

Occasion-sensitive semantics

Tamara Dobler

University of Amsterdam,
Institute for Logic, Language, and Computation
t.dobler@uva.nl

Abstract. In this paper we provide a formal model for occasion-sensitive semantics motivated by so called ‘Travis cases’ (Travis 1978, 2000, 2008, 2009). We suggest that understanding of an utterance of φ (knowing its truth conditions) can be modelled as a twofold partitioning of worlds in logical space, where the initial partition is induced by the context-invariant meaning of the sentence uttered and the latter on the basis of context-dependent goals. Our model uses only a single parameter to capture occasion-sensitivity of sentences: *practical goals*.¹

1 Introducing occasion-sensitivity

Context-sensitivity is a phenomenon that affects natural language expressions. Frege already pointed out that some sentences express thoughts only when certain features of context are known, such as the time of an utterance or the speaker. The case in point are sentences containing context-sensitive terms such as indexicals or demonstratives (see Frege 1956: 296). For Frege, as for many others, context-sensitivity seemed a fairly isolated phenomenon, affecting only a small class of expressions. More recently, however, *radical contextualists* (e.g. Travis 2008, Recanati 2004, Carston 2002) have suggested that context-sensitivity is pervasive, that it affects all open class expressions (i.e. not just pure indexicals), and that it does so in a way that classical truth-conditional semantic theories may find difficult to track.

¹ This paper will be enclosed in LNAI volume Modelling and Using Context (10th International and Interdisciplinary Conference, CONTEXT2017, Paris, France, June 20-23. Proceedings) and should be quoted from there.

1.1 Motivation: Travis cases

Radical contextualism is motivated by examples such as following:²

The leaves are green³

State of affairs: Pia's Japanese maple is full of russet leaves. Pia paints them green.

Zoe needs some green leaves for her decoration:

Zoe: Do you have green leaves?

Pia: These leaves are green.

Pia's botanist friend seeks green leaves for a study of green-leaf chemistry:

Botanist: Do you have green leaves?

Pia: # These leaves are green.⁴

Sid has a desk⁵

State of affairs: Sid, an impoverished student, uses a door over two stacks of milk crates as a desk to write on.

Concerned if Sid has a desk to write on:

Pia: Does Sid have a desk?

Max: Sid has a desk.

On the look out for high end furniture:

Pia: Does Sid have a desk?

Max: # Sid has a desk

The shoes are under the bed⁶

State of affairs: Pia is looking for her shoes. Sid sees them, heels protruding from beneath the bed.

Retrieving shoes:

² In the literature these are known as *Travis cases* (sometimes referred to as context-shifting arguments, see Cappelen and Lepore (2008). For critical discussion see Hansen (2011), Hansen and Chemla (2013), Predelli (2005), Vicente (2012), Vicente (2015), Kennedy and McNally (2010), Rothschild and Segal (2009), MacFarlane (2009)

³ Travis (2008): 111

⁴ We use the sign # to indicate that the response is in some sense inadequate or infelicitous, and that what is said is intuitively false.

⁵ Travis (2000)

⁶ Travis 2009: 119-120

Pia: Are the shoes under the bed?
Sid: The shoes are under the bed.

Making sure the shoes are well hidden. Sid's response is meant to reassure Pia that the shoes are well hidden.

Pia: Are the shoes under the bed?
Sid: # The shoes are under the bed.

1.2 What Travis cases are taken to show

Travis cases are taken as evidence for a number of claims. First, they demonstrate what we deem to be an uncontroversial fact that the interpretation and evaluation of utterances are somehow sensitive to contextual factors. However, the reason why Travis cases are particularly interesting is because the relevant sentences contain *no* pure indexicals (such as *I, you, here, today*) or other standard context-sensitive expressions (e.g. gradable adjectives, predicates of personal taste, modals, quantifier phrases), and yet their interpretation and evaluation varies across contexts of utterance. Second, in each example, the sentence type used in both contexts is the same, and it is assumed that *its conventional meaning remains invariant*. One important consequence of the fact that the semantics of sentence types is invariant is that it would be impossible to avoid the denial of the law of non-contradiction, since the same sentence appears to be evaluated as true and false of the same state of the world, *if*, that is, the sentences used in these examples are taken at their face value (namely, as not being context-dependent in some sense). Third, the fact that the interpretation of utterances varies in this way is typically taken as evidence for the variability of intuitive *truth-conditions* or *truth-conditional content of utterances*.⁷

A number of analyses of Travis cases have been proposed, many whose primary concern is to preserve a compositional account of truth-conditions.⁸ With the exception of semantic minimalists who reject that the intuitions about truth-value of utterances are good evidence for truth-conditions of sentences, almost all other solutions assume that at least some of the constituents in these sentences have context-dependent interpretations. Thus the appeal to context-dependence seems by far the most common way to avoid the attribution of contradictory

⁷ See Recanati (2004) and Travis (2006) for a defense of this thesis; see Cappelen and Lepore (2008), Borg (2004) and Fodor (2003) for criticism

⁸ The phenomenon exhibited in Travis cases has thus far been analysed as a form of lexical ambiguity or polysemy (Fodor 2003; Vicente 2015; Carston 2010), structural ambiguity (Kennedy and McNally 2010), conversational implicature (Cappelen and Lepore 2008; Fodor 2003), circumstance of evaluation dependence (MacFarlane 2009; Predelli 2005), pragmatic modulation (Pagin 2005, Recanati 2012, Carston 2010) and (hidden) indexicality (Szabó 2001; Rothschild and Segal 2009)

beliefs to speakers or the denial of the law of non-contradiction, whilst also preserving truth-conditional compositional semantics.

1.3 Non-indexical contextualism

In this paper we aim to take a different approach because we believe that occasion-sensitivity is different from indexicality, and also that it is not a threat to compositional semantics. Thus, although we will propose what is, broadly speaking, a *contextualist* treatment of occasion-sensitivity, unlike Rothschild and Segal (2009) for instance, we won't treat individual lexical items (e.g. *green*) as (non-standard) full-blown *indexical* predicates. We fully agree with Kennedy and McNally (2010) that an indexical account which specifies no rule for determining the contextual value of an indexical (e.g. the speaker in context for *I*) places *no constraint* on its possible interpretations, and potentially creates the overgeneration problem (see Stanley 2005). Furthermore, unlike Szabó (2001), we will refrain from postulating hidden variables in the logical form of colour adjectives, mainly for the reasons already pointed out in Rothschild and Segal (2009) and Kennedy and McNally (2010), namely, because such an analysis lacks the required generality.

The approach defended in this paper is formally and conceptually closest (albeit not identical) to the position known as *non-indexical contextualism* developed in MacFarlane (2007, 2009). We agree with MacFarlane that context-sensitivity is a broader phenomenon than indexicality and it shouldn't be conflated with the latter. We also think that occasion-sensitivity is a good example of non-indexical context-sensitivity, where (i) semantic content of a predicate is invariant, (ii) the world of evaluation is invariant, and yet (iii) the extension of the predicate (and the truth-value of an utterance) varies. This is possible, MacFarlane argues, because the circumstance of evaluation includes *more parameters* than just a possible world (where these parameters are initiated by the context of use rather than the context of assessment).

With regards to Travis cases, MacFarlane's suggestion is that predicates such as *green* or *weighs 80 kilos* invariantly express a property of *being green* or *weighing 80 kilos* yet such a property has an intension which is a function from the circumstance of evaluation – consisting of a *pair* of parameters (counts-as parameter and world parameter) rather than one – to extension. A non-standard, all encompassing *count-as parameter* “settles what things have to be like to have various properties: e.g. the property of weighing 160 pounds, or of being tall” (MacFarlane 2009:) such that the extension of a predicate varies when the value of the counts-as parameter varies, even when the world of evaluation remains the same.

We believe that this picture, generally speaking, adequately captures the phenomenon illustrated in Travis cases. However, we also contend that because of its generality it potentially suffers from similar problems as

Rothschild and Segal (2009), in particular to do with lack of constraint and over-generation. Our aim, therefore, is to further develop this approach, and revise it where necessary, in order to explain which features of occasions of use guide our judgements as to when something would count as having a certain property, and how these features constrain possible interpretations and evaluations of sentences.

Our proposal departs from MacFarlane's on several points (although the first two may be only terminological departures): first, we take it that properties that feed into the counts-as function can also be *zero-place properties* such as the property of *these leaves being green* or *Sid's having a desk*; second, we don't consider a zero-place property to be the same as the proposition expressed by an utterance on occasion (i.e. intuitive truth-conditions); we consider the proposition expressed by an utterance on occasion (i.e. the Austinian proposition)⁹ to be an *output* of, rather than an input to, the counts-as function. Third, whereas MacFarlane leaves it open which feature of a context drives modifications of properties and why such modifications would ever take place, we suggest that a specific contextual feature mandates shifts in when things count as having certain properties.¹⁰ By considering pragmatic processes as not linguistically mandated (rightly so), existing non-indexicalist approaches fail to put enough emphasis on why, in the course of communication, we nevertheless *do* systematically modify standing semantic contents of our expressions. We suggest that Travis cases point towards one specific contextual factor as playing part in fixing when things would count as having a certain property: relevance (of the information) for the *goals* that speakers pursue within particular practical projects. Even though objective predicates have their context-invariant semantic contents, these contents are *uninformative* unless they are restricted so that they become valuable for achieving contextually salient goals.

2 The role of standing meaning

We will assume that each declarative¹¹ sentence φ has a context-invariant, compositional meaning in virtue of syntactic and lexical structure and linguistic

⁹ See Recanati (2007)

¹⁰ Reference is usually made to the speaker's intentions or the topic of conversation, without going into details of how or why these are able to determine correctness of selecting a certain function and not other. According to MacFarlane, "the counts-as parameter will be determined in complex ways by other features of the context, including the topic of conversation and the speaker's intentions" (MacFarlane 2007: 246).

¹¹ The model we are about to propose can be extended to interrogatives and imperatives, although this is beyond the scope of this paper.

conventions. On our view, standing meaning has a substantial role to play in the process of utterance interpretation. However, in the light of pervasive context-sensitivity of natural language, we believe it is crucial to capture the idea that the meaning of φ only *constrains* the intended interpretation, or the proposition(s) expressed, without fully determining it, i.e. the meaning of φ *underspecifies* the proposition that some utterance of φ expresses.

To capture this notion of semantic constraint on what is said, we suggest that the conventional meaning of an (atomic) declarative sentence can be viewed as the set of worlds which are *compatible* with its meaning, that is, those worlds where the sentence is true *on some of its uses* but false on others. For any atomic (NP+VP) sentence (used to state a piece of information) there will be: (i) some possible worlds or situations which are incompatible with what it means, and which can thus be eliminated purely with respect to its conventional meaning, and (ii) others which are compatible with the meaning of the sentence but in which an *utterance* of the sentence is not always true.

In standard truth-conditional frameworks, a world is considered to be compatible (incompatible) with the meaning of φ only when φ is true (or false) at that world. Accordingly, the set of worlds in \mathcal{W}_φ would thus correspond to the proposition expressed by φ , modelled as the set of worlds in which φ is true. However, as evidenced in Travis cases, this common identification of meaning with truth conditions is problematic insofar as the truth valuation of an utterance of φ may shift even when the world of evaluation and conventional meaning remain unchanged. So we believe it is crucial to resist the common assumption that conventional meaning is sufficient to determine when a sentence uttered in context would be true. Hence, on our view, the meaning function does not partition the worlds into those at which φ is true and those at which φ is false. Rather φ is considered to be truth-evaluable and worlds are considered to be truth-makers only after taking into account a further factor: an *occasion* on which φ is used.

Take **The shoes are under the bed** Travis case. The sentence *The shoes are under the bed* may be true if the shoes are well hidden under the bed, or if their heels are poking out. This sentence could also be true if the shoes are three floors down aligned with the bed. Or if they are buried in cement under the bed etcetera. However, an *utterance* of the sentence on an occasion won't necessarily be true at all these worlds: this is the lesson we learned from Travis cases. Even though all these very different conditions of the world are compatible with what the sentence means not all of them will make an utterance of the sentence on some occasion true since an occasion places some further constraints on success (i.e. truth) of the utterance.

Let us call this set generated by the meaning function *the compatibility set of φ* , or \mathcal{W}_φ . If a world is semantically *incompatible* with the meaning of φ , then there will be no possible understanding on which φ is true at that world (where an *understanding* of φ is a refinement of \mathcal{W}_φ). If a world is semantically

compatible with φ then there will be some understanding on which an utterance of φ is true at that world, but also some understandings on which it is false at that world.

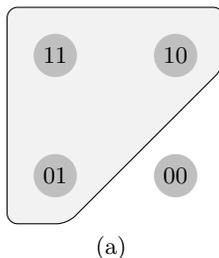


Fig. 1. Semantically compatible worlds, \mathcal{W}_φ

For simplicity's sake, let us assume there are only two possible ways of being green: *painted* green and *naturally* green. **Figure 1** represents the set of all alternatives, where (11) is the world where the salient leaves are both painted and naturally green, (10) is the world where they are painted green but not naturally green, (01) is the world where they are only naturally green, and (00) is the world where they are neither painted nor naturally green. Since in our simplified model painted and naturally green exhaust the ways of being green, in the (00) world the leaves are *not* green. The set of worlds which are semantically compatible with the sentence *The leaves are green* is thus $\mathcal{W}_{green} : \{11, 10, 01\}$, whilst the set of semantically incompatible worlds is $\overline{\mathcal{W}}_{green} : \{00\}$.

Travis cases show that not all the worlds in \mathcal{W}_{green} will make an utterance of *The leaves are green* on some occasion true: for instance, in the botanist context, the world (10) is ruled out as a truth-maker because it matters in this context that the leaves are naturally green (see **Figure 2**). By contrast in the decorator context, it does not matter if the leaves are painted or natural, so the only world ruled out is $\{00\}$.

Hence it depends on which understanding of green is relevant, which worlds will make an utterance of φ true. The resolution of the issue which understanding of φ is relevant for truth evaluation is here independent from its meaning. To fix which understanding of a predicate is correct we need to ask which information (about the world) would be relevant to an agent on some occasion. In the next section we suggest that the worlds in \mathcal{W}_φ are kept or eliminated on the basis of how *conducive* they are to the *goal* that the agent aims to achieve on an occasion.¹²

¹² In our diagram the world (01) is thus not fit for the botanist's goal.

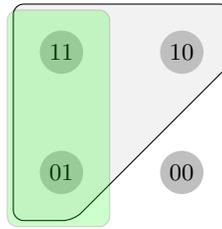


Fig. 2. Understanding of green in the botanist context

3 Occasions and goals

So far we used the term *occasion* and *occasion-sensitivity* without properly defining them, or showing what the relation between an occasion and a world is supposed to be. In this section we will try to give a more detailed account of these notions by identifying a feature that we consider necessary for determining which proposition is expressed by an utterance of a sentence on occasion. We shall suggest that a key feature of some occasion of using φ that gives rise to the phenomenon of occasion-sensitivity is a *practical goal* which is salient of that occasion and which agents strive to achieve. Recognising which goal is salient is necessary to determine which proposition someone expressed by an utterance of φ .

3.1 Communication goals understood narrowly and broadly

What does a practical goal have to do with determining when a sentence is true? In this section we try to motivate the connection between goal directed actions and correct understanding of utterances.

To get the idea that practical goals constrain interpretation off the ground, we take the following remark made by Robert Stalnaker (who cites Grice) as our guiding principle:

“[W]e should see speech as action to be explained, like any other kind of action, in terms of the beliefs and purposes of the agent” (Stalnaker 1999: 2).

To Stalnaker’s principle we add the following observation about the success-conditions of actions: namely, that an action (including a speech action)

is successful when it successfully achieves its goal. But what constitutes the goal of a speech act?

Ever since Frege (1956) it is common to distinguish two elements in a speech act: its *content* (truth-evaluable thought) and its *force*. The idea of (illocutionary) force is normally understood as capturing what a speech act aims at; an assertion characteristically aims at truth or saying something true, a question characteristically aims to raise an issue to be resolved by a (truthful) answer, and an order aims to issue an order to be successfully fulfilled. Saying something true, raising an issue, issuing an order etc. are thus *goals* that characterise different speech acts. Are *these* goals then sufficient to determine (together with the sentence meaning) when a particular speech act is successful?

Our answer to this question is negative because we believe that goals (and hence successes) that agents aim to achieve in the course of communication are typically *much broader* (and much more practical) than the characteristic goals of assertions, questions, or commands. That is, by virtue of performing an *assertive* speech act, an agent does aim to say something true and, if true, add the proposition to the common ground (see Stalnaker 1999). Nonetheless, as a rule, this is *not the only thing* that agents do whilst communicating: an issue raised and proposition asserted are always turned towards fulfilments of other *practical goals* such as hiding shoes from kleptomaniac friends, doing green leaf chemistry, decorating, putting kids to sleep, building houses and bridges, and so forth.¹³ So, besides narrowly defined goals that characterise (conventional) illocutionary forces, there are various practical goals we aim to achieve on particular occasions. Our suggestion is that practical goals are another necessary factor in determining when a speech act (of assertion) is successful, and, in particular, when its content is true.

If communication, in general, is embedded in other goal-directed, practical activities, then we can also assume that what someone communicates on a particular occasion is directed at fulfilling some such goal. To take a Travis case again as our example, if I know that my friend is a botanist and is searching for green leaves for her experiment (assuming I am not totally ignorant about this topic) then I also should know that her request for green leaves bears on that particular goal and that so does my response. If I sincerely say to my friend, “Search no more, the leaves on my tree are green”, then what I say is that the world is indeed such that her project can be successfully carried out. In other words, by asserting this, I suggested that the world is *conducive* to my friend’s goal. If it turns out, however, that I have the leaves that are painted green but are naturally red, then her reasonable expectations as to the state of the world won’t be satisfied and her project won’t take off. That is, *this* state of the world is *not* conducive to the salient goal and so what I said to my friend was false (although it could have been true were her goal decorating rather than chemistry).

¹³ This idea of communication is based on Wittgenstein et al. (2009).

It is in this way that we take practical goals to be an important factor in interpreting and evaluating what someone has said.

3.2 Goal-conduciveness of a world

Agents communicate to exchange or gain the information about the world which helps them achieve their goals, which, it was argued, outstrip narrow goals that define illocutionary forces of particular speech acts. We also suggested that the state of the world can be such that it is more or less conducive to fulfilling some practical goal, and so, when we provide some information to someone about the state of the world, it goes without saying that this will have certain practical consequences, i.e., some positive or negative effect on salient plans and goals. For instance, the world where a room is such that the shadows of objects can be seen in it is *not* conducive to the goal of opening a camera and developing a film since such a project won't be successfully achieved in this state of the world. Consequently, if someone says that the room is dark, where this particular goal is salient in the conversation, then he said something false. If, however, he says this when the salient goal is to put a child to sleep, then he might well be saying something true.

Our main suggestion is that the worlds which are compatible with the meaning of a sentence (the members of \mathcal{W}_φ) may include those with very different potentials in terms of enabling a successful action (some very low, some high). For a particular goal g and world w we can take it that w will be more or less conducive to g (i.e. achieving g in w is more or less difficult). We will give a formal definition of this notion below.

We propose that, among the worlds that are semantically compatible with the sentence uttered (those belonging to the compatibility set), only those worlds that are conducive to the salient goal on an occasion constitute the proposition expressed by an utterance of the sentence on that occasion.

4 A model

In order to determine which proposition is expressed by an utterance of the sentence on some occasion we propose to first evaluate all possible worlds in terms of how goal-conducive they are, and to assign each a numerical value, thus producing a ranking over worlds. The set of worlds with the highest rank that is also semantically compatible with the sentence will correspond to the proposition expressed by the utterance of that sentence on that occasion.

4.1 Language

Because we take simple NP+VP sentences as our atom of analysis, we will use

a language L for *propositional logic*, with its standard set of propositional letters and connectives. We assume the set of atomic sentences A and define formulas as follows:

$$\varphi := p, \varphi \vee \varphi, \neg \varphi, \varphi \wedge \varphi, \varphi \rightarrow \varphi$$

where $p \in A$, the set of atomic sentences.

4.2 Semantics

A *model* M for L is a tuple $M = \langle W, G, V, \mathcal{C}, \rangle$ in which W is a non-empty set of possible worlds, G is a non-empty set of goals, V is a valuation function which assigns to a sentence a numerical value 1 for semantically compatible or 0 for semantically incompatible, \mathcal{C} is a goal-conduciveness function which maps a pair of worlds and goals to a value in the interval $[0, 1]$ (\approx the c-value of the world w given the goal g).

Definition: conduciveness to goal

$$\mathcal{C} : G \times W \rightarrow [0, 1]$$

where $g \in G$ is a goal that agents are aiming to reach on occasion.

Definition: conduciveness ordering

$$\forall w, w' \in W, w' \preceq_g w \text{ iff } \mathcal{C}(w', g) \leq \mathcal{C}(w, g)$$

We can produce a ranking of all worlds in logical space with respect to how conducive they are to some goal. For any pair of possible worlds $\{w, w'\}$ the world with the lower conduciveness value will be ranked lower than the world with the higher conduciveness value. If they have the same value they will also have the same rank.

Definition: the top tier

$$w \in \text{MAX}_g \text{ iff } \forall w' \in W, w' \preceq_g w$$

The set of worlds with the highest ranking, MAX_g , will form the top tier of the worlds wrt to their goal-conduciveness. For any world w with a given ranking one can conceive of another world w' that has the same ranking despite differing in some way from w . So, in principle, there will also be more than one world with the MAX_g value.

We were able to order the worlds in terms of their conduciveness to a given goal quite independently of language. This is because a successful action that brings about some goal does not necessarily presuppose a use of language (i.e. a performance of a speech act). However, since our primary interest is in the

interpretation and evaluation of sentences used in communication (i.e. speech acts), our model should capture how pragmatic constraints on interpretation we defined by means of contextual goals interact with semantic constraints we defined by means of the initial partition induced by conventional meaning.

Definition: valuation function

$$V : A \times W \rightarrow \{0, 1\}$$

where A is the set of atomic sentences, and values 1 and 0 designate *semantically compatible* and *semantically incompatible*, respectively (NOT: true and false).

Definition: valuation of formulas

$$\begin{aligned} M, w \models_g \varphi & \text{ iff } w \in \text{MAX}_g \wedge V(\varphi, w) = 1 \\ M, w \models_g \neg\varphi & \text{ iff } M, w \not\models_g \varphi \\ M, w \models_g (\varphi \wedge \Psi) & \text{ iff } M, w \models_g \varphi \text{ and } M, w \models_g \Psi \end{aligned}$$

The proposition expressed by an utterance of an atomic sentence p is defined in model M as the intersection of \mathcal{W}_φ , the set of worlds that are semantically compatible with the sentence, and the set of worlds that have the MAX_g value.

5 Examples

We can now analyse Travis cases using our present semantics in the form of the following conditional T-sentence:

- (T) The utterance u of the sentence φ is true at world w relative to goal g iff w is conducive to g and w is semantically compatible with φ .

The leaves are green: In this example we have a shift of goals between the two occasions: on one occasion the salient goal is to decorate a table centrepiece, in which case the natural property of the leaves is not relevant; on another occasion the salient goal is to examine the green leaf chemistry, in which case the natural property is relevant. The stipulated world of evaluation is such that the leaves are only painted green (but are naturally red). This world is semantically compatible with the uttered sentence on both occasions, but is only conducive to the former goal and not to the latter. Therefore, the utterance is true at w only wrt to the goal of decorating and is false wrt the goal of doing chemistry.

- (O1) u_{green} is true at w wrt $g_{decorate}$
(O2) u_{green} is false at w wrt $g_{chemistry}$

The shoes are under the bed: Again we have a shift of goals on the two occasions. On one occasion the goal is to retrieve the shoes and the extent to which they are visible is not relevant to a successful achievement of this goal. On another occasion the goal is to make sure the shoes are well hidden, so not visible at all. The world of evaluation is such that the heels of the shoes are visible from under the bed. Saying that the shoes are under the bed is conducive only to the former goal, whilst saying this on the second occasion would prevent the achievement of the latter goal. Therefore, the information conveyed by the sentence is true with respect to the former goal and false wrt to the latter.

(O1) u_{shoes} is true at w wrt $g_{retrieve}$

(O2) u_{shoes} is false at w wrt g_{hide}

Sid has a desk: In this scenario, each occasion of utterance is paired with a different goal such that on one occasion the goal is to have something to sit and write on, and on another the goal is to find a piece of expensive office furniture. The world of evaluation is such that Sid has a door over two stacks of milk crates. Saying that Sid has a desk in the first case facilitates the fulfilment of the goal, whilst in the latter case it probably doesn't. Hence, the information conveyed by the utterance is true on the former occasion and false on the latter.

(O1) u_{desk} is true at w wrt g_{write}

(O2) u_{desk} is false at w wrt g_{office}

6 Conclusion

This completes our model for occasion-sensitive semantics. It is important to bear in mind that we do *not* claim that *all* meaning is occasion-sensitive insofar as we assumed that the conventional meaning of a sentence creates an initial partition which is not sensitive to any *particular* occasion. The outcome of the meaning partition corresponds to what is in the literature known as *minimal proposition*.¹⁴ Unlike proponents of the notion of minimal proposition, we take intuitions behind Travis cases as evidence for the variability of intuitive truth-conditions across occasions. Although we decided to sidestep the terminological side of this debate by stipulating that occasion-sensitive utterances of φ are successful (and not merely true), the substantial point of Travis cases and occasion-sensitivity nonetheless remains: 'minimal proposition' is not only something with no 'psychological reality' (Recanati 2004) but moreover it plays no significant role in decision making, planning and action. And if we are interested

¹⁴ See Cappelen and Lepore 2008, Borg 2004, Recanati 2004, Pagin and Pelletier 2007, Travis 2006

to capture what propositions our utterances express in the course of communication, rather than in their structural semantic properties, then the analysis which stops at the initial meaning partition is incomplete, to say the least. We have shown that only once we take into account particular goals (and also agents' expectations) we are able to specify the subset of semantically appropriate worlds which corresponds to the proposition expressed on occasion.

Bibliography

- Austin, J. L. (1975). *How to do things with words*. Oxford university press.
- Borg, E. (2004). *Minimal semantics*. Cambridge University Press.
- Borg, E. (2012). *Pursuing meaning*. Oxford University Press.
- Cappelen, H. and Lepore, E. (2008). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*. John Wiley & Sons.
- Carston, R. (2002). *Thoughts and utterances: The pragmatics of explicit communication*. John Wiley & Sons.
- Carston, R. (2010). Explicit communication and freepragmatic enrichment. In *Explicit communication*, pages 217–285. Springer.
- Fodor, J. A. (2003). *Hume variations*. Cambridge Univ Press.
- Frege, G. (1956). The thought: A logical inquiry. *Mind*, pages 289–311.
- Hansen, N. (2011). Color adjectives and radical contextualism. *Linguistics and Philosophy*, 34(3):201–221.
- Hansen, N. and Chemla, E. (2013). Experimenting on contextualism. *Mind & Language*, 28(3):286–321.
- Kennedy, C. and McNally, L. (2010). Color, context, and compositionality. *Synthese*, 174(1):79–98.
- MacFarlane, J. (2007). Semantic minimalism and nonindexical contextualism. *Context-sensitivity and semantic minimalism: New essays on semantics and pragmatics*, pages 240–250.
- MacFarlane, J. (2009). Nonindexical contextualism. *Synthese*, 166(2):231–250.
- Pagin, P. (2005). Compositionality and context. *Contextualism in philosophy*, pages 303–48.
- Pagin, P. and Pelletier, J. (2007). Content, context and composition. *Content and Context. Essays on Semantics and Pragmatics*, pages 25–62.
- Pietroski, P. (2005). Meaning before truth.
- Predelli, S. (2005). Painted leaves, context, and semantic analysis. *Linguistics and Philosophy*, 28(3):351–374.
- Recanati, F. (2004). *Literal meaning*. Cambridge University Press.
- Recanati, F. (2007). *Perspectival Thought: A Plea for (Moderate) Relativism: A Plea for (Moderate) Relativism*. Oxford University Press.
- Recanati, F. (2012). Compositionality, semantic flexibility, and context-dependence. *Oxford Handbook of Compositionality*, pages 175–192.
- Rothschild, D. and Segal, G. (2009). Indexical predicates. *Mind & Language*, 24(4):467–493.
- Stalnaker, R. (1999). *Context and content: Essays on intentionality in speech and thought*. Oxford University Press.
- Stanley, J. (2005). Semantics in context. *Contextualism in philosophy: Knowledge, meaning, and truth*, pages 221–253.
- Szabó, Z. G. (2001). Adjectives in context. *PRAGMATICS AND BEYOND NEW SERIES*, pages 119–146.
- Travis, C. (1978). Meaning versus truth. *Dialogue*, 17(03):401–430.

- Travis, C. (2000). *Unshadowed thought: Representation in thought and language*. Harvard University Press.
- Travis, C. (2006). Insensitive semantics. *Mind & Language*, 21(1):39–49.
- Travis, C. (2008). *Occasion-sensitivity: Selected essays*. Oxford University Press.
- Travis, C. (2009). *Thought's footing: a theme in Wittgenstein's Philosophical investigations*. Oxford University Press.
- Vicente, A. (2012). On travis cases. *Linguistics and philosophy*, 35(1):3–19.
- Vicente, A. (2015). The green leaves and the expert: Polysemy and truth-conditional variability. *Lingua*, 157:54–65.
- Wittgenstein, L., Anscombe, G. E. M., Hacker, P. M. S., and Schulte, J. (2009). *Philosophische Untersuchungen: Philosophical investigations*. Wiley-Blackwell.