions are *occasion sensitive*. Travis's ingenious examples can be found almost everywhere in recent debates in the philosophy of language, or at least in those debates that focus on what kind of content (if any) is carried by sentences in context. Here I want to formulate a rough taxonomy of the responses given to Travis's challenges and to propose my own diagnosis. The first part of the paper briefly introduces some of Travis's examples and states what can be taken to be his argument against meaning determining truth-conditions. The second section presents the different reactions that the examples and the argument have motivated. In the final sections, I argue for a view which differs substantially from all of those reactions.

2 The examples and the argument

Travis's examples are very well known by now. They typically involve very simple non-ambiguous and indexically non-problematic utterances, that turn out to have different truth-conditions on different occasions, or in different contexts or scenarios. Probably the example that has attracted most attention is that of the green leaves (Travis 1996). Travis asks us to think of two different occasions on which the expression type 'the leaves are green' is uttered by someone called 'Pia'. On the first occasion, Pia is talking to a photographer who needs some green leaves for her picture. Pia finds some leaves that have been painted green, though they are originally red. Still, in that context, her utterance of 'the leaves are green' is judged to be true. However, then along comes a botanist looking for green leaves. Referring to the very same leaves, Pia says again: 'the leaves are green'. This time her utterance is judged to be false. Nothing has changed either regarding the leaves or in the meaning of the utterance, and yet, what was true before is now false.

In another example, we are asked to imagine a kettle that has been burnt all over (Travis 1985). Upon finding it, Pia says, alarmed: 'that kettle is black!' We seem to judge that she has told the truth. But the days pass, and Pia and her boyfriend see the kettle in a shop. Again, Pia tells him: 'that kettle is black'. Her boyfriend, however, responds: 'no, it seems it was burnt and they have not bothered to clean it up. See', he adds passing his finger across its surface, 'the kettle is silver coloured'. In this scenario, we are invited to think that what Pia said was false. A third example, developed at length in Travis (2000), revolves around the utterance 'the ink is blue'. Here Travis claims that there are two ways in which the ink can be judged to be blue, and so, that there are two different sets of truth-conditions that an utterance of 'the ink is blue' can have, depending on the occasion when it is uttered. In one context, the ink can properly be said to be blue if and only if it produces blue writing when used; but in another context, it can only be said that the ink is

¹ By 'indexically non-problematic' (without qualification, from now on), I mean either that the expression type contains no indexicals at all or that the value of the indexicals within the expression type do not vary with the relevant contextual variations (i.e., with the relevant scenarios). The only expression type considered here will be sentential expressions.

blue if and only if it looks blue (irrespective of whether it produces blue writing when used). Perhaps there are more possible readings of ‘the ink is blue’, but Travis seems content with these two.

Famously, Travis concludes from all this that “[w]hat words mean plays a role in fixing when they would be true; but not an exhaustive one” (1996, p. 451). That is, he takes his examples to show that, though meaning has to constrain truth-conditions in some way, it falls short of determining them. It is not easy to see why Travis interprets his examples in this way. In principle, he could take them to show that both word meaning and sentence meaning are occasion sensitive or unstable. That is, he could use his examples to argue for the view that only utterances (of words and of sentences) are bearers of meaning, denying that there is such a thing as the meaning of a sentence type or of a word type. If he went down that path, he could preserve the idea that the meaning of a token of a sentence determines (or is identical to) its truth-conditions. Yet, he does not want to go that way. Rather, he thinks that two factors intervene in the determination of the truth-conditions of an utterance: the meaning and the circumstances. It is certainly difficult to get a clear idea about what meaning consists of according to Travis, but he apparently assumes that meaning is quite a stable feature of words, i.e., that there is something that could be said to be the word’s type meaning (see the quote above and also Travis 2000, pp. 221–227). The minimal interpretation of what Travis holds would be that the meaning of a word is just a set of constraints on what that word can be used for.² The role of meaning, he claims, is to set some constraints on what can be said: a word cannot simply be used to refer to anything (Travis 2000). That is, there is variability in word use, but there is also stability, a normative stability.

So, we may think that word types have meaning in a very thin sense: all the utterances of a word have something in common, namely, all the utterances of a word set the same general constraints on what can be said when using that word. If we move from words to sentences, we may assume that non-ambiguous sentence types also have meaning, according to Travis. This further assumption only requires that compositional rules are stable, something which Travis, on the face of it, does not call into question.

Now, if this interpretation of what Travis has in mind is right (I do not claim that it is, only that it might be), his claim that meaning does not determine truth-conditions may stem from his endorsement of the following two theses:

- (i) Any two non-ambiguous, indexically non-problematic, utterances, *u* and *u'*, of the same linguistic type, can have different truth-conditions on different occasions;

² I find it difficult to go further than this minimal interpretation of Travis. That meanings somehow embody constraints on the correctness of word use is something he clearly holds: e.g. “[s]uch constraints are meaning contributions to the collaborative enterprise of making the content of our statements what it is” (Travis 2000, p. 222). However, he apparently wants to explain where these constraints come from as well: “Meaning’s contribution to the cooperative enterprise of fixing content is to make the correctness of one’s words a matter of how the things one speaks of fit into that system [of categories]—a matter of the way the system one thus invokes, so invoked, would categorize” (ibid., p. 224). I do not dare to suggest an interpretation of what Travis says here and in other similar passages. However, the purpose of this paper is not exegetical. What I’m interested in is not so much Travis’s account—he would surely reject the representationalist approach I am taking—as his examples.

- (ii) Any two non-ambiguous, indexically non-problematic, utterances of the same linguistic type have the same meaning.

The examples seem to show that the first thesis is true. It is perhaps more controversial to ascribe the second thesis to Travis. Yet, as I have just argued, if meanings are indeed stable features of words, and compositional rules are stable as well, (ii) has to be true. It is important, however, not to place too much content into (ii). The thesis should be interpreted as the claim that any two non-ambiguous utterances of the same type set the same constraints on the conditions that can make them true. The assumption that both (i) and (ii) are true leads us to conclude that meaning does not determine truth-conditions. That is, the constraints brought about by the meaning of a sentential utterance are not sufficient to determine its truth-conditions.

In the next section, I briefly present the main reactions to this argument developed in recent years by other authors, before I develop and defend a new approach.

3 Minimalism, contextualism, indexicalism and relativism

As it happens, many philosophers of language seem to find the argument unobjectionable, even though there is radical divergence as to what it actually shows. Minimalists such as Borg (2004) or Cappelen and Lepore (2005) clearly reject the argument, trying to block it by denying one of its premises. For, according to minimalists, (i) is false. If we have two non-ambiguous utterances which contain no overt indexicals, or contain indexicals whose values do not change with the proposed shift of context, then those two utterances must have the same truth-conditions. In their terms, they must express the same minimal proposition (which either is just the set of truth-conditions or determines those truth-conditions). According to minimalists, the reason why we may think otherwise is that we mistake “knowledge of truth-conditions” for “knowledge of when those truth-conditions are satisfied” (see Borg 2004). Knowledge of truth-conditions is a purely semantic knowledge and can be captured in a simple T-schema. In the example of green leaves, such knowledge tells us that ‘the leaves are green’ is true if and only if the contextually salient leaves are green: that is all. The truth-conditions Travis talks about are not truth-conditions, but rather, they look like verification conditions.

Indexicalists are onboard with the minimalists, as they would also reject premise (i). In contrast to minimalists, indexicalists can concede that the truth-conditions of the different utterances of ‘the leaves are green’ are those that the contextualist prefers. What they want to defend is that all variations have to be traced back either to some indexical expression (Rothschild and Segal 2009) or to some covert indexical attached as a hidden variable to some expression (Szabó 2001; Stanley 2002). Thus, if two utterances of the same linguistic type differ in their truth-conditions, it must be due to the presence of some indexical constituent or other: it cannot be that indexically non-problematic expressions differ in their truth-conditions.

As I mention above, many other philosophers of language take the argument to be good. Only perhaps Recanati (2004) could be taken as objecting to premise (ii). Recanati distinguishes two layers of meaning: standing meaning and modulated meaning. Standing meaning is not sensitive to contextual shifts and does not determine truth-conditions, whereas modulated meaning does determine truth-conditions. Modulated meaning is the meaning interpreters obtain after taking contextual (pragmatic) information into account. It seems, moreover, that according to Recanati's view, the meaning of an utterance *is* its modulated meaning. Thus, (ii) would be false on this account. However, if by meaning we understand *semantics*, Recanati would not be objecting to (ii): he agrees that the semantics of an utterance does not fix its modulated meaning and, a fortiori, that it does not fix its truth-conditions.

Relevance theorists share this overall view. They also distinguish between encoded meaning and explicatures: the explicature has determinate truth-conditions, but it incorporates pragmatic knowledge. The level of encoded meaning, on the other hand, is purely semantic, but it falls short of determining truth-conditions. There are two strands in Relevance theory, though. According to the first (Sperber and Wilson 1986/1995), (categorematic) words encode concepts. According to the second (Carston 2008), word meaning is not conceptual; rather, the semantic information attached to a word is thin, and its role is simply to put some constraints on what concepts the word may express, or to indicate a part of conceptual space (see Carston 2002). The semantic level, following this construal, is very close to the syntactic level.

Now, even though both strands in Relevance theory want encoded meaning to be non-propositional, the "conceptualist" strand may have more problems accounting for Travis cases in their preferred way. There seems to be no difference at the level of encoded conceptual meaning between two different (indexically fixed) utterances of 'the leaves are green'. So, the conceptual meaning of the expression type 'the leaves are green' does not change from context to context. However, if its conceptual meaning does not change, then it seems to follow that its propositional meaning will not change either (a proposition being, according to this view, an adequately structured combination of concepts). Thus, if we assume that propositions fix truth-conditions, it means that two different utterances of 'the leaves are green' have to have the same truth-conditions, contrary to premise (i) of the argument. In short, it seems that if Relevance theorists want to endorse Travis's argument (instead of sharing the minimalist take on it), they had better opt for Carston's thin construal of semantics.

The last group of authors that would not object to the argument is the relativist group. There are profound disagreements between those authors and radical contextualists such as Travis, to be sure. However, the disagreements do not touch the truth of the premises or the correctness of the argument. Relativists would have no quarrel with the premises at all; and they would agree that meaning does not determine truth-conditions. Their disagreement with contextualists lies in the lessons the latter purport to extract from the argument. As Predelli (2005) puts it, what they deny is that there is something exciting about all this, something that should make us rethink the semantics enterprise. For, according to the relativist,

Travis cases can be tackled perfectly well with the apparatus of formal semantics. Travis cases show that intensions alone do not determine truth-conditions. However, intensions plus points of application, or circumstances of evaluation, do determine truth-conditions. Formal semantics, thus, runs no risk from acceptance of the argument.

In what follows, I develop a different approach to Travis cases, which resorts to a “thick” conceptual semantics for lexical items. In the last section, I will discuss the virtues of this approach vis-a-vis the responses to Travis cases that I have just reviewed. To advance: the thick semantics approach is, overall, clearer and more conservative than its rivals. It is clearer than some of its rivals in that it does not invoke anything other than conceptual meanings for constituents and semantic rules that have such meanings as their inputs. In this, it contrasts with the “thin” semantics approach adopted by Carston (2008) and suggested by Travis himself, where meanings are not concepts, but some other kind of representations which encode constraints on conceptual expressions. Thin semanticists have yet to explain what these representations are and how they work. The approach is also clearer than indexicalism, which accounts for Travis cases in an unnecessarily complicated manner (and is most probably unsuccessful: see Clapp forthcoming). Furthermore, as I will suggest, the indexicalist would have to postulate the presence of indexicals in all sorts of lexical items (i.e., not only predicates, but also verbs) to cover all the data.

The approach I advocate is more conservative than other approaches in that it takes its *explananda* to be “intuitive” truth-conditions, and not some arbitrary and/or flat truth-conditions, the so-called ‘minimal’ truth-conditions. It is also more conservative than the thin semantics approach in that it assumes that utterance meaning is truth-conditional. Finally: thick semantics forms part of an explanatory programme, conceptual semantics, which, among other things, has provided successful accounts of phenomena as diverse as argument structure and (some cases of) polysemy. Minimalists, indexicalists and thin semanticists have yet to provide explanations for those phenomena. (It seems that minimalism is especially ill suited to that task.)

4 Denying (ii)

The view of meaning that I want to argue for provides another way of rejecting the argument above; by blocking it at premise (ii). Thus, I propose that we work under the assumption that (i) is true; that is, that neither minimalists are right in claiming that contextualists misfire in their attribution of truth-conditions in Travis cases, nor is it plausible to postulate the existence of as many indexical constituents as the indexicalist programme would require.³ In short, I take it for granted that two indexically non-problematic, non-ambiguous utterances of the same linguistic type

³ For an effective critique of the indexicalist programme as applied to colour predicates, see Clapp (forthcoming). Minimalism, in contrast, besides refusing to take *intuitive* truth-conditions as its *explananda*, has poor explanations for many other linguistic phenomena which will be mentioned here, such as polysemy, argument structure, impossible words, or coercion.

can have different truth-conditions. My claim, however, is that this does not imply that meaning does not determine truth-conditions, nor that the meaning of a sentential utterance is not fully propositional. In order to argue for this claim, I suggest that, instead of embracing a thin semantics, as Carston (2008) proposes, we should go for quite a thick semantics, in the spirit of what has been labelled as “conceptual semantics” (Jackendoff 2002).

Before going into the details of the proposal, there is a precision to be made: I have been speaking freely about non-ambiguous utterances. Ambiguity, however, is not a totally regimented concept. In the present context, it is usually taken to cover both scope, and, in general, syntactic, ambiguities and homonymy. Thus, an utterance can be ambiguous either because it is syntactically ambiguous or because it contains a homonymous term. In what follows, I only use the term ‘ambiguity’ to refer to syntactic ambiguity and homonymy. In accordance with this understanding of ‘ambiguity’, my proposal is rightly seen as a denial of premise (ii) of the argument. Yet, on a broader understanding of the term, it is probably best seen as a rejection of (i): just as indexicalists are able to reject (i) by postulating the existence of many more indexicals than the classic ones, it is possible to reject the premise by postulating the existence of more sources of ambiguity, broadly understood, than the classic ones.⁴

Another caveat: premise (ii) states that any two indexically non-problematic, non-ambiguous utterances of the same linguistic type must have the same meaning. One could think that this is obviously false. Take a simple sentence such as: ‘the cat is on the mat’. I can use this sentence to mean that a given animal is on the mat, but also to mean that a certain toy is on the mat. Referential expressions can be used to refer either to objects (or sets of objects) or to toys which represent those objects or sets of objects. That is, all referential expressions are polysemous, so there is no reason to think that two utterances of the same sentence type must have the same meaning.⁵ Alternatively, think about metaphors as sources of polysemy: as is well known, many metaphors become conventional and some even become dead (the metaphorical element of the meaning is not even retrievable). So, there is a point when a metaphor is properly treated simply as another meaning of a certain expression (which has some kind of resemblance to the historically former meaning of that expression). That is, the retrieval of the metaphoric content is no longer a pragmatic issue, but a matter of polysemy resolution. This shows again that two utterances of the same expression can have different meanings (in this case, one “literal”, the other metaphorical).

⁴ In fact, if we take ‘ambiguity’ in this broad, unregimented, sense, it seems that it is tautological that the variability in the truth-conditions of indexically non-problematic linguistic types is due to ambiguity.

⁵ In fact, it seems that referential expressions can be used to refer either to objects or to just any iconic representation of those objects. Culicover and Jackendoff (2004) tell us that there is a “statue rule”, according to which any NP can refer either to some determinate x or to a physical/iconic representation of x . Their talk suggests that this “statue rule” is a conventional linguistic rule. Cognitive linguists would oppose this way of thinking about what goes on in such cases: they take it that we think of representations in terms of represented things (Panther and Thornburg, 2003). Be that as it may, this kind of (cognitive or linguistic) rule is a permanent source of polysemy.

However, this obviously misses the point. Polysemy is a deep problem for a *literalist*, that is, someone who holds that there is a one-to-one relationship between words and concepts and that sentences have one, and only one, propositional content. Thus, polysemy is a problem for those who place much more into claim (ii) than is actually asserted. For instance, it is a problem for minimalists (witness Fodor's quarrel with polysemy: Fodor 1998; Fodor and Lepore 2002); but it is not a problem for Travis or, in general, for any of the authors who hold that linguistic meaning is thin. A polysemic expression can share the same meaning across its different uses if one thinks about meaning in terms of constraints. In fact, what distinguishes polysemy from homonymy is that the different uses of a polysemic expression have something in common: one of the constituents of polysemy is that there is some stability across uses. A thin semanticist would identify the stable feature of a polysemic expression with its meaning.

In any case, polysemy could eventually give us reasons to suspect that claim (ii) is not true, but polysemy is not what seems to be at issue in Travis cases. The variation in truth-conditions detected by Travis is not due to ambiguity, but seems not to be due to polysemy either: *prima facie*, 'leaves', 'are' and 'green' do not behave as polysemic expressions across the contexts we are invited to consider. So, just what is going on here, then?

Here is a story a conceptual semanticist can tell. First of all, let us accept that lexical entries are conceptually rich, as conceptual semanticists tend to think. For instance, Pustejovsky (1995) has famously argued that lexical entries are highly complex representations which include information about things such as the argument and event structure (in verbs) and the *qualia* structure—a structure which specifies knowledge about the constitution, the specific properties, the causal origin and the purpose (if there is one) of the *designatum*. Thus, a lexical entry for a noun, according to his view, gives us as information about the origins of its referent, about its purpose, about its constitution and about its distinguishing properties, i.e. those properties that help to single it out within a superordinate category. For instance, the meaning of the noun 'novel' tells us that novels are constituted by a narrative, that they are books, that they are for reading and that they are written by agents. All this information is meaning-constitutive, according to Pustejovsky.

Pustejovsky's approach has famously been put to service to explain polysemy (see below); which many authors see as a failed attempt (see, e.g., Carston 2010). It is not my purpose here to defend the attempt against its critics, but only to borrow the very general idea that word meanings may be richer than previously thought. In particular, that they can be so rich that they include knowledge that many would see as world (and not semantic) knowledge (more on this below).

So, I propose to borrow a good part of the overall idea and claim that, typically, a lexical entry for, e.g. a noun, consists of a complex conceptual representation that collects rich information about whatever entity the noun is intended to designate, such conceptual representation being its meaning. For instance, if the noun is intended to designate a set of objects such as pencils, its lexical entry will include information such as: that they are used for writing or drawing, that they typically produce grey writing, that they have a certain stereotypical shape, that they are composed of two parts, that they can be sharpened, etc. Also, the lexical entry may

include the information that pencils can have an apparent and a hidden, original, surface (as, in general, all physical objects can have), so that there may be a difference between appearance and reality even in surface properties. This last piece of information can be directly stored in the lexical entry 'pencil', or it can be stored only indirectly. For instance, it can be that the lexical entry 'pencil' just tells us that pencils are physical objects, and it is the general concept PHYSICAL OBJECT that informs us that physical objects can present to us surface properties (such as colours) which are not their original properties.⁶

With this in place, we are in a position to claim that sentences such as 'the leaves are green' are indeterminate. All that we know upon hearing an utterance of it, is that the property greenness applies to the leaves. Yet, this is indeterminate. For greenness can apply to leaves in at least two ways: given that leaves are physical objects and that physical objects can have apparent and original, hidden, surfaces, greenness can apply to the apparent surface of the leaves or it can apply to their original surface. Putting it in other, more direct words: the lexical entry 'leaf' tells us: "I'm a physical object, so I can have an apparent and a hidden original surface. Which of these two surfaces is supposed to be green? That is a question that one has to answer in order to determine the meaning of the utterance (that is, without such an answer, the meaning—not only the truth-conditions—is indeterminate)."⁷

The case of the blue ink is not exactly the same. Here what we have is that an utterance of 'that ink is blue' can have two different truth-conditions: the set of conditions where the ink writes blue and the set of conditions where it looks blue. That is, the variability is not due to the ink's having original and apparent surfaces (though one can think of a further case where, e.g., 'that's my blue ink' is said truthfully of ink which is now black but was originally blue). However, the case is similarly explained. 'That ink is blue', per se, gives us very scarce information: again, indexicals aside, it only tells us that the concept BLUE applies to the ink. However, when we try to apply BLUE to INK, we find we have a complex concept (INK) that tells us, basically, that ink can be blue in different ways: is the predicate supposed to apply to the external aspect of the ink, or is it supposed to apply to what the ink does when it is put to its standard use, i.e. writing?

The explanation here follows in part what Pustejovsky (1995) has to say about polysemy. Let me briefly present two examples. First: arguably, 'book' does not mean the same in 'I'm writing a book' as in 'I'm reading a book' or 'the book is on the shelf'. According to Pustejovsky, 'book' has different meanings in these three

⁶ This piece of knowledge is an expression of the more general knowledge that the original properties of objects may be altered.

⁷ Here is another "Travis case" one might think about: imagine that Pia has not painted some maple leaves, but some autumn leaves fallen from the tree long ago. Then we get the photographer first and the botanist afterwards. Perhaps, with this variation, it is tempting to say that "the leaves are green" may stand either for the proposition that the leaves appear to be green or for the proposition that the leaves are fresh, or alive. If that were the case, instead of the explanation related to apparent and hidden surfaces, the explanation could be: the lexical entry for 'leaf' tells us that leaves are physical objects, but also that they are living things which have a determinate colour while alive—typically, they are green. The property of being green thus can apply either to the leaves considered as living things or to the leaves considered as physical objects. If we want to know the meaning of an utterance of "the leaves are green", we have to know in which way GREEN is supposed to apply to LEAVES.

sentences because the part of the *qualia* structure that gets selected is different in each case. In the first case ('I'm writing a book'), the book is presented as content which is nowhere but in the mind of the author—at most. In the second case ('I'm reading a book'), the book is presented also as content, but as content that has been put into written words. In the third case ('the book is on the shelf') the book is presented as a physical object. The idea that the book is presented in one way or another is to be understood as the claim that what comes to the fore is one aspect or another of the information associated with 'book', information which constitutes its meaning. That is, the claim is compatible with the idea that the rest of the information is also activated, though in a weaker way.

Second example: take the polysemy of 'window'. It seems that different aspects of windows are highlighted in each of the following sentences: (1) 'he crawled through the window' (opening), (2) 'the window is broken' (glass), (3) 'the window is rotten' (frame), and (4) 'my favourite part of the bedroom is the window' (whole). Each sentence makes us think about windows from a particular perspective, to the extent that one can say that different notions are expressed by the different occurrences of 'window', thus yielding different truth-conditions in each case. What can be plausibly claimed in this case is that the word-type 'window' stands for a certain complex which tells us what windows are, and that different parts of that complex come to the fore on different occasions. Again, it may be that the rest of the information is also activated; but a particular aspect is more active than the others.

Now, the explanation of Travis cases I am pointing to is that rules such as "BLUE modifies INK", leave indeterminate which part of the conceptual structure associated with 'ink' should come to the fore. That is, it does not specify how the ink is to be thought of, how it is presented (whether primarily as a liquid or as something that writes). In general, I want to claim, the information we begin with when extracting sentential meanings is typically unspecific. To illustrate this point, let me come back to Pustejovsky's ideas. First let us once again consider the case of 'book'. Pustejovsky (1995) controversially claims that 'I am starting a book' has the default meaning that I am starting *to read* a book. His idea is that, as 'start' requires an event as a complement, when interpreters come across a noun phrase, they look at its *telic quale* and construct an event. As the *telic quale* of 'book' specifies that books are for reading, the resulting event is "reading a book". Now, as many critics have pointed out, things are much less systematic than Pustejovsky supposes. 'I am starting a book' can perfectly well mean that I am starting *to write* a book, that I am starting *to print* a book, etc. That is, we have to construct an event, but there seems to be no systematic way of doing so. If we accept that what goes on in these cases is, as Pustejovsky claims, that NPs are type-coerced into eventive structures, what we have is that the coercion rule we follow must be highly unspecific. At least, it does not specify which part of the conceptual structure of BOOK has to be activated.

Something along the same lines can be said about the 'window' case. It is true that the most habitual reading of 'the window is broken' is that the window pane is broken, but of course 'the window is broken' can also mean that the window has come off its hinges. The concept BREAK does not necessarily activate WINDOW PANE. So, once again it seems wiser to think that the composition rule is less specific than Pustejovsky proposes.

Going back to Travis cases: as claimed above, what we have in the example of ‘the ink is blue’ is that the composition rule only tells us that BLUE modifies INK. The context has to fix which aspect of INK is modified (its appearance as a liquid or the writing it produces), and so it has to fix how BLUE is going to modify INK. In the case of ‘the leaves are green’, the rule tells us again that GREEN applies to LEAVES. The context has to determine how this rule is to be specified, as GREEN can apply to the apparent surface of the leaves or to their original surface. The rule is unspecific because our concept LEAF includes the information that leaves are physical objects whose original surface may have been altered. In the same way as we cannot know which part of the window is broken when we hear ‘the window is broken’, we do not know which surface of the leaves is said to be green when we hear ‘the leaves are green’. We may usually suppose that the speaker is talking about the original colour of the leaves, just as we may suppose that the speaker is talking about the window pane when telling us that the window is broken, but such suppositions are just that: suppositions.

Now, one could think that all that this shows is that *semantic* construction rules are often indeterminate, and that this is as far as semantics go: it just states general and probably unspecific construction rules that serve as constraints on truth-conditions. It is even possible—though hardly plausible, I think—to read Travis as indicating such an idea in the following passage: “[t]o say ‘It’s blue’ of some ink, speaking English, is to describe the ink as blue. Meaning decides that much and, *by itself*, no more” (Travis 2000, p. 221). In short, it can be claimed that meaning does not determine truth-conditions because composition rules are unspecific. Two utterances of the same sentence type do have the same meaning (their constituents have the same meaning and the unspecific rules are the same); yet, they may have different truth-conditions. Now, it just is not plausible to ascribe this position to Travis because it is not plausible to ascribe it to any thin semanticist—and Travis certainly seems to be one. For, in order to hold that rules are indeterminate, you need thick word meanings: the indeterminacy of the rules stems from the fact that a property concept can apply to an object concept in different ways due to the object concept’s having a certain structure.⁸

However, the idea that semantic composition rules might be unspecific seems implausible on its own. There are two levels of composition rules in the picture suggested here: first, we have the, e.g., unspecific rule “BLUE applies to, or modifies, INK”. Then, we have the specific rule: say, “BLUE applies to the appearance of the ink, considered as a liquid”. This second rule emerges when knowledge of the first rule meets knowledge of the meaning of ‘ink’. Contextual factors help select it as the rule that should be used in understanding a particular utterance of ‘the ink is blue’. So, there are pragmatic factors involved in retrieving the meaning of the utterance, but what the pragmatic system is doing is not “intruding” on the semantic system or adding something that was not there: it is just helping to select one composition rule out of a set of possible rules which are entirely determined by semantic facts.

⁸ It is possible to both be a thin semanticist and claim that construction rules operate on thick concepts, once these are “reached” or “constructed” in accordance with the constraints of word meanings; but this seems to multiply things unnecessarily.

It is natural to think about the first kind of rule as only compositional templates, and to think about the second kind as proper semantic composition rules. The main reason is that the first kind of rule does not produce anything that we would call meaningful. Semantic rules are in the business of providing the meaning of wholes, when they have as their input the meaning of constituent parts: the first kind of rule does not give us the meaning of the whole, and it can even be argued that such rules do not really take the meaning of the constituents as their input—they are not sensitive to what the constituents really mean. Another reason to think that the more specific rules do deserve to be called semantic, regardless of what status we assign to the unspecific rules, is that compositional rules of the second kind are, as suggested, entirely determined by semantic facts, namely, by the semantic/conceptual information included in lexical entries. What is not determined is which information we should apply in a particular case.

In summary, what we have in Travis cases is that it is undetermined which one of the several possible composition rules we have to use in order to capture the meaning of the utterance. That is, composition rules are not indeterminate; what is indeterminate is which construction rule we have to use. If we do not have the relevant information, we will not understand the meaning of the utterance. As it happens, 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.

To put it another way: the meaning of the constituent pieces determines the possible meanings that a certain complex may have. Thus, if we have the construction 'new book', its possible meanings are, say, NEW TOME and NEW TEXT (see Cruse 2004). Pragmatic factors select one of these possible meanings, thereby yielding the meaning of the whole. In the case of sentences, what we have is that the meaning of the constituents plus the syntactic information generate a number of possible meanings for the whole, that is, a number of possible truth-conditions. It is up to pragmatics to select one of these possible meanings, i.e., to pick out some specific (but pre-existent) truth-conditions.

As I say, it is lexical entries which generate the indeterminacy of composition rules. As there is quite a lot of information contained in a lexical entry, a rule that works only with syntactic information cannot determine, say, how a property concept applies to an object concept. It may be surprising to see that, in particular, it is lexical entries *for nouns* which generate the indeterminacy of composition rules, at least in the cases considered here. That means that the source of variation in Travis cases is not the predicate, but the noun. That is, the value of 'green' remains constant across contexts, while what 'leaves' contributes to the truth-conditional content of the utterance varies from context to context (i.e., a different part of the structured concept can become differentially activated in different contexts).

On second thoughts, however, this should not be that surprising. A colour concept is typically flat, whereas kind concepts are informatively rich. As Carey (2009) puts it, kind concepts are "inductively deep". If, as is customary, concepts are thought of as bodies of information stored in long-term memory that are used for categorization and induction, it is predictable that a kind concept will contain much more information than a colour concept: kind concepts are projectible, they pick out interesting categories in the world and, above all, they are the basis for our inductive

inferences. The information that they contain must be rich: otherwise, they will not be suitable to form the basis of our inductive inferences. That said, it should not be odd to claim that the information contained in a kind concept can be differentially activated; that is, depending on the context, some parts of the information stored in the concept become more activated than others.

5 Objections and discussion

There seems to be an obvious problem with this account, a problem inherited from the conceptual semantics programme which inspires it. It is the problem of conflating world knowledge and semantic knowledge. Before exploring that problem, however, let me speak in favour of the account proposed here.

The first virtue of the account is its comparative clarity. The account is simple: lexical meanings are conceptual and lexical concepts have a rich structure.⁹ As the semantic information in lexical entries is so rich (or at least, as rich it is), syntax alone cannot determine which rule we should use in generating the meaning of whole utterances: syntax underdetermines semantic rules. This last point, moreover, is the only point where the account departs from orthodoxy. The rest of its profile is conservative: word meanings are concepts, utterance meanings are truth-conditional, the truth-conditions of an utterance are the *intuitive* truth-conditions (i.e. the data for a theory of meaning are the speaker's intuitions), etc. Other approaches are more elusive (and less conservative to boot). For instance, it seems sensible to categorize Travis as a "thin semanticist". Now, there is something elusive in thin semantics. We do not get a clear sense of what meanings are if they are not concepts: we are told that the meaning of a word embodies some constraints on the range of things (concepts) the word can be used to express (Travis 2000), that it contains instructions for building concepts (Pietroski 2010), or that it points to a region of conceptual space (Carston 2002). These ways of speaking do not really explain what meanings are and they leave it unclear how meanings relate to concepts and sentence meaning to truth-conditional contents. In fact, postulating a layer of meanings, or semantic representations, which provide access to conceptual structures, seems to be an unnecessary complication.

In contrast, the approach I present here is more informative than, for instance, the indexicalist approach advocated by Rothschild and Segal (2009). Rothschild and Segal hold that predicates are indexicals whose content is fixed according to unspecified and variable contextual factors (that is, they are not fixed by a specific set of contextual parameters).¹⁰ This does not explain what constrains the correct

⁹ One could think of this structure as a "cluster" structure, instead of a definitional structure (Vicente 2010).

¹⁰ On the other hand, Rothschild and Segal's approach is notoriously complicated; and it is difficult to believe that colour predicates are indexical, anyway (see Clapp forthcoming). They may behave as indexicals in certain circumstances, such as when they behave as covert comparatives ('blue' meaning BLUER THAN X), but they are not behaving as indexicals in these typical Travis cases. Finally, the indexicalist approach would have to claim that all sorts of predicates besides colour terms (see, e.g., 'new' in 'new book') and even verbs ('break' in 'break a window') are indexicals.

use of a certain predicate. In the present account, it is possible to explain and predict the correct uses of a predicate. Roughly: a predicate is correctly used because (will be correctly used if) the property concept which is expressed by the predicate applies to one of the (adequate) features of the object concept of which it is predicated.

Beside from the relative merits of this approach vis-a-vis its actual or possible rivals, I think that its main strength is that it is part of a programme that seems to be highly fruitful. The idea that meanings are conceptual and that lexical concepts have structure has been put to use in the explanation of phenomena as diverse as argument structure (Levin and Rappaport Hovav 2005; Pinker 2007), impossible words (Johnson 2004), polysemy (Jackendoff 1992; Pustejovsky 1995) and coercion (Jackendoff 2002). For instance, it seems that no other theory (of meaning or otherwise) is as predictive and explanatory as conceptual semantics when it comes to explaining alternations (see Levin and Rappaport Hovav 2005; Pinker 2007).¹¹ Again, it seems that no theory other than conceptual semantics can explain why 'blik', a verb that enters into constructions such as 'the desk blikked Mary', meaning that Mary broke the desk, is an impossible word in English (see Johnson 2004). The answer the conceptual semanticist provides is that the event structure of MARY BROKE THE DESK is MARY [CAUSED [THE DESK BECOME <BROKEN>]]. If that is the case, then 'Mary', and not 'the desk', has to occupy the subject position in the sentence, for it denotes the unembedded argument, which in turn is the causal origin of the causal event, while 'the desk' denotes the patient.

Finally, whereas there are various accounts of polysemy (an arduous matter if ever there were), conceptual semantics seems to have at least some simple explanations for some simple cases. Thus, Jackendoff (1992) famously proposed that the stability, on the one hand, and the variability, on the other, in the uses of 'keep' can be accounted for by postulating that 'keep' is related to a structure like [CAUSE [ENDURANCE [STATE OF X]]]. What is common to all the uses of 'keep' is that a state is caused to endure; what varies is the kind of state that endures (it can be a state of location, a psychological state, etc.).

Now, the idea that meanings are conceptual, that lexical concepts have structure, and that this structure is *informationally rich* has also been put to use in the explanation of some phenomena. As mentioned above, Pustejovsky (1995) applied it to the study of (some cases of) polysemy. Another phenomenon that seems to be accountable for in terms of this general idea is coercion (Jackendoff 2002). What 'John slept for three hours' means differs from what 'John jumped for three hours' means, apart from in other respects, in that 'John slept for three hours' tells us about a *continuous* event, while 'John jumped for three hours' speaks about a *repeated* event. That is, 'John jumped for three hours' means that John jumped repeatedly for three

¹¹ Thus, Pinker (2007) has this to say about why 'fill' does not enter into the locative alternation while 'load' does: 'fill' "is all about the state of the container... It is apathetic about how the container became full. By specifying the change of state of a container, it is compatible with a construction that is about a state-change... Because it says nothing about a cause or a manner of motion of the contents, it isn't compatible with a construction that is all about motion". On the other hand, 'load' "simultaneously specifies something about how a content has been moved *and* how the container has changed" (p. 50; his italics).

hours. Now, where does this information come from? A plausible answer is that it is generated by the knowledge that we have about what jumping amounts to. The only way to make this knowledge compatible with the information that John was jumping for 3 h is to assume that what he did for 3 h was make a long series of jumps.

It seems that such an account of this and similar examples resorts not only to structured concepts, but to something that could be thought of as being world knowledge (as opposed to semantic knowledge). It seems that knowing that jumps take a short time is not part of the meaning of 'jump', but something that we know when we learn about jumps. Whereas the explanations of impossible words, argument structure, etc. do not call into question the intuitive division between world knowledge and semantic knowledge,¹² Jackendoff's explanation of these cases of coercion make the division problematic, as it is claimed that some aspects of what seems to be world knowledge affect meaning facts.

Thus, what I am proposing here as an explanation of Travis cases belongs to a fruitful research programme which can explain a variety of phenomena that, e.g., minimalist and thin semantics have yet to face. However, it belongs to that part of the general programme that crosses a line, the line which separates the semantic from the encyclopaedic, as Relevance theorists would put it. This is *prima facie* troublesome, but on the other hand it is also good news: for the present proposal is clearly not the first to resort to so-called world knowledge to account for meaning facts.¹³

Now, is the division between semantic and world knowledge neat enough? Is it sustainable? Am I really blurring it? The answer to these three questions depends on the way we tell the story. I propose to tell it thus: when doing semantics (or when talking about meaning) we are doing cognitive semantics, that is, we are trying to uncover some facts about how meaning is processed. Concepts and propositions that are easily and directly retrieved should count as meanings of, respectively, words and sentential utterances. There is, *ceteris paribus*, a strong correlation between conventionality and the speed of processing: the more conventional a use of a word is, the faster we will process it. So, what counts as semantic knowledge can be put either in terms of cognitive access or in terms of conventionality. Now, both cognitive access and conventionality come in degrees: a piece of information can be more or less accessible and a practice can be more or less conventional. This means that it is going to be difficult to draw the distinction between semantic knowledge and world knowledge: if the elements that we have to use to build the distinction are gradable, the distinction cannot track a "joint of nature".

There are clear cases of knowledge that should count as semantic, namely, those cases where a use of a word remits to a highly conventional pairing of word to conceptual structure which is extremely easily accessed. There are also clear cases

¹² Though a full explanation of alternations might have to resort to what we intuitively regard as world knowledge. For instance, as Falkum (2011) points out, 'cut' enters into the conative alternation in 'John cut the rope'/'John cut at the rope', but it does not admit such alternation in 'the bank cut its interest rates', as 'the bank cut at its interest rates' is clearly odd. This seems to have to do with what cutting interest rates amounts to.

¹³ The most radical proposal is probably Langacker's (see, e.g., 2008). In general, cognitive linguists—probably the most common defenders of conceptual semantics—are suspicious of the division between semantic and world knowledge (see Cruse 2004).

of knowledge that should not count as semantic. Idiosyncratic pairings or information which is accessed only after retrieving a (more) conventional pairing should not count as semantic. But there are surely cases that lie in between these two extremes that will be debatable. My proposal is to move up the line a bit in those cases with respect to what seems to be the orthodoxy. Thus, I think that a piece of knowledge such as that objects can have two surfaces should count as semantic knowledge; and applying surface property predicates, such as colours, to original surfaces is conventional enough (as of course is applying them to apparent surfaces). It is true that, for instance, some children will not understand an utterance of 'the leaves are green' as meaning that their original surface is green.¹⁴ They will have to learn that things can be painted so that they may have two layers of surface. But that does not mean that the knowledge has to be forever non-semantic. The concept of, say, physical object, grows and absorbs what was previously world knowledge.

6 Conclusion

Travis' examples have been seen as a problem for the development of a truth-conditional semantics. I have, however, denied that Travis cases show that meaning does not determine truth-conditions. I have argued that the variability in the truth-conditions that two utterances of the same type can express may be due to semantic facts alone. To this end, I have proposed that we think about word meanings (in particular, noun meanings) as rich conceptual structures; so rich that it may not be determined in which way a property concept applies to an object (or class) concept. The proposal entails that there is a role for pragmatics in interpreting utterances. However, pragmatic processing is triggered by semantic facts, much in the way that indexical saturation involves pragmatic processing triggered by semantic facts. In the present account, what the pragmatic system has to do is select a composition rule.

The main problem of this approach is that it apparently blurs the distinction between semantic knowledge and world knowledge. However, I have defended that, if the criteria for establishing the distinction are features such as conventionality and accessibility, then: (a) the distinction cannot be sharp and (b) it is possible to maintain that word meanings are as rich as is required by the proposal sketched here.

Acknowledgments My gratitude to Begoña Vicente and to Fernando Martínez Manrique. Research leading to this work has been supported by the Spanish Government, Research Project FFI2011-30074-C02-02.

References

- Borg, E. (2004). *Minimal semantics*. Oxford: Oxford University Press.
 Cappelen, H., & Lepore, E. (2005). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*. Oxford: Blackwell.
 Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.
 Carston, R. (2002). *Thoughts and utterances*. London: Blackwell.

¹⁴ They show some preference for this understanding, however, as soon as they become "essentialists".

- Carston, R. (2008). Linguistic communication and the semantics/pragmatics distinction. *Synthese*, 165, 321–345.
- Carston, R. (2010). Explicit communication and “free” pragmatic enrichment. In B. Soria & E. Romero (Eds.), *Explicit communication: Robyn Carston's pragmatics*. London: Palgrave.
- Clapp, L. (forthcoming). Indexical color-predicates: Truth-conditional semantics vs. truth-conditional pragmatics. *Canadian Journal of Philosophy*.
- Cruse, A. (2004). *Meaning in language: An introduction to semantics and pragmatics*. Oxford: Oxford University Press.
- Culicover, P., & Jackendoff, R. (2004). *Simpler syntax*. Oxford: Oxford University Press.
- Falkum, I. L. (2011). *The semantics and pragmatics of polysemy: A relevance-theoretic account*. Ph.D. Thesis, University College London.
- Fodor, J. (1998). *Concepts*. New York: Oxford University Press.
- Fodor, J., & Lepore, E. (2002). The emptiness of the lexicon: Reflections on Pustejovsky. In J. Fodor & E. Lepore (Eds.), *The compositionality papers* (pp. 89–119). Oxford: Oxford University Press. (Revised from the version published in *Linguistic Inquiry*, 29, 269–288.)
- Jackendoff, R. (1992). *Languages of the mind*. Cambridge, MA: MIT Press.
- Jackendoff, R. (2002). *Foundations of language*. Oxford: Oxford University Press.
- Johnson, K. (2004). From impossible words to conceptual structure: The role of structure and processes in the lexicon. *Mind and Language*, 19, 334–358.
- Langacker, R. W. (2008). *Cognitive grammar: A basic introduction*. Oxford: Oxford University Press.
- Levin, B., & Rappaport Hovav, M. (2005). *Argument realization*. New York: Cambridge University Press.
- Panther, K.-U., & Thornburg, L. L. (2003). Introduction: On the nature of conceptual metonymy. In K.-U. Panther & L. L. Thornburg (Eds.), *Metonymy and pragmatic inferencing*. Amsterdam: John Benjamins.
- Pietroski, P. (2010). Concepts, meaning and truth: First nature, second nature and hard work. *Mind and Language*, 25, 247–278.
- Pinker, S. (2007). *The stuff of thought*. London: Penguin Books.
- Predelli, S. (2005). Painted leaves, context, and semantic analysis. *Linguistics and Philosophy*, 28, 351–374.
- Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: MIT Press.
- Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
- Rothschild, D., & Segal, G. (2009). Indexical predicates. *Mind and Language*, 24, 467–493.
- Sperber, D., & Wilson, D. (1995). *Relevance: Communication and cognition* (2nd ed.). Oxford: Blackwell.
- Stanley, J. (2002). Making it articulated. *Mind and Language*, 17, 149–169.
- Szabó, Z. G. (2001). Adjectives in context. In I. Kenesei & R. M. Harnish (Eds.), *Perspectives on semantics, pragmatics and discourse: A festschrift for Ferenc Kiefer*. Amsterdam: John Benjamins.
- Travis, C. (1985). On what is strictly speaking true. *Canadian Journal of Philosophy*, 15, 187–229.
- Travis, C. (1996). Meaning's role in truth. *Mind*, 105, 451–466.
- Travis, C. (2000). *Unshadowed thought*. Cambridge, MA: Harvard University Press.
- Vicente, A. (2010). Clusters: On the structure of lexical concepts. *Dialectica*, 64, 79–106.