

# Are Utterance Truth-Conditions Systematically Determined?

Claudia Picazo (UGR)

(Forthcoming in *Inquiry*, S.I. on Semantic Variability)

Abstract: Truth-conditions are systematically determined when they are the output of an algorithmic procedure that takes as input a set of semantic and (optionally) contextual features. Truth-conditional sceptics have cast doubts on the thesis that truth-conditions are systematic in this sense. Against this form of scepticism, Schoubye and Stokke (2016) and Dobler (2019) have provided systematic analyses of utterance truth-conditions. My aim is to argue that these theories are not immune to the kind of objections raised by truth-conditional sceptics. In particular, I argue that the use of Questions Under Discussion (Schoubye and Stokke) and ways of being (Dobler) is problematic.

Keywords: underdeterminacy, truth-conditional scepticism, utterance truth-conditions, QUD, ways of being, semantics, modulation.

## 1. Systematic determination

Are the truth-conditions of our utterances<sup>1</sup> systematically determined? In order to address this question, a couple of clarifications are needed. First, the truth-conditions of an utterance *u* of sentence *s* can be different to (although constrained by) the linguistic meaning of *s*, even if *s* is non-indexical. It has been repeatedly noted that utterance truth-conditions are often the result of contextual modulation, a process whereby the meaning of *s* is adjusted to the context of utterance (Travis, 1997, Carston, 2002, Recanati, 2004). Here are a couple of well-known examples<sup>2</sup>:

(1) The leaves are green.

Travis (1997) imagines two occasions of use for this sentence. In the first, Pia decides to paint the leaves on her tree green because she does not like their original colour (brown). After finishing, she utters (1). In the second, a botanist calls, asking for green leaves for a

---

<sup>1</sup> By ‘utterance’ I mean the content uttered, what is said by a use of a sentence.

<sup>2</sup> The second example is usually considered an instance of incompleteness, not modulation. Incomplete sentences express incomplete propositions that need to be contextually completed in order to determine a truth-evaluable content. Completion is a mandatory process. By contrast, modulation is optional. As some authors argue that the contextual supplementation in (2) is not mandatory and equivalent to other cases of modulation (see Cappelen and Lepore 2005, 59-68), I treat it as modulation.

scientific study. Pia utters (1) as an answer. The intuition elicited by the example is that (1) is now false<sup>3</sup>. Philosophers have used this kind of example to argue that meaning can be modulated (Recanati, 2004), i.e., adjusted so as to fit the conversational context, and that therefore the truth-conditions of sentences as (1) can shift across contexts (keeping the referent of ‘the leaves’ fixed). In the first occasion ‘green’ is used to describe the observable colour of the leaves, whereas in the second it applies to their original colour. Another example:

(2) Tipper is ready.

As Bach (1994) notes, (2) does not make explicit what Tipper is ready for. What is said with the sentence can vary across contexts: that Tipper is ready for dinner, that Tipper is ready for an exam, etc.

The question about the determination of utterance truth-conditions concerns these context-relative truth-conditions. I will not address here the question whether sentences express truth-conditional content or whether, by contrast, utterances are the proper bearers of truth-conditions. What is important here is that, in these and other similar examples, the content that is intuitively communicated goes beyond the linguistic meaning of the sentence uttered.

Second, what does ‘systematically determined’ mean here? By this expression I mean that truth-conditions are a function of a set of features. In particular, the truth-conditions of an utterance *u* are systematically determined when there is an algorithmic procedure that takes as input a set of semantic and (optionally) contextual features of *u* and delivers its truth-conditions as output. The same output must always correspond to the same input. If, keeping the set of features fixed, there is variation in truth-conditions, then that shows that truth-conditions are not systematically determined by that set after all. An example of a non-systematic procedure in this sense would be that of an aesthetic all-things-considered judgement. In this kind of judgement, various things can be weighed, with no precise principles determining a verdict. For example, when judging a painting, we can consider the composition, colours, content, etc. People with different values, or even the same person on different occasions, can rank these things differently. So they don’t determine an aesthetic verdict. Although it is not standard, all-things-considered judgements have been used in the Philosophy of Language. Gauker (2008) argues that the referent of a demonstrative is the

---

<sup>3</sup> See Hansen and Chemla (2013), and Grindrod, Andow and Hansen (2019) for empirical evidence.

object of an all-things-considered judgement in which the things to be considered are salience, relevance, prior reference, charity, and pointing in location in a series. If he is right, then the referent of a demonstrative is not systematically determined. Something similar can perhaps be said about the determination of what Tipper is ready for. A second example of non-systematic determination is sketched in Travis (2000). According to Travis, our judgements concerning utterance truth-conditions depend on what we find reasonable in particular occasions, and, he adds, ‘so far as we know, there is no algorithm for reasonableness’ (Travis 2000, 209). That there is no algorithm for reasonableness means that there is no systematic procedure responsible for the attribution of truth-conditions to utterances.

Natural language semantics has traditionally assumed that sentential truth-conditions are systematic in the previous sense. Semantic theories that include a compositionality principle as the following are a good example:

*Compositionality:* The truth-conditions of a sentence *s* are determined by the semantic value of the simple expressions in *s* and their mode of composition.

According to this principle, truth-conditions are determined by a set of syntactic and semantic features. The semantic theory provides the rules to calculate, so to speak, the truth-conditions of the sentences within its scope on the basis of their linguistic meanings and modes of composition.

The phenomenon of modulation and Travis-style examples introduce some complexities. Besides semantic and syntactic features, utterance truth-conditions (as opposed to sentential truth-conditions) seem to be contingent upon contextual features, such as the topic of conversation. For these cases, semanticists have put forward formal theories that can be seen as procedures that take as input syntactic, semantic and contextual features and deliver utterance truth-conditions as output. For instance, it has been argued that the logical form of ‘green’ contains a context-sensitive variable for the part of the object that is supposed to exhibit the colour (Szabó, 2001). If this theory works, the truth-conditions of utterances of (1) can be said to be systematically determined (ignoring the complexities introduced by the expression ‘the leaves’).

There are, however, reasons to be sceptical about the project. That truth-conditions are systematically determined is something that has long been called into question. Several philosophers and linguists have endorsed a view that I will call ‘truth-conditional scepticism’. As an example, authors as Travis and Pietroski suggest that judgements about what we find

reasonable in particular occasions play a central role in determining the truth-conditions of an utterance, which casts doubts on truth-conditional semantics:

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, 68)

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, 218) (Emphasis added.)

Here I will be concerned with one particular argument that can be found in Searle and Travis<sup>4</sup>. I will call it the ‘Searle-Travis regress’. The argument goes as follows. We start with a representation of an utterance’s content as a set of semantic and (optionally) contextual features<sup>5</sup> f. For example, we start with a representation of (1) as a set of semantic features. Second, it is shown that utterances that share all the features in f can nonetheless have

---

<sup>4</sup> See Searle (1978, 1980 and 1992, 182) and Travis (1997, 2000, 35-36, and 2008, 6).

<sup>5</sup> By ‘contextual feature’ I mean the value of a parameter.

different truth-conditions. Travis does this by describing two occasions in which utterances of (1) have different truth-conditions. Third, one can add more features to  $f$  (presumably, contextual features). For example, a contextual feature corresponding to the part of the object that exhibits the colour. But again, another example is put forward showing that the new set is also insufficient. For example, one in which the illumination conditions are relevant to the application of the colour term. And so on.

There is one reason to think that the situation will always repeat itself. According to philosophers as Travis, the new features that we add to  $f$  can themselves be responsible for further variation in truth-conditions. Suppose that we add a parameter ‘part’ to the set of semantic features of (1). Then the question will arise as to what the admissible values of this parameter are. In particular, the problem is that plausible values for ‘part’ are compatible with variation in truth-conditions. For example, we can represent the utterance of (1) in the first occasion described by Travis as *set1*:

(set1) Semantic features *Sem(1)*

Contextual feature *Part: surface*.

The problem with *set1* is that two utterances can share all these features, yet be true under different conditions. For example, one can be true if and only if the green paint covers at least 90% of the leaf’s surface whereas the other can be true if and only if the paint covers some relevant spot. If this line of reasoning is correct, then there are serious doubts that utterance truth-conditions are systematically determined by a set of semantic and contextual features.

In order to better appreciate the sceptic’s argument, it can be useful to distinguish semantics and metasemantics. Semantic theories are theories about the composition of so-called semantic values. Semantic systematicity, in this sense, amounts to compositionality. Metasemantics is about how the values that go into composition are fixed. In the case of context-sensitive expressions, metasemantics is about how the value of the expression is fixed in context. A metasemantic theory is systematic iff there is a systematic procedure determining the value of the expressions it is about. The notion of systematicity at stake here can be said to combine semantics and metasemantics. It is semantic in the following sense: it is about a set of features (values) getting composed and yielding truth-conditions. But it is also metasemantic in the sense that it takes as input contextual features and therefore requires the identification of a set of parameters responsible for the variation in truth-conditions. Travis’ target is not the

compositionality of semantic values, but rather the more ambitious project of identifying a set of parameters such that the value of these parameters together with the context-invariant properties of the sentence uttered (semantic and syntactic properties) unambiguously determines truth-conditions.

This notion of systematicity is different from Recanati's (2010, 9-10). Recanati claims that modulation is compatible with systematic semantics. Although he acknowledges that modulation itself might be unsystematic, Recanati points out that we can define a *mod* function that takes as input an expression  $e$  and a context of use  $c$  and yields as output the modulation function  $f$  that is relevant for the interpretation of  $e$  in  $c$ . The output of the modulation function is the value of the expression in context. This value, together with the semantic values of the other expressions, goes into composition thus yielding the truth-conditions of the utterance<sup>6</sup>. Here, systematic semantics takes as input whatever is the result of modulating 'green', or, to use Travis' terminology, whatever counts as 'green' in a context of use. By contrast, the notion of systematicity that I am using here, and that I take at least some of Travis' arguments to target, is one that points towards the relation between some features of the context of use and the truth-conditions of the utterance. Thus it goes beyond compositionality in that it imposes a metasemantic requirement. Part of what Travis seems to attack is the idea that there is something that determines what counts as 'green' in a particular occasion other than our judgements about what is reasonable to take as 'green'—other than our going case by case deciding what is the most reasonable choice given the circumstances.

This more demanding notion fits well with a traditional approach to truth-conditional semantics exemplified in Szabó's theory (mentioned above). On this kind of approach, often called indexicalism, meaning (semantics) is supposed to determine truth-conditions. In the case of context-sensitive expressions, meaning cannot do it by itself, but it is conceived as a rule that picks a referent in context. Indexicals as 'I' are the paradigmatic example. The meaning of 'I' seems to identify a parameter (speaker) and can be conceived as a character that determines a referent in context. So meaning can still be said to determine truth-conditions. This machinery sets the basis for a nice explanation of our ability to interpret speech: we are able to interpret an utterance of 'I'm tired' because we know the meaning of the expressions in it, and meaning tells us what to look for in context—namely, a speaker. Travis' arguments, including saliently the regress, call into question the possibility to extend this approach to other words (as colour predicates). Thus he casts doubts on (at least some

---

<sup>6</sup> See Rothschild and Segal (2009) for a similar theory.

versions of) truth-conditional semantics. This can (and has) be seen as a problem, as, quoting Rothschild and Segal, ‘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). Here is where Recanati’s notion of systematicity enters the picture. Taking modulation seriously, one of the aims of Recanati is to show that this phenomenon does not lead to a complete rejection of systematic semantics. In particular, Recanati shows that modulation is compatible with what I have called ‘semantic systematicity’, i.e., compositionality. The modulated meaning of an expression in context is not determined by linguistic meaning (as opposed to what happens with indexicals), and perhaps it is neither determined by a set of parameters as the ones envisioned by Szabó, but this does not prevent all forms of systematic semantics.

Now, for the purposes of this paper I will use the more demanding notion of systematicity. The reason is twofold. First, I take this notion to be the target of sceptics as Travis. Second, new theories that identify a contextual parameter responsible for the shift have been put forward, and it is interesting to assess how these new theories fare with respect to sceptical arguments. These new theories are not instances of indexicalism. They do not complexify the meaning of the expressions used in Travis-style examples by adding new, context-sensitive variables corresponding to contextual parameters. However, as indexicalism, they identify a contextual parameter that explains the shift.

Here is the plan. In the next section I present and assess the Searle-Travis regress in more detail, distinguishing three versions of it. I then examine two recent proposals that deal with the examples above, as well as other cases of modulation or Travis-style examples: Schoubye and Stokke (2016) (section 3) and Dobler (2019) (section 4). Both theories follow Roberts (2012) and use partition semantics (Groenendijk and Stokhof, 1984). Although the regress argument does not apply as such to these theories, I argue that similar lines of reasoning cast doubts on these approaches.

## 2. Three versions of the regress argument

The aim of this section is to present and assess three versions of the Searle-Travis regress, focusing especially in Travis’ use of the argument. These versions are not clearly distinguished on Travis’ work. However, addressing them separately provides a better grasp of the sceptic’s position.

### 2.1. Number of parameters

The first version of the argument concerns the number of contextual features that affect truth-conditions. This version calls into question the assumption that there is a limited number of parameters on which truth-conditions depend. The idea here is that there is no end to the number of parameters that one would need to include in her theory—we start with a part parameter, then we add another parameter for illumination conditions, a new example is imagined showing that a third parameter is needed, and so on. Although it might work against some specific proposals (as Szabó's), this line of reasoning is not very promising. Some theories, saliently the ones discussed in the next sections, posit generic parameters, as goals, that can be used to deal with many examples. So it seems that, once we have the right kind of parameter, this version of the regress argument no longer applies.

### 2.2. Sentences

The second version targets sentences. Searle (1992) and Travis (1997) develop their argument by showing that more explicit sentences than (1) can have different truth-conditions on different occasions. Travis (1997) briefly considers the possibility of substituting (1) with (3):

(3) The leaves are painted green.

As with (1), Travis takes it that we can imagine two occasions of use, one in which (3) is true and another in which it is false. Imagine that the leaves are painted in pointillist style and only look green at a certain distance. Although Travis himself does not develop the example, we can think of two photographers looking for green leaves for a picture: a landscape and a close-up of a leaf. Whereas in a conversation about landscapes an utterance of (3) would be appropriate, it seems it would be false in a conversation about a close-up. Something similar can be said about (4):

(4) The leaves are green on the surface.

Imagine now a dead tree whose leaves are brown but covered with green dust. Someone interested in decorating the garden could utter (4) and thereby say something true. By contrast, in a different conversation, for example if someone, after being told that the tree



is dead, utters (4) surprised that the leaves have not turned brown, it could be objected that they are not green on the surface, but covered with green dust.

This second version of the argument is problematic for some semantic theories. In particular, it is problematic for theories that are supposed to state truth-conditions, as do disquotational clauses in Davidsonian semantics. The problem is that the sentence that is used in the metalanguage to state the truth-conditions of the target sentence can itself be interpreted in multiple ways, and so is itself in need of interpretation. Consider, for the sake of the argument, a theory that delivers the following instance of the T-schema:

(1\*) ‘The leaves are green’, as used to describe the superficial colour of the leaves, is true if and only if the leaves are superficially green.

As the sentence ‘The leaves are superficially green’ can be interpreted in multiple ways, Travis (2006) holds that we have failed to state the truth-conditions of the target utterance. Now, what is the moral of this argument? Although it creates problems for a certain form of semantic theorizing, I think that it fails to establish the metaphysical conclusion that utterance truth-conditions are not systematically determined. I will illustrate why with a modification of Perry’s Z-land example (Perry, 1986).

Z-landers are a group of people who have no representation for Z-land, the place where they live. Neither in language nor in thought. They have no word or concept for it, and they do not travel. Except for this, they are just like us. In particular, Z-lander semanticists work on theories of truth-conditions and state the truth-conditions of their favourite sentences. One sentence they theorise about is (3):

(5) It’s raining.

Given that she lacks the concept of Z-land, the Z-lander semanticist states the truth-conditions of (5) with 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 include a variable for a location. The Z-lander semanticist is simply blind to the fact that the truth-value of (5) is relative to a location. However, when a Z-lander says ‘It’s raining’ the content of her utterance concerns Z-land. It is true if and only if it’s raining in

Z-land, regardless of whether or not it rains at other places. We could specify the truth-conditional content of the 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.

Having a concept and a word for Z-land, we are able to do what the Z-lander semanticist cannot—we can identify and specify the (let us suppose) complete set of contextual features that determines the truth-conditions of Z-landers’ utterances of ‘It’s raining’. As a conclusion, the fact that the Z-lander semanticist is unable to identify the set of contextual features relevant for the analysis of (5) does not mean that there is no such a set. An English semanticist can very well add a location parameter to her semantic theory for the language spoken in Z-land, thus making explicit the relevant contextual feature.

As applied to sentences, the Searle-Travis regress suggests that we are unable to encode in a sentence the truth-conditions of our utterance. The right side of the biconditional in (1\*) can be interpreted in multiple ways, for Travis-style examples can be imagined for ‘The leaves are superficially green’. With respect to our own languages, we are on a par with the Z-lander semanticist in the sense that the instrument we use in disquotational clauses (natural language) is not completely well suited to the purpose. In the case of the Z-lander, it is because the semanticist is blind to the existence of a place-parameter. In ours, it is because the words we use to state truth-conditions are open to a variety of interpretations<sup>7</sup>.

Two remarks about the scope of the argument are in order. First, as I said above, the analogy with the Z-lander semanticist suggests that the sceptic’s point should be seen as concerning the limitations of our theorizing on truth-conditions rather than as a metaphysical claim about truth-conditions themselves. Second, the argument does not apply to all semantic theories alike. A careful distinction between linguistic entities and contents might avoid the challenge, as I show in section 3. Here is where the third version of the argument becomes relevant.

### 2.3. Content

---

<sup>7</sup> As Gross (2005) notes, semanticists tend to endorse a thesis of Meta-Insensitivity according to which the meta-language used in a truth-theory must be context-insensitive. A way to escape the regress argument would be to put forward theories that do without the Meta-Insensitivity thesis. However, Gross discusses some problems for this move.

Together with sets of parameters and sentences, Travis (1989, 23-24) applies the argument to contents and, in particular, to properties<sup>8</sup>. As applied to (1), the argument would go as follows. In order to explain the variation in truth-conditions, we can postulate that the predicate expresses different properties in the two scenarios described, P and Q. However, Travis argues, P and Q are not immune to the kind of argument used for sentences. We can see this if we ask questions about whether other objects have these properties. As before, we can ask whether a leaf painted in pointillist style would make Pia's utterance in the first scenario true—whether the leaf is P. For Travis, we might be unable to answer these questions, and this shows that it is has not been determined<sup>9</sup>. Even though what Pia says is more specific than what (1) encodes, her communicative intentions can fail to determine a total function from possible worlds to truth-values. For many objects, it can be indeterminate whether they count as green or not. In this sense, what she says can still be made more precise, in the same sense that (3) is more precise than (1). The point of the argument is that properties are analogous to words. Our interpretation of (1) depends on the contextual information we have. The more information we have, the more specific our interpretation will be (that the leaves are green in some sense or other, that they are painted green, that they are painted green with such-and-such technique...). According to Travis, the same happens with properties. We can posit that, in a given occasion of use, 'green' expresses the property P. But, Travis' objection goes, the property can be made more specific.

How challenging is this argument for theories of truth-conditions? Travis seems to point towards some kind of instability in truth-conditions, so to speak: he seems to suggest that content can always be refined by adding more contextual information. However, we need not go that far. As I mentioned above, I think that what this version of the argument calls into question is the viability of conceiving content as a total function from possible worlds to truth-values. Rather, for many possible worlds it can be indeterminate whether that possible world makes the utterance true<sup>10</sup>. For instance, in the first scenario described for (1) it is plausibly indeterminate whether a leaf painted in pointillist style would make the utterance true—the truth-value at those worlds is indeterminate. Thus utterance content is best seen as a partial function. When further contextual information is added, i.e., when the sentence is interpreted in a richer context, as one in which the possibility of painting leafs in pointillist style has been

---

<sup>8</sup> To my knowledge, he does not apply it to a sentence together with the value of a parameter, as e.g., the contextual value *surface*. However, we could apply the same line of reasoning as for properties. See also Borg (2012, 36).

<sup>9</sup> One could reply that our being unable to answer certain questions does not show that there are no answers. I discuss this counterargument and potential objections in Picazo (2020).

<sup>10</sup> See Picazo (2019).

raised, we get a new (probably also partial) function—one that does have a truth-value at worlds where the relevant leaf is painted in pointillist style. Since what we have here are two different inputs and two different outputs, it is compatible with systematicity as defined in section 1.

As a conclusion, each version of the argument targets a different aspect of systematic approaches to truth-conditions. The first version casts doubts on approaches that try to identify all dimensions of incompleteness, as Szabó (2001) calls them, i.e., approaches that try to list all the contextual dimensions that shift truth-conditions, as part or illumination conditions for colour terms. The second version casts doubts on disquotational clauses and the possibility to paraphrase or encode content in a context-insensitive sentence. Finally, the third version introduces some indeterminacy in truth-conditions in the sense of calling for partial functions. In the next sections I examine two recent theories that deal with examples as (1) and (2). These two theories posit generic parameters, thus escaping the first version of the argument. However, I will argue that either other versions of the regress argument or related lines of reasoning challenge some assumptions of these theories.

### 3. Questions Under Discussion (QUD)

In this and the next section I examine Schoubye and Stokke's (2016) and Dobler's (2019) proposals respectively. The aim of these authors is to put forward a systematic theory of utterance truth-conditions. In order to achieve this aim, they identify a contextual feature (Question Under Discussion in the case of Schoubye and Stokke and goals in the case of Dobler) that plays an indispensable role in the determination procedure. Both proposals follow Roberts (2012) and Groenendijk and Stokhof's partition semantics (1984).

Schoubye and Stokke's (2016) explicit goal is to provide a systematic account of the intuitive truth-conditions of utterances ('what is said', in their terminology). Contexts, according to these authors, contain Questions Under Discussion (QUD), a technical counterpart of the informal notion of topic of conversation. The goal of the conversation is to answer the relevant QUD (Roberts, 2012). Utterance truth-conditions are determined by the literal truth-conditions of the sentence uttered (called 'minimal content') and the QUD of the context. Minimal content is, as usual, represented as a set of worlds, and the set of answers to a question is represented as a partition on a set of possible worlds. Here is the determination procedure, in their terminology: what is said by an utterance *u* is the weakest proposition entailed or entailing the minimal content of *u* that is an answer to the QUD of the context. One of the examples they discuss is (2). Imagine the following conversation:

Julie: Is Tipper ready for the interview?

Rebecca: She's ready.

In this example, Julie explicitly introduces a QUD. The set of answers to this QUD is a partition including two sets: the set of worlds where Tipper is ready for the interview and the set of worlds where Tipper is not ready for the interview. The minimal content of Rebecca's answer is in turn represented as the set of worlds where Tipper is ready (for something). The weakest proposition that entails the minimal content of 'She's ready' and is an answer to the QUD introduced by Julia is *that Tipper is ready for the interview*, i.e., the set of worlds where Tipper is ready for the interview. This fits most people's intuitions about the content of Rebecca's utterance. We can represent the set of features used in this proposal as follows:

(set2) Semantic feature *Minimal content of 'She's ready' (set of worlds)*

Contextual feature *QUD: Is Tipper ready for the interview?*

Schoubye and Stokke put forward a systematic procedure determining utterance truth-conditions. Their formal apparatus takes semantic features (minimal content) and contextual features (QUDs) as input and delivers utterance truth-conditions as output. Importantly, it makes the right predictions. The choice of QUDs as the contextual feature that explains the shift in truth-conditions allows them to account for different examples of modulation. However, the determination of the relevant QUD brings in important problems. In their presentation, it seems as if QUD can be explicitly or implicitly introduced<sup>11</sup>. Let's consider first explicitly stated QUDs.

In the example, the QUD is introduced via an explicit question. Now, it is important to note that questions, qua sentences, are as open as declarative sentences to modulation. Just as the content of Rebecca's answer is more specific than the minimal content of the sentence uttered, the same can happen with Julie's question. A speaker can mean different things with 'ready for the interview': that Tipper is already dressed and ready to drive to the location where the interview takes place, that Tipper has prepared her resumé, that Tipper has prepared a presentation of her skills and done some research about the company, etc. Explicit questions are subject to a version of the Searle-Travis regress: two utterances of (2) can answer

---

<sup>11</sup> See also Roberts (2012).

the same explicit question, yet have different truth-conditions because the question is interpreted in different ways. Explicitly stated questions as the one in the example cannot be the whole story, as, in order to use them to calculate the content of the target utterance, we need first to interpret them.

It can be objected that QUDs, as opposed to questions, are contents, not linguistic entities, and that, as such, they are not open to different interpretations, so to speak<sup>12</sup>. Linguistic entities are the things that we use in order to convey content. Contents are what linguistic entities express. Linguistic entities can be ambiguous (think of ‘Cora went to the bank’), but contents cannot—contents disambiguate linguistic entities. This objection undermines the Searle-Travis regress as it has been applied here, as it explicitly concerns the specification of the QUD (thus a linguistic entity). Now, the problem with this reply is that it leaves us with no explanation as to how to determine the relevant QUD—whether it is the content of an explicit, articulated question or an implicit QUD. In what follows I argue that, because of that, there are reasonable doubts that QUDs can do the work they are supposed to do. First, as Grindrod and Borg (2019) argue, theories that appeal to QUDs in order to deal with modulation and related phenomena face a circularity threat. I will point out that the threat is especially relevant for Schoubye and Stokke’s theory. Second, the theory assumes that contexts contain a single, well-defined QUD that enters into the determination of the utterance’s truth-conditions. However, it should be doubted that ordinary conversations contain a unique QUD.

Grindrod and Borg (2019) argue that Schoubye and Stokke’s proposal, as well as other proposals that use QUD in order to deal with context-sensitivity broadly understood, face a circularity threat: either there is an independent way of determining the QUD of the context (something not offered in Schoubye and Stokke’s paper), or it seems as if identifying it implicitly depends on what we take the utterance to say. An example might help clarify this point. Imagine that Julie, Rebecca and Tipper are going to a party. Tipper, who is supposed to drive them to the party, is late. While waiting for her, Julie and Rebecca chat about an interview Tipper will have in a few days and Julie asks ‘Is Tipper ready for the interview?’. Suddenly, Tipper appears, car keys in hand. Rebecca says, while pointing at Tipper: ‘Look. She’s ready.’ The intuition here is that Rebecca’s utterance says that Tipper is ready to go to the party, not for the interview. But how is the relevant QUD determined? Unless a plausible story is told about the determination of implicit QUDs, the order of explanation here seems to

---

<sup>12</sup> See e.g. Fodor (2003, 106-108) and Recanati (2007, 12-13) for this distinction.

be the opposite than the one postulated in Schoubye and Stokke's account. We seem to identify the QUD on the basis of a prior grasp of the (modulated) content of Rebecca's utterance, not the other way around. In what follows I argue that Schoubye and Stokke's discussion of implicit QUD actually makes things worse.

According to Schoubye and Stokke, the QUD at stake can be inferred from contextually available cues. They follow Roberts (2012) in this respect and use prosody as an example. When a sentence is uttered, prosodic focus provides crucial information about the QUD that the utterance addresses. Consider the next sentence:

(6) Fritz will go to Potsdam tomorrow.

(6) can be used to answer several questions, including: 'Where will Fritz go tomorrow?', 'When will Fritz go to Potsdam?', and 'Who will go to Potsdam?'. An assertion of (6) can answer any of these questions. Now, the intonational pattern will vary depending on which question is being addressed. Simplifying, we can distinguish utterances of (6) where the focus is on 'Potsdam' from utterances where it is in 'Tomorrow' or 'Fritz'. Because of that, prosody can play two interrelated roles: it provides hearers a clue to infer which QUD the speaker is addressing, and it allows the speaker to introduce an implicit QUD via accommodation. The QUD is presupposed by the prosodic focus; if the utterance is accepted by the interlocutors, it is accommodated and becomes the active QUD at that point of the conversation. This provides an example of how speakers can introduce implicit QUD via assertion. But can the QUD required to modulate (2) be introduced in an analogous way? It seems not. By definition, modulated content goes beyond what is encoded in the sentence. In Schoubye and Stokke's theory we get the extra bit (for the interview) from the QUD. This makes the analogy with prosody problematic: the only way in which an assertion of (2) could introduce the QUD *Is Tipper ready for the interview?* would be by having a modulated content from which the hearer could infer the extra bit. The reason is clear. In order for the hearer to be able to infer the QUD *Is Tipper ready for the interview?* from an utterance of (2), all the ingredients must already be there. This is what happens with utterances of (6), where the sentence plus the prosodic focus are sufficient to obtain the relevant QUD. By contrast, we cannot infer the relevant QUD if all one has in the minimal content of (2)—a plurality of QUDs are compatible with (2), some of which are not the relevant QUD (that Tipper is ready to have a shower, that Tipper is ready to have dinner, etc.). Thus, if the QUD was to be introduced via an utterance of (2), then the content of (2) would need to be already

modulated, and the explanation would be circular. As a conclusion, the possibility to introduce implicit QUD via prosodic focus cannot be used to motivate the possibility to introduce the kind of QUD required for the determination of modulated content.

The second worry concerns an assumption of singularity. Schoubye and Stokke assume that utterances address a single QUD. When an explicit question is raised it might make sense to posit a single QUD<sup>13</sup>. However, when the QUD has to be inferred from the contextually available cues it seems more plausible to admit a plurality of QUDs. In ordinary conversations, the contextual information constrains the questions that it is reasonable to take as QUD, but very often it fails to single one out. Imagine again the two friends waiting for Tipper. When Tipper appears and one of them utters (2) it can be perfectly clear for both of them that ‘ready’ is to be interpreted with respect to the party they are going to. However, previous discourse and other contextual information might fail to single out a unique QUD. Several candidates seem to be admissible: ‘Is Tipper ready to go to the party?’, ‘Is Tipper ready to drive to the party?’, ‘Is Tipper ready to leave?’. These QUD, together with the minimal content of (2), determine different utterance contents (what is said, in Schoubye and Stokke’s terminology). However, the interlocutors can be indifferent towards which of them actually is *the* QUD of the context<sup>14</sup>.

#### 4. Extralinguistic goals and ways of being

In this section I examine Dobler’s semantics (Dobler, 2019), another account that follows Roberts (2012) approach to discourse structure and applies partition semantics. Instead of QUD, i.e. conversational goals, Dobler makes use of domain goals, i.e. the non-linguistic purposes of a conversation, and provides an account in which utterance truth-conditions are determined by these goals together with the minimal content of the sentence uttered. Minimal content is again represented as the set of worlds that make the sentence true. The difference is that this set is partitioned according to alternative ways in which the sentence can be true (alternative truth-makers)<sup>15</sup>. Dobler uses (1) as an example. The minimal content of (1) is the set of worlds in which the leaves are green in some way or other. Now, there are different ways of being green. Simplifying, Dobler considers that this set can be partitioned into three

---

<sup>13</sup> For the sake of the argument, I ignore here the objection raised above about the possibility of having modulated questions.

<sup>14</sup> This argument is similar to an argument used in Buchanan (2010), among others. According to Buchanan, sentences containing certain context-sensitive expressions, as quantifiers and comparative adjectives, fail to determine a unique proposition in context. As he puts it, the communicative intentions of the speaker are indifferent towards a variety of candidates. My point here is that the same happens with implicit QUD.

<sup>15</sup> See Sainsbury (2001).



subsets that make use of two ways of being ‘green’: those in which the leaves are naturally green, those in which the leaves are painted green, and those in which the leaves are both naturally and painted green. In a conversation, some of these partitions, but usually not all, will be conducive to the domain goal. This is represented in the formal apparatus as a valuation map that assigns each partition value 1 (goal-conducive) or 0 (not goal-conducive). The truth-conditions of the utterance are precisely the set of partitions that are goal-conducive, i.e. a set of worlds in which the sentence is true and useful to the goal of the conversation. For example, for the conversation with the botanist only the sets in which the leaves are naturally green or both naturally and painted green are goal-conducive. Thus, these sets represent the content of utterances of (1) in that conversation. We can represent the set of features that are used to determine utterance truth-conditions as follows:

(set3) Semantic feature *Minimal content of ‘The leaves are green’ (set of worlds) and ways of being green (partition).*

Contextual features *Goal: Botanist’s goal, Valuation map for botanist’s goal*

Two notions are key for the proposal: domain goals and ways of being. Below I pose an objection for ways of being. Before that, some brief remarks about how domain goals fare with respect to the Travis-Searle regress are in order. As I have already noted, generic purposes as goals avoid the first version of the argument. The second version would apply to sentences making explicit goals. One natural option for the partitions on (1)’s minimal content that Dobler gives as an example is to go for a goal that can be encoded as ‘scientific research’. Now, utterances sharing the minimal content of (1) (used to talk about the same leaves) and the encoded goal ‘scientific research’ can have different truth-conditions. Imagine a leaf that is green by the effect of some mould. In some scientific researches this leaf will count as ‘green’, for example in a research about different kinds of mould. However, there will be other cases in which the leaf will not count as ‘green’, for example in studies about chlorophyll. Against this line of reasoning, in Dobler’s proposal goals are not linguistic entities, and so the proposal escapes the second version of the argument as well. As to the third version, we should take home the point that communicative intentions can be vague and do without total valuation functions. When it comes to domain goals, it can be indeterminate whether some ways of being are goal-conducive. For example, it could be unclear to the botanist’s colleagues whether leaves that are both painted and naturally green are goal-conducive. Thus the valuation map should be conceived as a partial function.

I move now to ways of being. A first worry concerns the granularity of the partition. Since domain goals can have different granularities, the structure of the partition needs to be more complex than in the simplified example discussed by Dobler. Going back to the previous example, a conversation can have a coarse-grained goal, as doing scientific research. Other conversations can have finer-grained versions of this goal, as doing a study on chlorophyll. Analogously, ways of being green can be conceived as more or less fine-grained. Being naturally green is something that can be split into being green by the effect of chlorophyll, being green by the effect of some mould, etc. In order for the proposal to work, the granularity of the ways of being green (thus, of the partition) should correspond to the granularity of the domain goals. An example might help here. Imagine a group of botanists working on a study on chlorophyll. As part of their work, they discuss where to get the leaves they need for the study and say things like ‘Pia’s leaves are green’. The domain goal in the conversation is to get leaves that they can use in their study, that is, leaves that are green because of the presence of chlorophyll. Therefore, only possible worlds where Pia’s leaves are green because of the presence of chlorophyll make the utterance true. Possible worlds where Pia’s leaves are green because of some mould do not make the utterance true. In order to make the right prediction about truth-conditions, the theory must discriminate possible worlds where the relevant leaves have chlorophyll from those where they have mould. In order to account for different examples we need ways of being with different granularities. Which level of granularity is used in the interpretation of a particular utterance will depend on the domain goal that the utterance concerns. This can be seen as an unpalatable consequence, as the selection of a semantic property (the level of granularity of the partition) depends on a contextual feature (domain goals). However, the partition itself does not depend on the goal.

Nonetheless, a sceptic can object to the notion of ways of being on different grounds—namely, that Travis’ examples challenge the semantic structure of the proposal. I will put the objection in terms of competence. Although not presented as such, Dobler’s proposal distinguishes semantic and pragmatic properties. Meaning and its partition (ways of being) belong to semantics, whereas domain goals obviously fall on the pragmatic side. Thus being semantically competent involves two things—knowledge of meaning, understood as a set of possible worlds, and knowledge of ways of being, i.e., a partition on the set of worlds determined by meaning. In this framework, semantically competent speakers should know, in virtue of their semantic competence, how to partition the meaning of words as ‘green’ into alternatives. This means knowing that being painted green, naturally green, and so on, are ways of being green. Now, Travis-style examples challenge precisely this. Part of Travis’ point

is that competent speakers do not know how to classify things independently of an occasion of use (or a domain goal). We know that, for certain purposes, painted green leaves count as ‘green’, that is, that for certain purposes we must include being painted green as a way of being green. We also know that, for other purposes, painted green leaves do not count as ‘green’, that is, that for those purposes we must not include being painted green as a way of being green. But competent speakers can fail to know or even disagree about whether, out of context, independently of any purpose, being painted green is a way of being green (rather than, strictly speaking, a way of not being green)<sup>16</sup>. Travis’ examples cast doubts on the ability of competent speakers to occasion-independently see which ways of being are, e.g., ways of being green. This being so, and in absence of reasons to take competent speakers to be blind to semantic properties, it cannot be assumed that words as ‘green’ have the semantics that Dobler posits.

To reinforce this point, let me note that we can find in the literature more challenging examples. Clapp (2012) imagines a group of scientists studying colour blindness. Each week they show various swatches to a different subject and adapt their use of ‘green’ to the subject’s perceptual judgements. In this context, a use of ‘This is green’ means roughly the same as ‘This looks green to this week’s subject’. In order to apply Dobler’s explanation to this case, we would have to say that looking green to a colour-blind subject is a way of being green. Now, although the swatches that look green to the subject might count as ‘green’ for the scientists’ purposes, it would be odd to conclude from this that it is part of the semantics of ‘green’ that being perceived as green by a colour-blind subject is a way of being green.

As a conclusion, the assumption that the meaning of ‘green’ is a set of worlds that is partitioned into alternative ways of being green turns out to be problematic from. First, the two uses of (1) described in section 1 raise doubts that competent speakers would count being painted green as a way of being green independently of the occasion, as Dobler assumes. Second, other examples analogous to the one imagined by Travis have been put forward, and it should be doubted that in all of them it makes sense to speak of ways of being green and, therefore, that Dobler semantics can be applied to them. In particular, it would not be very natural to assume without further justification that being perceived as green by a colour-blind subject is a way of being green.

---

<sup>16</sup> As an anecdote, when I present the green leaves example to non-philosophers or philosophers not working on language an answer I frequently get is that the leaves are not *really* green, because they are painted. This presupposes that being painted green is not a way of being green.

## 5. Conclusions

Truth-conditional scepticism targets the thesis that utterance truth-conditions are systematically determined by a set of semantic and contextual features. In section 2 I have presented three versions of an argument that sceptics have used against truth-conditional semantics—the Searle-Travis regress. Although it's not a knockdown argument, the regress poses a challenge to some semantic theories. However, two recent proposals that follow Roberts (2012) and Groenendijk and Stokhof (1984) seem to avoid the challenge: Schoubye and Stokke's account of modulation using QUD and Dobler's explanation of Travis' examples using domain goals. I have argued that these proposals are not without problems. QUDs and domain goals have the advantage of affording a general account of modulation and Travis-style example. However, there are doubts that the relevant QUD is determined independently of (and prior to) the interpretation of the target utterance. Moreover, ordinary conversations might contain not one but a plurality of QUDs, similar but not identical, that would give rise to different truth-conditions. Domain goals might fare better, as they can avoid the circularity threat. Nonetheless, in Dobler's proposal, their role is to select among alternative ways of being, and this notion is challenged by Travis-style examples. As a result, truth-conditional scepticism has not been ruled out by these theories. When carefully examined, both theories reveal themselves to be subject to objections similar to those raised by proponents of truth-conditional scepticism<sup>17</sup>.

### References:

- Bach, K. 1994. Conversational Implicature, *Mind & Language* (2): 124-162.
- Borg, E. 2012. *Pursuing Meaning*, Oxford University Press.
- Buchanan, R. 2010. A puzzle about meaning and communication. *Noûs* 44 (2): 340-371.
- Cappelen, H. and E. Lepore. 2005. *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*, Blackwell.
- Carston, R. 2002. *Thoughts and utterances: The pragmatics of explicit communication*, John Wiley & Sons.
- Clapp, L. 2012. Indexical Color Predicates: Truth Conditional Semantics vs. Truth Conditional Pragmatics, *Canadian Journal of Philosophy* 42(2): 71-100.
- Dobler, T. 2019. Occasion-sensitive semantics for objective predicates. *Linguistics and Philosophy* 42 (5): 451-474.
- Fodor, J. 2003. *Hume variations*, Cambridge University Press.
- Gauker, C. 2008. Zero tolerance for pragmatics, *Synthese* 165(3): 359-371.
- Grindrod, J., J. Andow and N. Hansen. 2018. Third-person knowledge ascriptions: A crucial experiment for contextualism. *Mind and Language* 34(2): 158-182.
- Grindrod, J. and E. Borg. 2019. Questions Under Discussion and the Semantics/Pragmatics Divide. *Philosophical Quarterly* 69 (275): 418-426.

---

<sup>17</sup> For very helpful comments on previous drafts of this paper, I am grateful to Manuel Pérez Otero, Inés Crespo, Andreas Heise and two anonymous referees. Funding: project FFI2016-80088-P, Spanish Ministry of Economy, Industry and Competitiveness.

- Groenendijk, J. and Stokhof, M. 1984. *Studies on the semantics of questions and the pragmatics of answers*. Ph.D. thesis, University of Amsterdam.
- Gross, S. 2005. Context-sensitive truth-theoretic accounts of semantic competence. *Mind and Language* 20 (1): 68–102.
- Hansen, N. and E. Chemla. 2013. Experimenting on Contextualism. *Mind and Language* 28 (3): 286–321.
- Perry, J. 1986. Thought Without Representation, *Proceedings of the Aristotelian Society* 60(1986): 137–151.
- Picazo, C. 2019. Are mental representations underdeterminacy-free? *Synthese* 196: 633–654.
- 2020. Occasion-Sensitivity and What Is Said, in T. Ciecierski and P. Grabarczyk (eds.) *The Architecture of Context and Context-Sensitivity*. Springer. pp. 227–243.
- Pietroski, P. 2003. The character of natural language semantics, in A. Barber (ed.), *Epistemology of Language*, Oxford University Press, pp. 217–256.
- Recanati, F. 2004. *Literal Meaning*, Cambridge University Press.
- 2007. *Perspectival Thought: A Plea for (Moderate) Relativism*, Clarendon Press.
- 2010. *Truth-Conditional Pragmatics*, Oxford University Press.
- Roberts, C. 2012. Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics. *Semantics and Pragmatics* 5: 1–69.
- Rothschild, D. and G. Segal. (2009). Indexical predicates, *Mind and Language* 24(4): 467–493.
- Sainsbury, M. 2001. Two ways to smoke a cigarette, *Ratio* 14(4): 386–406.
- Searle, J. 1978. Literal Meaning, *Erkenntnis* 13: 207–224.
- 1980. The background of meaning, in J. R. Searle, F. Kiefer and M. Bierwisch (eds), *Speech act theory and pragmatics*, Springer, pp. 221–232.
- 1992. *The Rediscovery of the Mind*. MIT Press.
- Schoubye, A. J. and A. Stokke. 2016. What is said?, *Noûs* 50(4): 759–793.
- 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. 1989. *The Uses of Sense: Wittgenstein's Philosophy of Language*, Oxford University Press.
- 1991. Annals of Analysis, *Mind* 100(2): 237–264.
- 1997. Pragmatics, in B. Hale and C. Wright (eds), *A Companion to the Philosophy of Language*, Blackwell, pp. 87–107.
- 2000. *Unshadowed Thought: Representation in Thought and Language*, Harvard University Press.
- 2006. Insensitive semantics, *Mind and Language* 21(1): 39–49.
- 2008. *Occasion-Sensitivity: Selected Essays*, Oxford University Press.