

Occasion-sensitive semantics for objective predicates*

Tamara Dobler

April 6, 2017

Abstract

In this paper we give a semantics for objective predicates that takes into account the phenomenon of occasion-sensitivity associated with so called *Travis cases*. A key proposal is that the set of truth-makers for a declarative sentence consists of worlds (or situations) with different capacities to support a particular goal-directed course of action. This capacity of goal-conduciveness decides whether any given world is *accepted* as a truth-maker for some utterance of the sentence on a particular occasion. We argue that a world will not be accepted in case it is not conducive to a contextually salient goal, despite it being a member of the truth-maker set for a sentence. Occasion-sensitivity in the interpretation and evaluation of sentences containing objective predicates is then explained by reference to practical goals of activities that are salient on an occasion of utterance. We propose a formal model for Travis cases inspired by partition semantics for questions (Groenendijk and Stokhof 1984), and we consider two applications: to disagreement and to cancellability.

1 Introduction

1.1 Objective and discretionary predicates

This paper presents a novel approach aimed at explaining the irreducible context-sensitivity of objective predicates. By objective predicate, we refer to a broad range of predicates such as *green*, *desk*, or *under the bed*, which can be contrasted with *subjective* or *discretionary* predicates (predicates of *personal taste*) such as *tasty* or *fun* (cf. Lasersohn 2005, Kölbel 2004, Coppock 2016). To properly appreciate the difference between these types of predicates, consider the following sentences:

- (1) The leaves are green.
- (2) Sid has a desk.

*Acknowledgements to be added later

- (3) The chilli is tasty.
- (4) Roller coasters are fun.

Sentences (1) and (2) contain the objective predicate *green* and *desk*, while sentences (3) and (4) contain the discretionary predicates *tasty* and *fun*. (1) and (2) are used as statements concerning “matters of fact” whilst (3) and (4) concern “matters of opinion”. One important intuitive difference between them is that two contradictory statements containing discretionary predicates can both be true at the same world of evaluation. Consider the following two dialogues:

- (5) The chilli dialogue
 - (a) John: The chilli is tasty.
 - (b) Sid: No, the chilli is not tasty.
- (6) The green leaves dialogue
 - (a) John: The leaves are green.
 - (b) Sid: No, the leaves are not green.

Whereas both (5a) and (5b) can be true without either John or Sid being at *fault*, if (6a) is true at a world (6b) must be false, on pain of contradiction.¹ This is because, unlike discretionary predicates whose extension varies with respect to something more (or other) than a possible world parameter – a judge, individual, perspective, outlook, centred world, context of assessment, standard of taste etc. – the extension of an objective predicate is typically taken to vary only relative to a world of evaluation. Thus, in contrast to statements concerning personal taste or opinion, factual (objective) statements ought to be true at a world *simpliciter* without further relativisation to judges or other perspectival parameters that play a role in fixing the extension of discretionary predicates.

Another way to put the point is that knowing the meaning of a discretionary predicate is *not* enough to assure the agreement as to when a discretionary statement would be true at a given world of evaluation: it matters for its truth value who evaluates and from which perspective. By contrast, *anyone* who knows what the sentence *The leaves are green* means, and knows which leaves are being referred to, should agree what conditions must obtain for it to be true, *viz.* it is true if and only if the leaves are green.

1.2 Travis cases

And yet, there is a class of cases that might bring this intuitively correct classification into question. So called ‘Travis cases’ (Travis 1978, 2000, 2008, 2009) show that truth valuations of factual statements could vary even when the world

¹For discussion of faultless disagreement see Kölbel (2004), MacFarlane (2014), Stojanovic (2007), Huvenes (2014), Coppock (2016).

of evaluation remains the same, and that disagreements could arise as to whether a factual statement is true.² Consider the following dialogues taking place on different occasions:

The leaves are green³

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

Occasion 1: Zoe needs some green leaves for her decoration:

(7a) Zoe: Do you have green leaves?

(7b) Pia: These leaves are green.

Occasion 2: Zoe is a botanist seeking green leaves for a study of green-leaf chemistry:

(8a) Zoe: Do you have green leaves?

(8b) Pia: # These leaves are green.⁴

(8c) Zoe: No, these leaves are not green.

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.

Occasion 1: Concerned if Sid has a desk to write on:

(9a) Pia: Does Sid have a desk?

(9b) Max: Sid has a desk

Occasion 2: On the look out for high end furniture:

(10a) Pia: Does Sid have a desk?

(10b) Max: # Sid has a desk

(10c) Pia: No, Sid does not have a desk

So we have a dilemma. Although we can agree on what conditions must obtain for sentences containing objective predicates to be true at a world (based exclusively on our knowledge of meaning and of what leaves are being referred to), we can nevertheless get into disagreements as to whether the sentence is

²We assume throughout that a factual statement amounts to an utterance of a non-ambiguous, non-elliptical, non-indexical sentence containing an objective predicate.

³Travis (2008): 111

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

⁵Travis (2000)

true at a given world. Following Charles Travis let us call the phenomenon observed in these cases *occasion-sensitivity*. Can we preserve the intuitive contrast between discretionary and objective predicates if objective predicates are occasion-sensitive in this way? How can we explain the fact that on one occasion a response in (b) is correct (the speaker seems to get the facts right) and on another it's not (the speaker seems to be at fault) without the world of evaluation undergoing any change?

A key aim in this paper will be to try to better understand the processes that impact interpretation in such a way as to bring about the dilemma outlined above. It will be argued that the dilemma arises due to the fact that speakers refine their understandings of sentences uttered on particular occasions of use so as to maximise their informativeness in an attempt to fulfil certain practical goals. Furthermore, it will be shown that, notwithstanding initial appearances, the occasion-sensitivity of objective predicates does *not* bring the objectivity of factual statements (and the objectivity of disagreement concerning matters of fact) into question.⁶

1.3 The current approach and the roadmap

Let us briefly outline the main points that are crucial to the approach defended in this paper. In an attempt to accommodate the intuitions elicited by Travis cases we are going to suggest that despite having an important role in understanding when a sentence (containing an objective predicate) would be true, our knowledge of standing meaning must be supplemented by other extralinguistic factors which, besides meaning, play a role in truth evaluations of particular utterances.⁷

Knowing the meaning of a sentence or knowing when the sentence is true (being able to distinguish worlds where the sentence is true from worlds where it is false) leaves room for a great deal of variation as to what worlds or situations that make the sentence true are actually like.⁸ We can thus have many alternative truth-makers for a sentence; i.e. a broad and highly diverse set of worlds that all make the sentence true. Now this fact is significant because the successful fulfilment of the activities and projects we are engaged in on different occasions of utterance might not be deemed achievable at all those worlds that would make a sentence true.

⁶Namely, we argue that factual statements, unlike statements of opinion, are *not* sensitive to *subjective* standards of assessment. That is, it doesn't seem to matter for the truth value of a factual statement which individual evaluates it or from which (subjective) perspective.

⁷This idea has been defended for more than four decades by Charles Travis and others. See Travis 1978, Travis 2008, Moravcsik 1994, Recanati 2004, Carston 2002, Austin 1975, Searle 1980, Chomsky 2000, Pietroski 2005.

⁸Here and throughout when we talk about different truth-makers for a sentence, the intended distinction is not the usual one between worlds and partial worlds (situations) but rather between worlds (or situations, states of affairs, of what have you) which are all sufficiently similar to constitute the sentence's truth-maker, yet also sufficiently dissimilar for this to matter in considering them a truth-maker for an utterance of that sentence on occasion. We discuss this idea at length in section 3.

A key point we shall argue for is that, to communicate successfully and effectively, only conveying information that is relevant to the activities we are engaged in, it is necessary to turn a blind eye to certain alternatives. Which particular alternatives are disregarded, and which are not, is determined by the goal toward which the activity that motivates interlocutors on a particular occasion is directed. If a particular truth-maker for the sentence is incompatible with a successful fulfilment of an activity – if it is not *goal-conducive*, as we will call this property of it – then it won't be *accepted* by rational, well informed agents as a truth-maker of the sentence, given that context. In other words, we will evaluate the utterance of the sentence in a given context false despite the fact that the world *is* one of the truth-makers for the sentence. In this way, pragmatics decides how a sentence ought to be understood if we are to communicate successfully in order to carry on with our projects.

After reviewing some of the existing approaches to occasion-sensitivity in section 2, in section 3 we propose an approach, inspired by partition semantics of questions (Groenendijk and Stokhof 1984), according to which the set of worlds at which a sentence is true can be partitioned according to different ways that objects in those worlds have a given property. The conventional, standing or literal meaning of a sentence determines a very broad range of worlds at which the sentence is true. Still, and crucially, this initial valuation won't suffice to explain when a rational agent will accept an utterance of the sentence on an occasion as true. To specify when the utterance is true other factors besides standing meaning will need to be taken into consideration.

Among extralinguistic factors, contextually salient *goals* are crucial in narrowing down the initial set of ways for objects to have a certain property to those most likely to support a particular course of action. In section 3.2 we introduce the notion of *goal-conduciveness*, which is a feature that agents assign to worlds on the basis of whether they judge them to be conducive to their goals. The idea is that to understand correctly what someone has asserted or asked using a particular sentence presupposes being aware of their goal, because goals help restrict the initial truth-maker set generated by standing meaning. That is, they determine which particular ways that the world could be, out of many compatible with standing meaning, are in the agent's purview on an occasion. Pragmatics thus follows the 'less is more' principle since the informativeness of utterances is increased by reduction (see 3.3).

After introducing the formal model (section 4), in section 5 we discuss two applications, namely, to intuitions about genuine and faultless disagreement (5.1) and to cancellability (5.2).

2 This approach compared to other approaches

2.1 Indexicalist and non-indexicalist approaches

One way to explain the data in Travis cases is to ascribe the variability in extension to the variability in intension. There are two general ways of doing this:

indexicalist and non-indexicalist. On indexicalist approaches, objective predicates occurring in Travis cases behave just like pure indexicals (*I, now, here*) in that they don't have a stable, context-invariant content but express different properties in different contexts (Rothschild and Segal 2009). Or, alternatively, objective predicates are taken to contain some hidden indexical elements in their logical form whose values are provided by context (Szabó 2001). In both cases, the process of interpretation of predicates is automatic and mandatory, *viz.* indexical elements must be saturated for there to be interpretable contents.⁹

On non-indexicalist approaches, by contrast, objective predicates are unlike pure indexicals in that they *do* have context-invariant contents. In other words, there is nothing in the logical form of those predicates that mandates the saturation of gappy elements by contextual values. Nonetheless, this context-invariant content *may* be (and often is) modulated or enriched by optional pragmatic processes. So a predicate may (but need not) acquire modulated content in context.¹⁰

Another type of non-indexicalist analysis explains the variability in extension in Travis cases not by appealing to the variability in content but by appealing to some form of *circumstance relativity* (see Predelli (2005), MacFarlane (2007, 2009, 2014)). The method is to either introduce additional parameters to a circumstance of evaluation besides the world of evaluation (MacFarlane) or, alternatively, to introduce a function that modifies the standard world parameter (Predelli). In either case, non-indexical contextualism is not supposed to be a relativist position insofar as the values of the parameters at the circumstance of evaluation are fixed by the context of use rather than the context of assessment.¹¹

2.2 The current approach and other approaches

Whilst acknowledging that occasion-sensitivity is a real phenomenon that requires a systematic explanation, the current approach nonetheless departs from existing indexicalist and non-indexicalist analyses of Travis cases. First, the current approach differs from the indexicalist approaches is that it doesn't postulate overt or covert indexical elements in the logical form of objective predicates whose resolution is required for predicates to have interpretable contents. Intuitively, knowing the meaning of *green* seems to be enough to know which property is denoted by *green*, what its content is. By contrast, knowing the 'character' of *I* is not enough to know who *I* refers to (i.e. its content). Furthermore, Kaplan's informal definition of indexicals makes the key reference to a semantic rule associated with their meaning which determines their content

⁹Recanati (2004, 2012) has made a strong case for distinguishing linguistically mandated processes of *saturation*, characteristic of pure indexicals, from essentially optional pragmatic processes of free enrichment and modulation. See also Pagin and Pelletier (2007)

¹⁰See Recanati's and Pagin's pragmatically based *modulation analysis*, Recanati (2004, 2007, 2012); Pagin (2005), Pagin and Pelletier 2007

¹¹See MacFarlane 2014: 88.

in context.¹² There is no comparable rule that governs the interpretation of objective predicates on particular occasions of use.¹³

Second, although managing to circumvent the difficulties that indexicalist analyses face, existing non-indexicalist approaches remain inexplicit about where pragmatic pressures that cause modifications of semantic content exactly originate, and why such modifications would ever take place if these processes are indeed essentially optional.¹⁴ What compels us to shift our understanding of when what we say would be true? Why should a change of occasion compel us to alter our perspective as to when, say, the leaves would count as being green? 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 the contents our sentences semantically express.

We believe that being more specific about the kind of contextual feature to which the extensions of objective predicates are sensitive can help explain *why* their semantic content would *need* to be modified in context even if there is nothing in the logical form of these predicates that requires such modifications. We will 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 not sufficiently informative unless they are modified so that they become relevant for achieving contextually salient goals. So it is informativeness and relevance for goals that *mandate* the modification of semantic content – or so we shall argue.

2.3 Goals as parameter of interpretation

Several different theories have emphasised the significance of practical goals and purposes that guide action in interpreting speech. In the Gricean tradition, these extralinguistic factors are considered as not affecting the truth-conditional content of utterances, rather their role is merely ‘post-semantic’ in that they facilitate the inferential processes that generate conversational implicatures.¹⁵ In the Austinian tradition, however, as emphasised by philosophers and linguists such as Charles Travis, John Searle, Julius Morawschik, Francois Recanati, Paul

¹²“What is common to the words or usages in which I am interested is that the referent is dependent on the context of use and that the meaning of the word provides a rule which determines the referent in terms of certain aspects of the context.” (Kaplan 1989: 490)

¹³Since such a rule is missing in the case of predicates it is hard to see why a predicate rather than any other lexical item that combines with the predicate (e.g. noun) should be treated as an indexical and thus the source of observed variability in truth-valuations. In fact, Vicente (2015) and Giberman (2016) aim to explain Travis cases by locating the source of variability in the relevant noun phrase.

¹⁴Reference 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. See e.g. (Predelli 2005: 365) and (MacFarlane 2007: 264).

¹⁵See also Cappelen and Lepore (2008), Borg (2004), Fodor and Lepore (2004), Fodor (2003).

Pietroski, Robin Carston and many others, extralinguistic contextual factors (that are not linguistically mandated) are important not only for working out implicatures, but also, and primarily so, for fixing the intuitive truth-conditions of utterances. Travis cases are then frequently appealed to as primary evidence for this theory.

Yet, it is in decision-theoretic and game-theoretic approaches to utterance interpretation that goals and purposes have explicitly been given a proper explanatory role through the notions of outcomes and utilities of actions or preferences of agents. Here it is important to mention the appeal to goals and relevance for goals in the semantics of questions (Van Rooy (2003), Ginzburg 1995), and in the game-theoretic analysis of assertions (Parikh 2000, Parikh 2006).¹⁶ In Parikh (2006) different forms of semantic underdetermination (including some Travis cases) have been explicitly considered, and what is contextually relevant for the interpretation of assertive utterances has been linked to the value for a decision problem that the agent faces on some occasion (see Parikh 2006: 362).

Although we will appeal to decision- and game-theoretic notions, we will try and capture these insights in a standard model-theoretic framework using partition semantics (see Groenendijk and Stokhof 1984). The chief reason for this is to highlight the interaction of semantic and pragmatic factors on the interpretation of objective predicates.

3 Occasion-sensitive semantics

3.1 Standing meaning

What does one know when one knows the conventional meaning of an objective predicate and a declarative sentence containing it? What does knowing the meaning of a declarative enable one to do? A classical answer is that to know the meaning of a sentence is to be able to specify its truth conditions, i.e. when it would be true. Accordingly, one ought to know what the world is like that makes the sentence true if one knows that the sentence is true. Travis cases bring this idea into question by showing that knowing meaning and the world of evaluation doesn't guarantee an agreement as to whether the sentence is true. Our challenge is to adequately model the fact that meaning does place some restrictions on which worlds are candidates for making an utterance of a sentence true, but that it does not go so far as to determine, for any given world, whether we would accept it as a truth-maker for an utterance unless more information about context is available.¹⁷ Our approach consists in keeping the semantics of declarative sentences truth-conditional and assigning to pragmatics a job of adequately restricting semantic truth-conditions in order to facilitate a successful fulfilment of practical activities.

¹⁶A similar appeal has been made to so called QUD's – or questions under discussion – within formal pragmatics. See Roberts (1996), Roberts (2004), Roberts (2012), Schoubye and Stokke (2015).

¹⁷We leave it open here how meaning is fixed and whether certain contextual factors (e.g. a history of use) play a role in shaping the conventional meaning of a word.

3.1.1 Meaning as inducing partitions

We assume that a simple declarative¹⁸ sentence containing an objective predicate such as *green* or *desk* has a context-invariant, compositional meaning in virtue of syntactic and lexical structure and linguistic conventions. On our view, standing linguistic meaning has a substantial role to play in the process of utterance interpretation insofar as a particular sentence, meaning what it does, cannot be used to state just anything.¹⁹ However, in the light of the occasion-sensitivity of objective predicates, we believe it is crucial to capture the idea that this meaning only *constrains* the intended interpretation or the proposition(s) expressed by an utterance on some occasion without fully determining it.

Consider the sentence (1), *The leaves are green*. Just by knowing what this sentence means we are able to exclude many worlds in which the leaves are not green. Still, this sentence, meaning what it does, is compatible with a variety of alternative worlds: worlds in which the leaves are naturally green, worlds in which they are painted green, worlds in which they are lit by a green light source, worlds in which they are mouldy green, worlds in which they are dyed green, etcetera. There are also, of course, worlds at which various combinations of these ways of being green obtain of the leaves (e.g. worlds in which the leaves are both lit green and mouldy). So knowing what the sentence means, albeit excluding many possibilities, nonetheless leaves it unresolved what the actual state of the world is like, that is, in which particular way the leaves are green.

The meaning of a declarative sentence ϕ thus induces a *bi-partition* of logical space: one cell consisting of the worlds where the sentence is true and another of its complement.²⁰ Let us call the cell of the bi-partition where ϕ is true, *the truth-maker set for ϕ* or $|\phi|$ for short. Anyone who knows the meaning of ϕ thus expects the actual world to be somewhere in $|\phi|$, if ϕ is indeed true.

However, as intimated earlier, there are many different ways in which a given property can hold of a given object. For example, the leaves may be green whilst (only) painted green or whilst (only) naturally green or whilst both. To capture the fact that the sentence's truth-makers don't form a homogeneous class, and that one's knowledge of meaning is not sufficient to determine in *which way* some object has a given property, we suggest that, in addition to the initial bi-partition of logical space, there is a further fragmentation or *refinement* of the bi-partition induced by different ways in which a given property can hold of a given object.²¹ If, for instance, we assume there are two ways that greenness

¹⁸The model we are about to propose can be extended to interrogatives and imperatives, although this is beyond the scope of this paper.

¹⁹Thus we cannot, without enough stage-setting, use the sentence *The cat is on the mat* to say that snow is white. See Travis 2009, Recanati 2004.

²⁰To some extent this is in line with the view known as *semantic minimalism* (see Cappelen and Lepore 2008, Borg 2004, Borg 2012) As Emma Borg notes, “[the minimalist] holds that word meanings generalise over specific contextual understandings of them – that the meaning of ‘green’ leaves open the way in which an object might be green (on its surface, on the inside, etc.)” (Borg 2012: 46).

²¹The number of ways in which a property can be applied to some individual will be constrained by the cardinality of individuals, but in principle for any particular d denoted by β

may apply to the leaves, then the refinement of the bi-partition will contain three cells such that their union forms a *cover* over $|\phi|$, and one cell covering the complement of $|\phi|$.²² We illustrate the initial partition and its refinement in **Figure 1** and give a formal definition in section 4. The diagram in **Figure 1** represents the set of all alternatives, assuming only two ways for leaves to be green: painted and naturally green. The cell w_{11} consists of the worlds in which the leaves are both painted and naturally green, the cell w_{10} of the worlds in which they are only painted, and w_{01} contains the worlds in which the leaves are only naturally green. Since painted and naturally green exhaust the ways of being green, the w_{00} cell consists of worlds in which the leaves are *not* green.

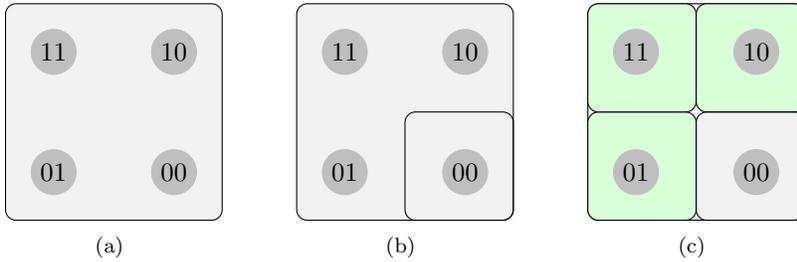


Figure 1: Logical space, bi-partition and its refinement

3.1.2 Partitions and occasion-sensitivity

We defined the meaning of a declarative sentence classically as the set of worlds in which the sentence is true. Crucially, however, we suggested that this is not a homogeneous set of truth-makers but it admits of further fragmentation in accordance with different ways that a given property can hold of a given object. This way of thinking of the truth-maker set makes us better appreciate the fact that one’s knowledge of meaning is not sufficient to determine in which *particular* refinement cell the actual world is located.

Travis cases show that not all the worlds where a declarative sentence is true will make its utterance on some occasion true. In general, due to pragmatic pressures, the set of truth-makers for an utterance on an occasion will include only a certain subset of refinement cells in at which the sentence is true. For instance, in the botanist context (see **Figure 2**) the set of worlds in which the leaves are not naturally green, w_{10} , is excluded from the set of truth-makers

and any particular property expressed by α , the number of distinct ways that the property applies to the individual is determined by what rational thinkers are prepared to recognise as a distinct way. Any way for things to be, regardless of how specific and complex it is, may have its generation of ways. Namely, *naturally green* can branch off into *naturally green*₁ and *naturally green*₂, and *naturally green*₁ can branch off into *naturally green*_{1(1)}} and *naturally green*_{1(2)}} and so on.

²²See **Definition 4**

for the utterance of (1) because it matters in this context that the leaves are naturally green (also w_{00} is excluded purely in virtue of meaning). By contrast, in the decorator’s context it’s required that they are painted green, so the cell ruled out is w_{010} (and of course w_{00}). In each context, $|(1)|$ is restricted to allow only certain pragmatically accepted refinement cells, whilst pragmatically “bad” cells are *excluded*. By asserting (1) our botanist selects one subset of the refinement cells in $|\phi|$ (and has one understanding of when an utterance of (1) would be true), and the decorator selects another subset of the refinement cells in $|\phi|$ (and has another understanding of when an utterance of (1) would be true).

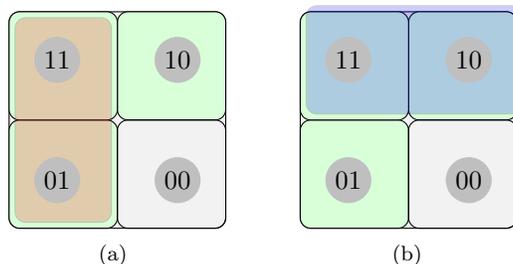


Figure 2: Botanist and decorator contexts

In the next section we look more closely at a pragmatic mechanism behind the exclusion of certain refinement cells which explains why speakers find some sets of worlds unacceptable as truth-makers for an utterance on some occasion.

3.2 Pragmatics

3.2.1 Goal-conduciveness

Our key observation from the previous section is that the set of worlds in which a sentence is true is a *very* large, highly diverse (in many cases *too* diverse) set of worlds. In this section we want to elaborate on the fact that not all truth-makers for ϕ are accepted by rational agents as truth-makers of some utterance of ϕ on an occasion. We suggest that this is because not all of them can successfully support some particular course of action, and so won’t be equally valuable when planning to achieve a certain practical goal. For a particular practical goal γ and world w we can take it that w will be more or less *conducive* to γ (i.e. achieving γ in w is more or less difficult).

Consider again a case where a botanist is preparing to do some experiment on chloroplasts. To achieve this goal she needs to find the leaves that are green in a suitable way, where the leaves that are only painted won’t do. The botanist knows she will not be able to go on with her experiment if the leaves were that way. Thus this particular way for the leaves to be green is excluded from the

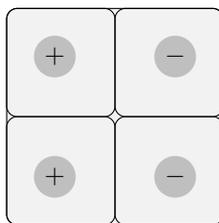


Figure 3: Conduciveness map for the botanist

botanist’s understanding of being green on this occasion.²³ Having a particular practical goal, therefore, requires that the way for a given object to have a given property is *goal-conducive*, and it requires that an agent knows this fact.

We will use the formal notion of *information states* to capture what some agent knows or believes. Let s be an *information state* modelled as a set of possible worlds which are compatible with the information available in the state. An information state of any semantically competent agent a who believes that the sentence ϕ is true (minimally) contains the truth-maker set for ϕ . For example, anyone who knows that the leaves are green knows that they are green in no particular way. Let us call the information state of an agent who would locate the actual world in *any* refinement cell of the truth-maker set, the *initial information state*.

Now, besides having their basic semantic competence *pragmatically competent* agents will also be able to assess a given situation from the point of view of a particular goal and decide whether the situation would facilitate a successful achievement of it or not. The information state of a pragmatically competent agent is *enhanced* compared to the initial information state in as much as she knows what it takes to achieve a given goal and is able to disqualify certain sets of worlds as not conducive to this goal.²⁴ Although this is a non-trivial simplification, we take it that possible worlds, considered as formal objects, differ only with respect to the way in which a given property applies to a given object: all else is equal, irrelevant or absent. It is all we therefore have to go on in deciding whether any given set of worlds is goal-conducive or not.²⁵

²³However, not only *her* understanding. It is an understanding that is required by this particular goal, project, activity, so it would be the understanding of anyone who pursues this or similar goal. See section 5.1.

²⁴For the simplicity’s sake, here we consider goal-conduciveness to be an absolute value of a world. So, in our framework, a world can either be goal-conducive or not. However, there are cases where goal-conduciveness is a matter of degree, and worlds are more or less goal-conducive. In those cases, pragmatically competent speakers are only required to be able to judge relative differences between the worlds *vis-a-vis* their capacity to support a goal.

²⁵A world in which the leaves are naturally green *and* experiments on chloroplasts are, say, banned would not be as, a whole, conducive to Sid’s goal of doing such an experiment notwithstanding the goal-conduciveness of the leaves and their greenness. Because potentially any contingency could have an impact on whether a goal will be successfully achieved we have to take into account only those facts about goal-conduciveness that are relevant for understanding a given expression.

For any practical goal γ there will be a *valuation map* over the cells of the refinement of the bi-partition, assigning each cell a *goal-conduciveness value* in accordance with the capacity of the worlds in the cell to support a successful achievement of a given goal.²⁶ We suggest that any pragmatically competent agent will have a particular goal-conduciveness valuation map as part of her information state. Which particular map is relevant is determined by which goal is salient on an occasion. **Figure 3** illustrates the goal-conduciveness valuation map for the botanist’s goal. Although there may be some variations among goal-conduciveness valuation maps between agents, in many cases, pragmatically competent agents will have similar or overlapping valuations as to which sets of worlds are conducive to a particular goal. As we shall see in section 5.1, a prerequisite for a *genuine disagreement* is that both agents have access to similar or overlapping valuations of worlds, otherwise they will not express the same proposition (i.e. they will communicate cross-purposes).²⁷

3.2.2 Communication

We suggested that a particular way the world is can be more or less conducive to some practical goal, and that worlds can be evaluated based on the extent to which they are conducive to a given goal. It also became clearer that not all the worlds at which a sentence is true are conducive to a certain goal. Thinking of possible worlds in terms of goal-conduciveness thus gives rise to a particular view of communication where agents communicate to seek and exchange the information about a world that helps them achieve their goals.

To take a Travis case again as our example, if I know that my friend is a botanist and is searching for some green leaves for her experiment (assuming I am not totally ignorant about this topic) then I should also know that her request for green leaves bears on that particular goal and that so should my response. If I then sincerely say to my friend, “Search no more, the leaves on my tree are green”, what I say is that the world is indeed such that her project can be successfully carried out (at least as far as the greenness of the leaves is concerned). In other words, by asserting this, I suggested that the way the world is *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).

On the current approach, not every utterance of a declarative sentence will be sufficiently informative. Only if an agent is in a pragmatically enhanced information state is a given assertion sufficiently informative, where such an enhancement presupposes that the agent has access to a relevant conduciveness valuation map. The goal will then determine the extent to which an assertion

²⁶See Definition 8

²⁷For discussion of the notion of disagreement by talking cross-purposes, see Davies (2017).

is informative.²⁸ To illustrate consider the following polar question in (11). Let us assume that A, B, C and D are semantically competent and know that there are different ways for the leaves to be green. Furthermore, they all know the contextually salient goal and have access to the goal-conduciveness valuation map.

- (11) A: Are the leaves green _{γ} ?
- (a) B: The leaves are green or the leaves are not green.
 - (b) C: The leaves are green (in some way or other).
 - (c) D: The leaves are green _{γ} .

Assuming A is a pragmatically competent agent, we know that question (11) relates to a particular goal. The informativeness of an answer increases as more alternatives are excluded. (11a) is not informative since it asserts something that any rational, linguistically competent speaker should already know. (11b) excludes *some* possibilities, and to this extent it is not entirely trivial and uninformative. Yet it is *not* sufficiently (pragmatically) informative to facilitate a particular course of action or support A’s goal because (11b) is compatible with *any* cell in the truth-maker set. Those who seek an information do that with a particular end in sight. (11b), however, does not resolve a particular goal-related issue; it does not contain an information based on which our agent A can act. The only informative answer that resolves the issue raised by the question in (11) is (11c) because A is informed that the actual world is compatible with the enhanced information state that A is in.

Let us now apply this to *Occasion 2* in **The leaves are green** Travis case. We suppose both speakers are pragmatically competent agents. The asymmetry between them consists in the following: by asking whether the leaves are green, *Zoe entertains an issue*²⁹ over the bi-partition of goal-conducive cells of worlds $\{w_{01}, w_{11}\}$ and their complement (namely, *Zoe* doesn’t know whether the actual world is in $\{w_{01}, w_{11}\}$ or in its complement), whilst her friend *Pia* knows (or ought to know) that the actual world (in the example, non-naturally, painted green leaves) is *not* in $\{w_{01}, w_{11}\}$. Therefore, on the proposed goal-directed model of communication, *Pia* gives *false* information about the world by asserting (1) on this occasion and is *not* resolving *Zoe*’s issue.

In the following section we define a formal model of occasion-sensitive semantics for a small fragment of English.

²⁸This is in line with some decision-theoretic approaches to questions where whether an answer resolves a question or not is determined by a salient decision problem. See Van Rooy (2003).

²⁹The notion of issues that an agent may entertain and what resolves them has been formalised in the framework of inquisitive semantics. See Groenendijk et al. 2009, Ciardelli et al. 2013, Ciardelli et al. 2015, Ciardelli and Roelofsen 2015)

4 A model

In order to determine when an utterance of the sentence on some occasion is true we propose to first evaluate all cells of the refined partition of worlds in terms of how goal-conducive they are and to assign each a value, thus producing a valuation map over those cells. A set of most goal-conducive worlds that are also semantically compatible with a sentence will correspond to the proposition expressed by an utterance of that sentence by an agent on an occasion.

4.1 Language

A small language fragment for occasion sensitive semantics \mathcal{L}_{os}

4.1.1 Basic expressions

- (a) Individual constants: **Sid, the leaves, the shoes**
- (b) Predicates: **green, desk, under the bed**
- (c) Logical constants: $\neg, \wedge, \vee, \rightarrow$

4.1.2 Formation rules

- (a) if α is a predicate and β is an individual constant, then $\alpha(\beta)$ is a sentence
- (b) if ϕ is a sentence, then $\neg\phi$ is a sentence
- (c) if ϕ is a sentence, then $\phi \wedge \phi, \phi \vee \phi, \phi \rightarrow \phi$ is a sentence

4.2 Occasion-sensitive semantics for \mathcal{L}_{os}

A model M for \mathcal{L}_{os} is a tuple $\langle D, W, I, \Sigma, \Gamma, \Theta \rangle$ where

- i. D is a set of *individuals*
- ii. W is a set of *possible worlds*
- iii. I is the *interpretation function*
- iv. Σ is the *ordering function*
- v. Γ is a set of *goals*
- vi. Θ is the *goal-conduciveness function*

4.2.1 Interpretation function

I is the interpretation function of the model M assigning an intension I_α to each predicate and I_β to each individual constant as follows:

Definition 1 (Interpretation function I_α)

If α is a predicate then I_α is a function such that for each $w \in W : I_\alpha(w) \subseteq D$

Definition 2 (Interpretation function I_β)

If β is an individual constant then I_β is a function such that for each $w \in W :$
 $I_\beta(w) \in D$

4.2.2 Truth

We write $M, w \models \phi$ for a sentence ϕ is true in model M with respect to the world of evaluation w

Definition 3 (Sentential truth $I_{\alpha(\beta)}$)³⁰

$M, w \models \phi$ iff $\forall w \in W : I_\beta(w) \in I_\alpha(w)$

4.2.3 Ordering function

Definition 4 (Ordering of $I_\alpha(w)$)

Let Σ_α be a function from W to n -tuples of sets of individuals, such that for any world w ,

- $\Sigma_\alpha(w) = \langle \sigma_\alpha^1(w), \dots, \sigma_\alpha^n(w) \rangle$, where $\sigma_\alpha^i(w) \subseteq I_\alpha(w)$ and $i \in \{1 \dots n\}$
- $\Sigma_\alpha(w)$ is a *cover* of $I_\alpha(w)$ such that the union of the sets in $\Sigma_\alpha(w)$ is equal to $I_\alpha(w)$, i.e. $\bigcup \sigma_\alpha^i(w) = I_\alpha(w)$

Example 1: Let us assume that $W : \{w_1, w_2, w_3, w_4, w_5\}$, and $D : \{a, b, c, d\}$. For each world w , let us further assume that the interpretation function I_α will assign the following values to the predicate α :

$$I_\alpha(w_1) = \{a, b, c\}, I_\alpha(w_2) = \{a\}, I_\alpha(w_3) = \{a, b\}, I_\alpha(w_4) = \{a, b, c\}, I_\alpha(w_5) = \emptyset$$

The ordering function Σ_α orders individuals in the extension of I_α according to different ways that individuals may have a property denoted by α . For a property denoted by α there will be n number of ways σ_α in which that property can be applied to individuals. To distinguish between the ways we assign each an index i . Thus $\Sigma_\alpha = \{\sigma_\alpha^1, \dots, \sigma_\alpha^i, \dots, \sigma_\alpha^n\}$, where $i \in \{1 \leq i \leq n\}$. The application of Σ_α to W , where $n = 2$ returns the following values:

³⁰Definitions for all complex clauses are classical.

$$\Sigma_\alpha(w_1) = \{\sigma_\alpha^1(w_1), \sigma_\alpha^2(w_1)\} = \langle \{a, c\}, \{a, b, c\} \rangle$$

$$\Sigma_\alpha(w_2) = \{\sigma_\alpha^1(w_2), \sigma_\alpha^2(w_2)\} = \langle \{a\}, \{\emptyset\} \rangle$$

$$\Sigma_\alpha(w_3) = \{\sigma_\alpha^1(w_3), \sigma_\alpha^2(w_3)\} = \langle \{a, b, c\}, \{b\} \rangle$$

$$\Sigma_\alpha(w_4) = \{\sigma_\alpha^1(w_4), \sigma_\alpha^2(w_4)\} = \langle \{b\}, \{a, c\} \rangle$$

$$\Sigma_\alpha(w_5) = \{\sigma_\alpha^1(w_5), \sigma_\alpha^2(w_5)\} = \{\emptyset\}$$

4.2.4 Equivalence relations

Definition 5 (Equivalence relations between worlds wrt $\alpha(\beta)$)

$$wR_{\alpha(\beta)}w' \text{ iff } I_\beta(w) \in I_\alpha(w) \Leftrightarrow I_\beta(w') \in I_\alpha(w')$$

Let \mathcal{W} be a bi-partition of W , and let $c \in \mathcal{W}$ be a cell of the bi-partition. Any two worlds are in the same cell c of the bi-partition \mathcal{W} iff $wR_{\alpha(\beta)}w'$.

We can define a refinement \mathcal{Z} of the bi-partition \mathcal{W} with respect to Σ_α by first defining equivalence relations with respect to the sentence $\alpha(\beta)$ and the ways Σ_α in which α applies to β .

Definition 6 (Equivalence relations wrt $\alpha(\beta)$ and Σ_α)

$$wR_{\Sigma_\alpha}^{\alpha(\beta)}w' \text{ iff } \forall i : 1 \leq i \leq n : I_\beta(w) \in \sigma_\alpha^i(w) \Leftrightarrow I_\beta(w') \in \sigma_\alpha^i(w')$$

$R_{\Sigma_\alpha}^{\alpha(\beta)}$ is an equivalence relation between worlds with respect to the set of ways Σ_α in which a property denoted by α applies to the individual denoted by β . It states that any two worlds are equivalent relative to a given sentence $\alpha(\beta)$ and a given way σ_α^i only if the individual denoted by β has a property denoted by α in the same way σ_α^i (or is in the extension of σ_α^i) at both worlds.

Definition 7 (Refinement \mathcal{Z} of the bi-partition \mathcal{W} of W)

Any two worlds w, w' are in the same cell of the partition \mathcal{Z} , where $\mathcal{Z} \leq \mathcal{W}$ (\mathcal{Z} is finer than \mathcal{W}), iff $wR_{\Sigma_\alpha}^{\alpha(\beta)}w'$.

Example 2: Suppose that $I_\beta(w) = a$ at all the worlds in W and $\alpha = \mathbf{green}$. It follows that:

$$I_{\mathbf{green}(a)}(w_1) = 1$$

$$I_{\mathbf{green}(a)}(w_2) = 1$$

$$I_{\mathbf{green}(a)}(w_3) = 1$$

$$I_{\mathbf{green}(a)}(w_4) = 1$$

$$I_{\mathbf{green}(a)}(w_5) = 0$$

Table 1: Partition \mathcal{Z} of W wrt $R_{\Sigma_{\mathbf{green}}}^{\mathbf{green}(a)}$

$c_1: \sigma_{\mathbf{green}(a)}^1$	$c_2: \sigma_{\mathbf{green}(a)}^2$	$c_3: \mathbf{Both}$	$c_4: \mathbf{None}$
w_2	w_4	w_1	w_5
w_3			

The worlds that are in the same cell of the bi-partition \mathcal{W} are: w_1, w_2, w_3 and w_4 . Namely, supposing $[w]_{R^{\mathbf{green}(a)}}$ is an equivalence class generated by $R^{\mathbf{green}(a)}$, then \mathcal{W} consists of two equivalence classes (w_1 and w_5 are class representatives):

$$[w_1]_{R^{\mathbf{green}(a)}} = \{w_1, w_2, w_3, w_4\}$$

$$[w_5]_{R^{\mathbf{green}(a)}} = \{w_5\}$$

Furthermore, suppose that there are two different ways in which **green** applies to a , i.e. $\Sigma_{\mathbf{green}} = \{\sigma_{\mathbf{green}}^1, \sigma_{\mathbf{green}}^2\}$. Then there will be a further fragmentation of the bi-partition \mathcal{W} wrt to the set of ways $\sigma_{\mathbf{green}}^i \in \Sigma_{\mathbf{green}}$ in which the property denoted by **green** applies to a . Assuming $\Sigma_{\mathbf{green}}(w)$ returns same values as in *Example 1*, then we have a further partition of worlds wrt to the equivalence relation $R_{\Sigma_{\mathbf{green}}}^{\mathbf{green}(a)}$ as illustrated in **Table 1**.

4.2.5 Goal-conduciveness

In our framework, given a particular goal $\gamma \in \Gamma$, an agent will select only some of the cells of the refinement \mathcal{Z} of the bi-partition \mathcal{W} as conducive to that goal.

Definition 8 (Goal-conduciveness function)

Let Θ be a goal-conduciveness function from a partition $P_{\alpha(\beta)}$ to assignments of value $(+, -)$ to each cell c in the partition.

$\Theta : \gamma(P_{\alpha(\beta)}) : \forall c \in P_{\alpha(\beta)}, \Theta_{\gamma}(c) \in \{+, -\}$, where γ is a particular contextually salient goal.

For all the worlds in the complement of $|\alpha(\beta)|$, the cell corresponding to $|\alpha(\beta)|^C$ is assigned a *negative* value since the sentence is false at those worlds.

Example 3: Suppose that a (refined) partition $\mathcal{Z} = \{c_1, c_2, c_3, c_4\}$ (see **Table 1**). Each cell corresponds to an equivalence class $[w]$ wrt $R_{\Sigma_{\mathbf{green}}}^{\mathbf{green}(a)}$ such that:

$$c_1 : [w_2]_{R_{\Sigma_{\mathbf{green}}}^{\mathbf{green}(a)}} = \{w_2, w_3\}$$

$$c_2 : [w_4]_{R_{\Sigma_{\mathbf{green}}}^{\mathbf{green}(a)}} = \{w_4\}$$

$$c_3 : [w_1]_{R_{\Sigma_{\mathbf{green}}}^{\mathbf{green}(a)}} = \{w_1\}$$

$$c_4 : [w_5]_{R_{\Sigma^{\mathbf{green}(a)}}} = \{w_5\}$$

Furthermore, suppose that for all w in $\{c_1, c_2, c_3\}$, $I_a(w) \in I_{\mathbf{green}(w)}$ and for all w in c_4 , $I_a(w) \notin I_{\mathbf{green}(w)}$. The goal-conduciveness function will assign each cell in \mathcal{Z} a value $(+, -)$ in accordance with the capacity of worlds in the cell to support a particular goal γ (where the cell corresponding to the complement of $|\mathbf{green}(a)|$ is assigned a negative value by definition). Let us assume Θ_γ returns the following values:

$$\Theta_\gamma(c_1) = -$$

$$\Theta_\gamma(c_2) = +$$

$$\Theta_\gamma(c_3) = +$$

$$\Theta_\gamma(c_4) = -$$

This valuation map for goal γ will be important for determining the truth value of an utterance of $\mathbf{green}(a)$ at w given a contextually salient goal γ .

4.2.6 Contextual truth

We write $M, w \models_\gamma \alpha(\beta)$ for an utterance of $\alpha(\beta)$ is true in model M with respect to the world of evaluation w and the goal of context γ

Definition 9 (Contextual truth)

$M, w \models_\gamma \alpha(\beta)$ iff

- i. $I_\beta(w) \in I_\alpha(w)$
- ii. $\Theta_\gamma([w]_{R_{\Sigma^\alpha}^{\alpha(\beta)}}) = +$

$M, w \models_\gamma \neg\phi$ iff $M, w \not\models_\gamma \phi$

$M, w \models_\gamma (\phi \wedge \psi)$ iff $M, w \models_\gamma \phi$ and $M, w \models_\gamma \psi$

$M, w \models_\gamma (\phi \vee \psi)$ iff $M, w \models_\gamma \phi$ or $M, w \models_\gamma \psi$

This definition states that an utterance of the sentence $\alpha(\beta)$ relative to the goal of context γ is true at a world w only if the sentence is true at w and w is in the equivalence class $[w]$ generated by $R_{\Sigma^\alpha}^{\alpha(\beta)}$ to which Θ_γ assigns value $+$, namely if the world is goal-conducive (with respect to the sentence).

Example 5: Let us assume that the actual world is w_2 and that the sentence is $\mathbf{green}(a)$ (i.e. $\beta = a$ and $\alpha = \mathbf{green}$ as in previous examples). Then, since w_2 is in the cell c_1 (see **Table 1**) and $\Theta_\gamma(c_1) = -$, $M, w \not\models_\gamma \mathbf{green}(a)$, i.e. an utterance of $\mathbf{green}(a)$ in model M with respect to the goal of context γ is false at w_2 .

5 Applications

5.1 Disagreement: discretionary and objective predicates

We started this discussion pointing out the intuitive difference between objective and discretionary predicates concerning factors to which their truth assessments are sensitive. There we noted that factual statements containing objective predicates, unlike statements of opinion containing discretionary predicates, are true at the world *simpliciter* without further relativisation to individual perspectives etc. However, in an attempt to account for Travis cases, we have shown that the evaluation of factual statements *is* also relative to further factors besides the world of assessment, namely, to contextually salient goals. So the question is whether our model preserves the intuitive distinction between objective and discretionary predicates (including the idea of genuine, factual disagreement) or not. To answer this question, let us first see what our model predicts about agreement and disagreement. Note that it is normally assumed that for genuine disagreement the *same proposition* needs to be asserted and denied (i.e. that two individuals genuinely *contradict* one another). A disagreement is *faultless* if two individuals are *genuinely* contradicting one another but neither is at fault (see Kölbel 2004).

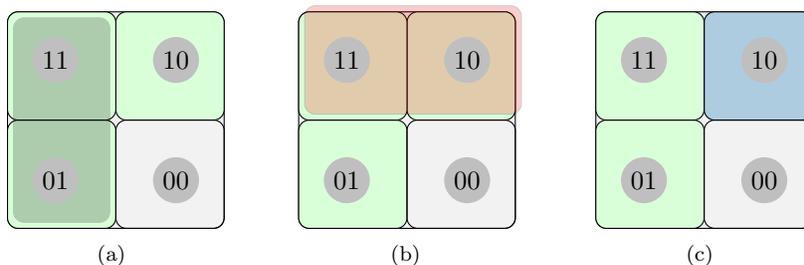


Figure 4: Utterance of (1) in three different contexts

In **Figure 4**, we have only two different ways of being green: painted green and natural green, and the worlds where the leaves are only naturally green, only painted green, both or none. We also have three different understandings of when an utterance of (1) would be true: the botanist’s (*dark green*), the decorator’s (*red*), and the painter’s (*blue*). At w_{00} no utterance will be true (because the sentence is false). All semantically competent agents are assumed to be in the initial information state, in which case they all *verbally agree* on the meaning of a sentence (*viz.* they agree on which worlds constitute the truth-maker set for the sentence).

The botanist and the decorator agree that the utterance of (1) is true at w_{11} , but they disagree about the truth-value of the utterance at w_{10} and w_{01} . Notice, however, that given that the botanist and the decorator don’t share

the same understanding of *green* we only have an *appearance* of agreement (or disagreement). In other words, when communicating, they talk *cross-purposes* and *fail to* genuinely agree or disagree. Call this variety of agreement and disagreement *pseudo-agreement* and *pseudo-disagreement*. The fact that there is an appearance of agreement between agents at a world about the truth-value of an utterance does not entail that agents entertain same propositions: i.e. although the botanist and the decorator pseudo-agree that (1) is true at w_{11} , they nevertheless express *different* propositions by their utterances, i.e. $(1)_\gamma$ and $(1)_{\gamma'}$. So an appearance of agreement between the agents about the truth value of an utterance of ϕ at a particular world is not enough to establish that they also share the same understanding of when an utterance of ϕ would be true. Finally, the botanist's and the painter's propositions never intersect: there will be no agreement (pseudo or genuine) between the agents who occupy these two understandings of when an utterance of (1) would be true, although they still verbally agree on the meaning of *green*.

So far we explained what verbal and pseudo agreement and disagreement consist in. When, according to our approach, would two agents *genuinely agree* or *disagree*? Consider again the green leaves dialogue concerning the leaves which are painted green but naturally red, i.e at the world w_{10} . Let us compare two versions of this dialogue, in (12*) and (12).

(12*) The green leaves dialogue

- (a) The decorator (relative to the decorator's goal): The leaves are green.
- (b) The botanist (relative to the botanist's goal): No, the leaves are not green.

In (12*) the decorator assesses the leaves from her context, and the botanist assesses them from his. As in the chilli dialogue in (5) with the predicate *tasty*, we seem to have a case of assessment sensitivity where the truth-value of utterances depends on which point of evaluation is occupied. Accordingly, both utterances in (12*) seem to be true relative to a particular perspective, and no one seems to be at fault. Still, there is a crucial difference between (12*) and the chilli dialogue (at least on relativist analyses of predicates of personal taste), in that in (12*) the decorator and the botanist are asserting and denying *different* propositions, $(1)_\gamma$ and $\neg(1)_{\gamma'}$ although the sentences they utter have the same standing meaning. That is, in (12*) we have a case of *pseudo-disagreement* where the agents are talking *cross-purposes*, although they *verbally agree*. Nonetheless, the conditions for a genuine agreement or disagreement are not in place.

Compare to this another version of the green leaves dialogue, which closely resembles **The leaves are green** Travis case.³¹

(12) The green leaves dialogue

³¹See footnote 32 for why not entirely.

- (a) The decorator (relative to the botanist’s goal): The leaves are green.
- (b) The botanist (relative to the botanist’s goal): No, the leaves are not green.

The difference between the dialogues in (12*) and (12) is accounted for by the difference in accessibility to *relevant goal-conduciveness valuation maps*. In (12*) we assume that either (i) the decorator’s information state does not include the conduciveness valuation map for the botanist’s goal γ' , or (ii) she’s not aware of the botanist’s goal and so is unable to evaluate goal-conduciveness of the world, or (iii) she might be mistaken about what is conducive to the botanist’s goal; in any case, she is unable to share the botanist’s understanding of when an utterance of (1) would be true, and so they communicate cross-purposes. However, assuming the decorator knows enough about the botanist’s goal and what’s required to reach it (*viz.* has access to the correct conduciveness valuation map for the botanist’s goal), she *ought to* recognise (if rational) that an utterance (12a) is false at w_{10} . Given that the decorator in (12) *could* and *should* evaluate the world (for goal-conduciveness) from the point of view of the botanist’s goal, she is *at fault* as regards the truth-valuation of the utterance. Thus here we have a case of *genuine, factual* disagreement because, with respect to w_{10} , one party (the decorator) asserts that the leaves are green in a way that is suitable for the botanist’s purpose, i.e. $(1)_{\gamma'}$, and the other party (the botanist) claims that the leaves are not green in a way that is suitable for her purpose, $\neg(1)_{\gamma'}$.³² **Table 2.** summarises the three varieties of disagreement for sentences containing objective predicates.

Let us now see how this classification compares to our intuitions about matters of taste. First, the possibility to occupy a different point of evaluation (and assess how things are from within there) seems *not* to be available for matters of taste. When John and Sid disagree whether the chilli is tasty, best they could do is acknowledge that they are evaluating the world differently; still no matter what they do (how much information about the world or each other they obtain) John *cannot* put himself in Sid’s perspective (and vice versa) and acknowledge that (and in which way) John’s standard of taste is wrong (false) and Sid’s is correct. Each of them has ‘epistemic privilege’ with regards to what they find tasty.³³ Thus, assuming John and Sid are competent speakers who know the standing meanings of *chilli* and *tasty*, if they do disagree genuinely, they disagree *faultlessly*.³⁴ Still they *cannot but* disagree faultlessly about matters of taste. Contrast this to matters of fact where there is always a possibility

³²**The leaves are green** Travis case is not that clear cut as the case in (12), since the decorator, despite (by assumption) being informationally and rationally equipped (and so required) to apply the experiment-relevant valuation map, seems to be biased towards her initial valuation map regardless of a new goal being introduced. So she is not appreciating that the occasion has changed (she’s stuck on her default occasion).

³³As Lasersohn points out, “we each have a privileged perspective on our internal affective states” (Lasersohn 2005: 655).

³⁴This point holds regardless of whether the account of personal taste predicates is contextualist or relativist, i.e. whether the judge is a parameter of content or a circumstance of evaluation.

Table 2: Varieties of disagreement: objective predicates

Disagreement	Shared meaning	Shared proposition	Faultless
verbal	no	–	–
pseudo	yes	no	–
genuine	yes	yes	no

for rational agents to agree or disagree *genuinely* (and with a possibility of one being at *fault*) once they acquire enough relevant information about a context (and what it takes to achieve some contextually salient goal).

Second, albeit relative to judges or personal standards, the evaluation of discretionary predicates seems not be relative to further contextual parameters such as practical goals. That is, it seems correct that if John evaluates (at a world) the chilli as tasty, then the chilli is tasty *simpliciter* regardless of what John’s further goals are in that context, e.g. whether the chilli is to be served at the Queen’s banquet or is prepared for a Sunday dinner.³⁵ Thus, even though the interpretation and valuation of sentences containing discretionary predicates is relative to judges, perspectives, personal standards (or what have you), these predicates don’t seem to exhibit a similar flexibility that objective predicates have, namely, to be interpreted or evaluated differently depending on which goal is salient in a context.³⁶

So, although our model relativises the evaluation of objective predicates to different perspectives (characterised by different contextually salient goals), and thus seems to temporarily blur the intuitive contrast between objective and discretionary predicates, the contrast is still there insofar as (i) objective perspectives are *accessible* to all rational, well-informed agents and subjective perspectives aren’t, and (ii) unlike discretionary predicates, the evaluation of objective predicates exhibits goal-relativity. Hence utterance truth, albeit relative to contextually salient goals, is still an objective (rather than a relativist) notion.

5.2 Cancellability

Cancellability is a feature of utterances usually associated with conversational implicatures where the intended content to be communicated by an utterance is something that is implied by uttering a sentence in a given context rather than something that the sentence would semantically express (meaning what it does). By saying that I am French I (may) imply in a context that I am a good cook (and a number of other things).³⁷ In the following example, (13b)

³⁵John could disqualify the chilli from being served at the banquet because of other consideration (e.g. its presentation) but not its taste if it is indeed tasty.

³⁶This assuming that gustatory and other personal experiences have a qualia-like nature.

³⁷See Recanati 2004.

con conversationally implicates (13c), which here spells out the content that the speaker intends to communicate in context.

(13a) Can you cook?

(13b) I am French.

(13c) \rightsquigarrow *I am a good cook*

Any implicated content such as the one in (13b) may be explicitly *cancelled* if the speaker wants to make sure that the hearer doesn't interpret (13b) as (13c).

(14) I am French but I am not a good cook.

Whilst expected implicatures may be cancelled in this way, we cannot cancel what the sentence semantically expresses without a contradiction as illustrated in (15).

(15) I am French but I am not French.

Even though we think it would be wrong to assimilate the phenomenon of occasion-sensitivity to that of conversational implicature, it is still important to note that in Travis cases the contextually expected interpretation of an utterance *is* cancellable without a contradiction. In (16) we assume the original set-up where the salient leaves are painted green and naturally red.

(16a) Botanist: Do you have green leaves?

(16b) Pia: These leaves are green, but they are not naturally green.

Unlike in the above mentioned Travis case where Pia's response (*viz.* "These leaves are green") is intuitively false (and we offered an explanation of why this is so), once the contextually expected material is made explicit then Pia's response seems *felicitous*: namely, what she says in (16b) is true. How does our theory account for this intuition?

On the approach defended in this paper, felicity of Pia's response in (16b) is to be explained in the following way. Pia is clearly aware of the botanist's goal and what it would take to achieve it. This is manifested by the fact that she is able to gesture towards the contextually expected reading of the expression *green leaves* by using a more specific form, namely, *naturally green leaves*. That is, Pia has access to a conduciveness valuation map that is adequate for this goal. However, instead of simply denying that the leaves are green in the way that would be relevant for the botanist's goal (saying that the leaves are not green), Pia also asserts that they are green in some other way where this is *not* the way that her botanist friend would expect the world to be given her pragmatically enhanced state and her goal. She is explicitly cancelling the expected interpretation of the sentence *These leaves are green* yet also saying that the leaves count as being green in some other way, for some other purpose.

The remark she is making is thus *metalinguistic* in as much as she is drawing the botanist’s attention to another proposition expressed by *The leaves are green*, which is not the one that the botanist would expect, and she is suggesting that that proposition would be true were they on another occasion.

Is Pia’s answer in (16b) informative according to our account? To see that it is compare the dialogue in (16) to the dialogue in (17).

(17a) Botanist: Do you have green leaves?

(17b) Pia: These leaves are green in some way or other

As pointed out in section 3.2.2 whether an assertion is sufficiently informative depends on whether the set of worlds at which a sentence is true is adequately restricted to only those sets of worlds that are conducive to the contextually salient goal. (17b) is *not* sufficiently informative since no refinement cells from the truth-maker set are excluded and so the response does not enable any further action by the botanist. By contrast, (16b) is *too* informative since Pia is saying that it’s not the case that the leaves are naturally green (so that the world is not conducive to the contextually salient goal) but *also* that the leaves are green in some way which is not relevant to the botanist. In (16b) Pia is thus violating Grice’s maxim of quantity by providing more information than is needed; assuming she is being cooperative, she is implying that were another goal salient, the leaves would count as being green.

6 Conclusion

In this paper we proposed occasion-sensitive semantics for objective predicates as a way to account for the phenomenon observed in Travis cases. The explanation for the phenomenon that we suggested exploits the fact that the set of worlds which make a sentence true has additional, internal structure as a result of different ways that a property expressed by a predicate applies to a given individual. We proposed that this structure can be formally represented as a refinement of the bi-partition of logical space induced by the standing meaning of a sentence. Travis cases then show that speakers accept only some of these refinement cells as containing truth-makers for an utterance on a particular occasion. Furthermore, it was suggested that the mechanism behind this selection process is utterly pragmatic and we singled out one particular pragmatic factor as driving this process, *viz.* practical goals (and a concomitant capacity of worlds to support them or not). Starting from the fact that rational agents may pursue different goals on different occasions and that they generally know what it takes to achieve them, it was suggested that their information states can be seen as being equipped with particular goal-conduciveness valuation maps, which in turn play a role in truth-evaluations of utterances on particular occasions. Thus, because in Travis cases on each imagined occasion a different practical goal is salient, the intuitive truth-valuation of an utterance varies, despite the same sentence being used. Additionally, we have shown how

the proposed account applies to different varieties of disagreement, and to intuitions about cancellability, and we argued that our approach preserves the intuitive distinction between discretionary and objective predicates.

Several avenues for future work are open to us at this point. The current account relies on two substantial simplifications, which allowed us to present a simple, perspicuous model of the occasion-sensitivity of the two Travis cases we have taken as our guide. First, in presenting the model, we assumed that the number of ways in which a property expressed by a predicate applies to an individual is relatively small, and also we disregarded the fact that for each such way, there may be a number of sub-ways, and sub-ways of sub-ways (and so forth). This feature of the interpretation of objective predicates would significantly complicate partition semantics that we used to formally model the phenomenon. A more sophisticated account may be able to better represent this feature of nested refinements and explain what constrains their generation. Second, we assumed that each world would get an absolute value for its conduciveness to a goal, which again is a simplification that could be removed in a more complex model. There are some borderline cases of occasion-sensitive predicates (e.g. *under the bed*, *dark*, *hot*, *open*) which are very close to cases of vagueness, and for which a more gradual ranking, taking into account the degrees of goal-conduciveness, would be more natural than the simple valuation we suggested above. Again, omitting this simplification would add more complexity yet also accuracy to the account of occasion-sensitivity.

For now, we hope to have shown that our understanding of sentences containing objective predicates crucially depends on wider, goal-directed projects: insofar as such projects vary across occasions, so do our expectations as to what our sentences say.

References

- 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.
- Chomsky, N. (2000). *New horizons in the study of language and mind*. Cambridge University Press.
- Ciardelli, I., Groenendijk, J., and Roelofsen, F. (2013). Inquisitive semantics: a new notion of meaning. *Language and Linguistics Compass*, 7(9):459–476.

- Ciardelli, I., Groenendijk, J., and Roelofsen, F. (2015). On the semantics and logic of declaratives and interrogatives. *Synthese*, 192(6):1689–1728.
- Ciardelli, I. A. and Roelofsen, F. (2015). Inquisitive dynamic epistemic logic. *Synthese*, 192(6):1643–1687.
- Coppock, E. (2016). Outlook based semantics.
- Davies, A. (2017). Elaboration and intuitions about disagreement. *Philosophical Studies*, x(x):xx–xx.
- Fodor, J. and Lepore, E. (2004). Out of context. In *Proceedings and Addresses of the American Philosophical Association*, pages 77–94. JSTOR.
- Fodor, J. A. (2003). *Hume variations*. Cambridge Univ Press.
- Giberman, D. (2016). Moving parts: a new indexical treatment of context-shifting predication. *Synthese*, 193(1):95–124.
- Ginzburg, J. (1995). Resolving questions, i. *Linguistics and Philosophy*, 18(5):459–527.
- Groenendijk, J., Roelofsen, F., et al. (2009). Inquisitive semantics and pragmatics. In *workshop on language, communication, and rational agency at Stanford*.
- Groenendijk, J. and Stokhof, M. (1984). *Studies on the semantics of questions and the pragmatics of answers*. Ph.D. thesis, University of Amsterdam.
- Huvenes, T. T. (2014). Disagreement without error. *Erkenntnis*, 79(1):143–154.
- Kaplan, D. (1989). Demonstratives. In Almog, J., Perry, J., and Wettstein, H., editors, *Themes from Kaplan*, pages 483–540. Oxford University Press.
- Kölbel, M. (2004). If-faultless disagreement. In *Proceedings of the Aristotelian society*, volume 104, pages 53–73. The Oxford University Press.
- Lasnik, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and philosophy*, 28(6):643–686.
- 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.
- MacFarlane, J. (2014). *Assessment sensitivity: Relative truth and its applications*. OUP Oxford.
- Moravcsik, J. (1994). Is snow white? *SYNTHESE LIBRARY*, pages 71–71.

- 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.
- Parikh, P. (2000). Communication, meaning, and interpretation. *Linguistics and philosophy*, 23(2):185–212.
- Parikh, P. (2006). Radical semantics: A new theory of meaning. *Journal of Philosophical Logic*, 35(4):349–391.
- 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.
- Roberts, C. (1996). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Working Papers in Linguistics-Ohio State University Department of Linguistics*, pages 91–136.
- Roberts, C. (2004). Context in dynamic interpretation. *The handbook of pragmatics*, pages 197–220.
- Roberts, C. (2012). Information structure: afterword. *Semantics and Pragmatics*, 5:7–1.
- Rothschild, D. and Segal, G. (2009). Indexical predicates. *Mind & Language*, 24(4):467–493.
- Schoubye, A. J. and Stokke, A. (2015). What is said? *Noûs*.
- Searle, J. R. (1980). The background of meaning. In *Speech act theory and pragmatics*, pages 221–232. Springer.
- Stojanovic, I. (2007). Talking about taste: disagreement, implicit arguments, and relative truth. *Linguistics and Philosophy*, 30(6):691–706.
- 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. (2008). *Occasion-sensitivity: Selected essays*. Oxford University Press.
- Travis, C. (2009). *Thought's footing: a theme in Wittgenstein's Philosophical investigations*. Oxford University Press.
- Van Rooy, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy*, 26(6):727–763.
- Vicente, A. (2015). The green leaves and the expert: Polysemy and truth-conditional variability. *Lingua*, 157:54–65.