

Predicates in perspective

Anthony Corsentino

Received: 9 November 2008 / Accepted: 5 December 2010 / Published online: 16 December 2010
© Springer Science+Business Media B.V. 2010

Abstract A familiar strategy of argument to the effect that natural-language predicates are semantically context dependent rests on constructing what I term *Travis cases*: different contexts for the use of a predicate are imagined in which its semantic (typically, truth-conditional) properties are claimed to differ. I propose an account of the semantic properties of predicates that give rise to Travis cases; I then argue that the account underwrites a genuine alternative to the standard explanations of Travis cases to be found in the literature; I close with a brief sketch of the connections, required by a fuller development of my account, among the semantic notion of a predicate's *content*, the metaphysical notion of a property's *instantiation*, and the cognitive notion of a language user's *perspective* in using a predicate.

Keywords Context-dependence · Predicate · Extension · Semantics-pragmatics distinction · Contextualism · Semantic content · Properties · Instantiation · Indexicality · Functional role

1 Rules and contexts

Human language is rule governed. The capacities we exercise in giving verbal utterance to our thoughts, and in comprehending others' utterances, involve the application of principles that connect form, meaning, and sound. Yet it is a devilishly difficult problem in philosophy and the cognitive sciences to explain how, in practice, we exploit these principles to find the appropriate words to convey our thoughts—to understand and to be understood.

A. Corsentino (✉)
Department of Philosophy, Boston University, Boston, MA 02215, USA
e-mail: corsent@bu.edu

Language is also context dependent. We must take into account various features of the situations in which we speak if we are to choose the right words to express our thoughts, or to grasp the thoughts expressed in others' words. Here too, language seems rule governed. An English speaker uses 'I' to speak of herself because she grasps the rule that a token refers to its utterer; she uses 'now' (in one of its uses) to speak of the present time because she grasps the rule that a token refers to the time of its utterance; she uses 'this' to speak of a salient object because she grasps the rule that a token refers to an object made salient in the context of the token's utterance. Since she may presume that her interlocutors grasp these rules too, she is entitled to rely upon them in choosing words that express the thoughts she intends to get across.

It is in fitting these two ideas together—that language is rule governed and that it is context dependent—that controversy reigns. One view that has gained increasing prominence in philosophy and the cognitive sciences is that language is *pervasively* context dependent: not only in such familiar cases as 'I', 'now', and 'this', but also among the *predicative* elements of language: adjectives, nouns, verbs, and others.¹ This view is best introduced by considering some examples. I term cases of the following kind "Travis cases," after Charles Travis, who has ingeniously constructed an array of them.²

1.1 Travis cases

Case 1 The house at the corner of Elm and Beech Streets, soon to be repainted, has just received a pressure wash. Years of grime have been washed from the siding. The interior, on the other hand, desperately needs a thorough spring cleaning; it hasn't seen a vacuum, mop, or duster in months. Is this house correctly characterized by the predicate 'dirty'? We can imagine different conversational interchanges:

- (1a) "I hear that it's a housekeeping nightmare over there; no one's bothered to vacuum in months." "You're right; I was inside a few days ago, and I can tell you that the house is dirty."
- (1b) "The folks at Elm and Beech are going to repaint. But they will need to prep the siding." "That's right. In fact, they did a pressure wash, and the house isn't dirty anymore."

Case 2 Suppose that a pane of clear glass has been covered with a transparent yellow plastic film. Is the pane correctly described by the word 'yellow'? We can imagine different contexts:

- (2a) "I want my skylight to cast a yellow glow. Is this pane yellow?" "Yes; the pane is yellow."

¹ By 'predicative' here I intend a *semantic* feature: a predicative element is a term of divided reference, which picks out an object or class of objects by expressing a *condition* on them. Thus 'red' picks out an object by virtue of expressing a condition—redness—that the object satisfies.

² See Travis (1981, 1985, 1989, 1997, 2000); and many others.

(2b) “I need a clear pane of glass. I see that some of these have removable films, such as this one. Is this pane yellow?” “No. Only the film is yellow; the pane isn’t.”

Each case exhibits alternate but coherent stretches of discourse in which, plausibly, the key words ‘dirty’ and ‘yellow’ are being used to effect different classifications. Speaker and hearer in (1a) understand that ‘dirty’ is being used to mark out a classification that includes the house, while in (1b) it is understood that the house is excluded. Similar points apply to the pane of glass as spoken about in (2a) and (2b).

Suppose that this is correct: speakers can use a single predicate to effect different classifications. Should we conclude that ‘dirty’ and ‘yellow’ are context dependent? To answer this question, we need a more precise characterization of context dependence.

Context Dependence: A word is *context dependent* just in case it is possible, consistently with its standing linguistic meaning, for the word to make different contributions to what a speaker (strictly, literally) *says*, and is properly understood to say, in using it.

Case 1 illustrates context dependence only if what the house is *said to be*, and properly *understood* to be said to be, is different in (1a) from what it is properly understood to be said to be in (1b). (For example, we might think that in (1a), the speaker is using ‘dirty’ to say, of the house, that its rooms are dirty, or that it has a dirty interior, while in (1b) the speaker uses the word to say of the house that its siding, or its exterior, is no longer dirty.) That is, it must be the case that the differences in the classifications that a speaker may use ‘dirty’ to effect reflect differences in what the word is used to say, and is properly understood to say.

If we accept this last assumption, then we face a serious difficulty in fitting the idea that language is rule governed with the idea that language is context dependent. One might suppose that the apparent lesson of our two cases generalizes, and that the literally expressed contents of utterances are *pervasively* context dependent. For if Travis cases are constructible for the words ‘dirty’ and ‘yellow’, they are constructible for almost any ordinary predicate in English—or in any other human language. Contextualist explanations of these cases come in two forms. Predicate context dependence is either to be viewed as governed by rules of the semantic component of the language’s grammar (as akin to the familiar context-dependent words like ‘I’ and ‘now’), or to be viewed as lying outside the scope of the grammar. Let us call these options the *grammatical* approach, and the *extra-grammatical* approach, to explaining Travis cases.

Both options have advocates. Jason Stanley and Zoltán Szabó, among others, pursue the grammatical approach;³ Charles Travis and many others pursue the extra-grammatical approach.⁴ Each option faces challenges. On the one hand, those who pursue the grammatical approach must assume that the range of context dependence displayed by ordinary predicates *is* governed by principles that are plausibly conceived

³ See, e.g., Stanley (2000) and Szabó (2001).

⁴ See, e.g., Atlas (1989), Bach (1994), Bezuidenhout (2002), Chomsky (2000), Moravcsik (1998), Recanati (2004), Reimer (2002), Sperber and Wilson (1986); and the articles cited in note 2.

as part of their lexical meanings. Given the seemingly open-ended possibilities for constructing Travis cases for ordinary predicates, one may suspect that predicate context dependence is not the right kind of phenomenon to be explained by reference to context-invariant grammatical principles. On the other hand, those who pursue the extra-grammatical approach must say what semantic theory does explain, if not facts about what speakers say, and are properly understood to say, in using words.

It is possible to deny that Travis cases like Case 1 and Case 2 illustrate context dependence. If the grammatical and extra-grammatical approaches do indeed face the challenges to which I have alluded, then we have a reason to explore this possibility thoroughly. There are two options here. One might think that we should explain *away* the intuition that an ordinary predicate is used to effect different kinds of classification among objects, hence rejecting the very reason that originally motivated postulating predicate context dependence; alternatively, one might maintain that we should accept and explain that intuition, but otherwise than by claiming that the predicates literally *say* different things in different contexts. Later in the paper, I outline some difficulties with taking the first of these options. My main aim, however, is to explore the second option: to find an alternative explanation of how a predicate can effect different classifications, if not by literally saying different things.

I propose such an alternative explanation, which I term the *specificational* approach. Sect. 2 sketches the specificational approach; Sect. 3 describes *contextualism*, and its grammatical and extra-grammatical versions; Sect. 4 outlines difficulties for contextualism, as well as for the best known anti-contextualist approach to Travis cases; Sect. 5 discusses some implications of the specificational approach for the role of the concepts of truth and extension in semantics; and Sect. 6 briefly concludes the paper.

2 Specifications

It will be useful, in describing the specificational approach, to introduce some minimal semantic and metaphysical terminology.

Semantic terminology. Let us say that a predicate—a term of divided reference—characterizes an object or class of objects by *expressing a property*. To say this is *not* to say that a semantic theory should assign intensional entities to predicates; even if a semantic clause ascribes an *extension*, such as a set, to a predicate, it does so by characterizing a condition for membership in that set—and the condition is individuated by a property. But I do assume a connection between *expressing* and *saying*: if a predicate expresses *F*, then it is used to say, literally, of one or more objects that it or they are *F*.⁵

Metaphysical terminology. Since we will speak of predicates as expressing properties, we will be speaking of properties themselves, and of what I term *specifications* of properties. I characterize the notion of a specification in a rough and intuitive way:

⁵ Let all this be taken for granted at this stage of the argument. Later, in Sect. 5, I consider some questions that arise in reconciling the specificational approach with referential semantics.

given a property F and a bearer (actual or possible) x of F , a specification of F is a way for x to instantiate F .⁶

Let us consider some examples. One way for a thing to be a writing instrument is to be a mechanical pencil; another is to be a green dry-erase marker. One way for a thing to be a fish is to be a gravid female neon tetra; another is to be an immature spanish mackerel. There is more than one way to skin a cat—that is, for an event to instantiate the property of *being a cat-skinning*.

What are these “ways of instantiating” a (say, first-order) property? They are themselves (say, first-order) properties. One way for a thing to instantiate the property of *being a writing instrument* is to instantiate the property of *being a mechanical pencil*; another is to instantiate the property of *being a green dry-erase marker*. A thing may, in turn, instantiate each of these properties in different ways.⁷ Let us say that any property P that constitutes a way for a thing to instantiate a property Q is a *specification* of Q . Then it is clear that most properties—certainly familiar properties of everyday objects—have multiple specifications.

Theorists of language occasionally note specificational multiplicity. It is sometimes referred to under such labels as “the efficiency of language” (Barwise and Perry 1983) and “generality of sense” (Atlas 1989). But specificational multiplicity is a feature of the metaphysical structure of properties; it is not (or not fundamentally) a feature of language. It is no requirement on a theory of the meanings of expressions like ‘writing instrument’ and ‘fish’ that it specify the variety of ways for things to instantiate the properties they express.⁸

⁶ Is the terminology of specifications merely a relabeling of a more familiar logico-metaphysical relation, such as that of genus to species, or of determinable to determinate? I shall not attempt a full treatment of that question here, but will make two brief comments. First, it seems that on standard accounts of the genus-species relation and of the determinable-determinate relation (such as in Sanford 2006), the “narrower” property is a *strong* specification of the corresponding “broader” property. (See Sect. 2.1 for the distinction between weak and strong specifications.) Since my proposal involves a plea for recognizing specifications that are not strong, it would appear that the concept of a specification is not readily reducible to either of these more familiar relations, as they are typically understood. Second, the examples of weak specifications given in the text do not meet a plausible ordering condition on the determinable-determinate relation: namely, that any given determinate cannot belong to two distinct determinables neither of which is a determinate of the other. (For example, *scarlet* is a determinate both of *red* and of *colored*, but of these latter two, the first is a determinate of the second.) The weak specifications that I discuss in the text, by contrast, are ways for things to instantiate two distinct properties, neither of which is a determinate of the other (as, e.g., a house’s having a clean interior and a dirty exterior is a way for that house to be clean as well as for it to be dirty). Thus, the relation between a weak specification and the properties it specifies fails to meet the ordering condition. I wish to thank Jacob Beck for discussion of the points in this note.

⁷ The same holds for relations. One way for x to bear the relation of grandmotherhood to y is for x to be the mother of y ’s mother; another way is for x to be the mother of y ’s father. Each of these relations, in turn, may be borne in different ways.

⁸ Now there can be no doubt that normal speakers of English know that ‘neighbor’ is not restricted in its application to individuals of a particular sex. But Atlas’s point is surely correct: we need not see this information as linguistically encoded. We could see a speaker’s possessing the information as the consequence of her possessing two separate pieces of information, one about the grammatical meaning of ‘neighbor’ (that it is true of all and only neighbors, or of nearby dwellers) and one about the metaphysics of neighbor-hood (that one can instantiate the property either (perhaps *inter alia*) by being a nearby male dweller or by being a nearby female dweller. What a speaker knows, then, in knowing that ‘neighbor’ is “sense-unspecified for gender,” is resolved into a piece of linguistic knowledge and a piece of metaphysical knowledge.

2.1 The vagaries of instantiation

Although specificational multiplicity is a metaphysical rather than linguistic phenomenon, I suggest that it discloses the key to understanding Travis cases. To see how, we need to draw a distinction between two different relations that properties can bear to their specifications.

Let us return to Case 1, and consider some evident facts about houses and the property *dirty*. It is a fact that some houses have that property. It is also a fact that some houses have dirty exteriors, while others have dirty interiors. It is clear that having a dirty exterior and having a dirty interior are both ways for a house to be dirty. So the properties *having a dirty exterior* and *having a dirty interior* are both specifications of the property *dirty*.

Now compare the following statements:

- (3) Necessarily, anything that is a dry-erase marker is a writing instrument.
- (4) Necessarily, anything that is a tetra is a fish.
- (5) Necessarily, any house that has a dirty exterior is dirty.
- (6) Necessarily, any pane of glass fitted with a film that filters yellow light is yellow.

Which of these is true and which false? (3) is evidently true, since a thing's being a dry-erase marker necessitates its being a writing instrument. (4) is true likewise. But what of (5) and (6)? It is clear that one possible way for a house to be dirty is for it to have a dirty exterior. A pane of glass *can* be yellow by being fitted with a film that filters yellow light. But does a thing's having a dirty exterior necessitate its being dirty? Does a pane of glass's bearing the property *being fitted with a film that filters yellow light* entail its bearing the property *being yellow*? It seems that the relation between specification and property in these cases is not one of necessitation, or entailment, or any relation the holding of which would guarantee the truth of the generalizations (5) and (6); for (5) and (6) are not true. Let us say, then, that the relation between a property and its specification the obtaining of which guarantees the truth of (3) and (4) is one of *strong specification*; whereas if the corresponding generalization does not hold, let us say that the relation is one of *weak specification*.

I do take it that, at least *prima facie*, (5) and (6) are not true. It may be objected, however, that there is an obvious way to explain that appearance away. One may argue that the difference I have indicated between the pairs (3)/(4) and (5)/(6) is revelatory not of a metaphysical fact about the properties *dirty* and *yellow*, but rather of a linguistic fact about 'dirty' and 'yellow': viz., that they literally express different properties in different contexts of use. A contextualist will take just that position. If that is correct, then the distinction I have proposed between weak and strong specifications cannot be motivated by reflection on statements like (3) through (6). The contextualist may claim, for instance, that (5), uttered in a context in which 'dirty' expresses the property *having a dirty exterior*, will have a reading on which it turns out true.⁹ But uttered in a context in which 'dirty' expresses some other property, (5) will turn out false. On

⁹ I presume that the reading in question would be a trivial truth: *Any house that has a dirty exterior has a dirty exterior*. Note that the contextualist's semantics for 'dirty', if it is to yield that reading, will apparently need to postulate distinct semantic contributions for each of the occurrences of 'dirty' in (5).

my interpretation, by contrast, (5) has one reading, on which (arguably) it does not express a truth. What can I say, against a contextualist, to motivate my interpretation of the difference between the pairs?

I will say two things. First, it is sufficient at the present stage of my argument to claim that it is *possible* to interpret statements like (5) and (6) as illustrating a particular sort of relationship among properties, rather than illustrating a feature of linguistic expressions. Grounds for preferring my interpretation over that of the contextualist emerge later in the paper, when I consider the difficulties that arise for explaining Travis cases along contextualist lines.

Second, it is not clear how a contextualist treatment of ‘dirty’ and ‘yellow’ helps to explain one’s intuitions about the truth values of (5) and (6). The contextualist proposal is that (5) has different readings in different contexts, depending on which property ‘dirty’ expresses: “Necessarily, any house that has a dirty exterior is dirty₁” (i.e., has a dirty exterior), versus “Necessarily, any house that has a dirty exterior is dirty₂” (i.e., has a dirty interior.) But it is not obvious that (5) has such readings. In particular, I doubt that it has a reading on which it (or, more precisely, the clause unmodified by the modal ‘necessarily’) expresses a trivial truth of the form “Any F that is G is G.” It is plausible, then, that the imagined contextualist treatment of (5) predicts nonexistent readings.¹⁰

I have suggested that statements like (3)–(6) provide a rough and ready criterion for distinguishing strong from weak specifications of properties, notwithstanding the contextualist’s temptation to treat the occurrences of the predicates in (5) and (6) as semantically context dependent. My claim, now, is that if F_1 is a strong specification of F , then a predicate that expresses F is not susceptible of Travis cases involving F_1 . But if F_1 is a weak specification of F , then a predicate that expresses F is susceptible of Travis cases involving F_1 . ‘Dirty’ is susceptible of “inside/outside” Travis cases, and ‘yellow’ is susceptible of “fitted with a yellow film/made of yellow glass” Travis cases. That is, whether a particular use of ‘yellow’ correctly classifies a particular pane of glass as yellow can depend on whether it is made of yellow glass, as opposed to being made of colorless glass but fitted with a yellow film. But ‘fish’ is not likewise susceptible of “tetra/mackerel” Travis cases. Whether a particular use of ‘fish’ correctly classifies a thing as a fish cannot depend on whether that thing is a mackerel, as opposed to a tetra. For otherwise, there could be a particular use of ‘fish’ that classified a tetra as a non-fish—i.e., that classified that tetra as a non-fish because it is a non-mackerel. But obviously, as (4) attests, there are, necessarily, no tetrads that are non-fishes.

2.2 Classifying things correctly

I have said that Travis cases involve *weak* rather than strong specifications. I suggest that in the different contexts described in a Travis case, different weak specifications become associated with the particular uses made of the predicate. If a weak

¹⁰ I am grateful to Bernhard Nickel and Jacob Beck for pressing me to confront the contextualist objection envisioned in the last few paragraphs.

specification F_1 becomes associated with a particular use of a predicate that literally expresses F , then F_1 sets a necessary and sufficient condition for the predicate's correctly classifying objects as F . A strong specification sets only a sufficient condition. (A thing's being a tetra is sufficient, but not necessary, for it to be correctly classified as a fish. That is true of any context in which one uses a predicate that expresses *fish*.) Weak specifications become associated with *particular uses* of a predicate, not with *every* possible use. There are two important points to note about the relation of association: (1) weak specifications are associated not with predicate-expression *types* but with predicate-expression *tokens*; (2) the relation of association is "post-grammatical," being neither linguistically represented in the structure of the predicate, nor fixed by the predicate's semantic interpretation—in short, it plays no role in fixing what is said.

We can clarify these points by considering a Travis case. Consider a house that is dirty on the outside, but not dirty on the inside. It is possible to imagine two cases in which a speaker uses the sentence 'The house is dirty' so that, in one case, 'dirty' correctly classifies the house just in case it has a dirty exterior (so that the classification is correct); and in the other case, 'dirty' correctly classifies the house just in case it has a dirty interior (so that the classification is incorrect). What does the speaker *say*, in each case, in uttering the sentence?

To answer this question, we should distinguish what the speaker literally says in using the sentence from the classification she uses the predicate 'dirty' to effect.¹¹ In both of the cases described, what the speaker says is simply that the house is dirty. But in one case, she uses 'dirty' to classify the house as having a dirty exterior; in the other, she uses it to classify the house as having a dirty interior. The standards for the correctness of each classification are given by different (weak) specifications of the property *dirty*. It is for this reason that we can say that each of these uses of the predicate 'dirty' is associated with a specification of the property that it literally expresses.

The distinction between saying and classifying may seem to embody an arbitrary choice of terminology. It does not. Let us call the *grammar* of a language L that family of principles that connect the forms, meanings, and sounds of the expressions comprising L . Then we can say that the difference between saying and classifying is to be explained by reference to the difference between those features of an utterance that are fixed by grammar, and those not so fixed. I assume—substantively—that what a sentence can be used to *say* is something fixed by grammar: if S is a sentence characterized by a grammar G , then G itself determines what S says (and if S contains context-dependent elements, then G will contain rules governing how context fixes the reference of those elements; see Sect. 3.1 below). It is for this reason that we can usually slip harmlessly between speaking of what a speaker uses S to say and what S itself says. But how a speaker might use S (or the relevant predicative constituent of S) to *classify* objects is not fixed by G , although it is constrained by the metaphysics of the property expressed by the predicative constituent of S .¹²

¹¹ I will henceforth assume it to be understood that 'say' means 'literally say'.

¹² I am assuming that the facts that determine which classification a speaker is using a predicate to mark out on a particular occasion, and therefore which specification becomes associated with the predicate's

I am claiming, then, that we should distinguish three facts involved in the foregoing Travis case.

(Fact 1) The speaker utters ‘The house is dirty’.

By virtue of Fact 1 (and given the rules of grammar *G*),

(Fact 2) The speaker says that the house is dirty.

By virtue of Fact 2 (and given the speaker’s aims and intentions),

(Fact 3) The speaker classifies the house as having a dirty exterior (/interior).

Fact 2 is something knowledge of which is guaranteed by linguistic competence, plus knowledge of the relevant contextual determinants of what is said. Fact 3 is something knowledge of which is *optional* with respect to linguistic competence: the hearer might be aware of this, but she might not, consistently with her exercising her linguistic competence to ascertain what the speaker said.

Participants in linguistic interchanges are attuned to the relevance of contextually associated weak specifications. Consider the dialectics that can occur when participants in a conversation justify or challenge an ascription of a property to a thing. You ask me if I’m quite sure that the house is dirty. I reply, “I’m sure it’s dirty. I saw: I was standing right outside it.” Or, if I tell you the house is dirty, you may ask, “How do you know it’s dirty? Did you ever have a look inside?” Dialogues like this reveal that particular weak specifications of *dirty*, though (according to the specificational approach) not literally expressed by ‘dirty’, can nevertheless be implicated in our philosophical description of speakers’ and hearers’ practices of evaluating or justifying particular uses of the predicate.¹³

Let me summarize the view I am proposing. The data to be explained are Travis cases, in which it appears that an ordinary predicate, like ‘dirty’, can be used to effect different classifications on different occasions of its use. I propose the *specificational approach*, which explains Travis cases in terms of the idea that what determines the condition for the correctness of a speaker’s use of a predicate to classify objects is a contextually associated *way of instantiating* the property the predicate expresses. Note that this is *not* a contextualist view, as contextualism was defined in Sect. 1.1. For what varies with context in a Travis case is *not* the property that the predicate expresses, and so not what it literally says about the objects that satisfy it. We should distinguish what a predicate is used to *express* from how it is used to *classify*. The former is fixed by the speaker’s grammar, while the latter is fixed by facts about the speaker’s point and purpose in using the predicate.

Footnote 12 continued

use on that occasion, are facts about the speaker’s cognitive state on that occasion. Cf. the discussion of “perspectives” in Sect. 5.

¹³ The point I make in this paragraph—that speakers can use the same word *form* to characterize and classify things differently—is one that I take to be common ground in the debate over predicate contextualism. Contextualists explain it by postulating semantic differences in literally expressed content; the anti-contextualists I call invariantists explain it by postulating pragmatic differences at an additional, post-locutionary level of communicated content. All of these approaches can appropriate my talk of “weak specifications.” The specificational approach is distinctive not in this choice of terminology, but rather in denying that the contextually associated weak specification of a lexically expressed property need be an ingredient of content at *any* level—semantic or pragmatic—in a communicative interchange. (I thank Jacob Beck for discussion here.)

To place these ideas within a broader theoretical setting, it will be useful to begin with a consideration of contextualist approaches, which explain Travis cases in terms of the idea that what the predicate literally expresses differs in each context. Then, after surveying some of the central difficulties that arise for contextualism, along with the standard form of anti-contextualism that I term “invariantism,” we will be able to appreciate the distinctive explanatory contribution that the specificational approach offers.

3 Contextualist approaches

There are two principal varieties of predicate contextualism. One, which we may call the *grammatical approach*, maintains that the context dependence of what a predicate expresses is governed by a linguistic rule, in the manner of familiar context-dependent expressions like ‘I’, ‘this’, ‘now’, and the rest. The other, which we may call the *extra-grammatical approach*, denies this. Here, the aim is to explain predicate context dependence as a “post-grammatical” phenomenon: linguistic rules fix a certain *meaning* for the predicate in context, and this meaning is what a competent speaker grasps in virtue of applying the linguistic rules; but what the predicate is used to *say* in context is an additional factor, one not fixed (solely) by the linguistic rules. But it is a common presupposition of all contextualist approaches that what a predicate contributes to what is said, in context, is dependent upon aspects of that context. They differ over whether that dependence is *linguistically rule-governed*.

To make these points less forbiddingly abstract, it will be helpful to introduce a few elements of a particular framework for accommodating context-dependence within semantic theory. I choose this framework for purposes of illustration only; what is crucial to the discussion is only the general principle that the framework illustrates.

3.1 Truth-conditional semantics in context

A semantic theory for a language aims to specify the grammatical information that a speaker/hearer draws upon in computing the meanings of utterances of sentences in that language. For sentences that contain no context-dependent elements (if any such exist), the theory will specify their meanings outright; for the rest, the theory will specify the grammatically determined constraints on how information about the context of utterance is to be incorporated into the computation of meaning.

Let us assume that to compute the meaning of a sentence is to compute its truth condition. Then, following Higginbotham (1988), we may impose as an adequacy condition on a semantic theory that it imply, for each sentence of the language, an instance (called a “conditionalized T-sentence”) of the following *conditionalized normal form*:

(7) If U is an utterance of S , and $C(x_1, x_2, \dots, x_n)$, then U is true iff $R(x_1, x_2, \dots, x_n)$.

The context-dependent elements of S will impose constraints, reflected in the antecedent of the conditional, on the values (x_1 etc.) their utterances may take on in context. As a rough illustration, and prescinding from various complicating factors,

including verb tense, consider a conditionalized T-sentence for ‘She called in sick today’:

- (8) If U is an utterance of ‘She called in sick today’, and x is referred to by ‘she’ in the course of uttering U , and x is female, and d is the day of the utterance, then U is true iff x called in sick on d .

The lexical entries for the words ‘she’ (on at least one standard use) and ‘today’ will contribute principles governing the contextual dependence of their truth-conditional properties to the derivation of the conditionalized T-sentence for ‘She called in sick today’.

If we assume, then, that to compute an utterance’s truth condition, as determined in context, is to compute what the utterance (strictly and literally) says and is properly understood to say, then the conditionalized normal form reflects the assumption that context-dependence is linguistically rule-governed. The conditionalized normal form makes clear that to compute the correct truth condition for an utterance of a context-dependent sentence, the speaker or hearer must integrate her knowledge of context-invariant grammatical meaning, which yields principles governing truth-conditional context dependence, with information about the relevant features of context.

I have used the concept of the conditionalized normal form to illustrate the idea that, on most orthodox approaches to semantic theory, the facts to be explained are facts about what sentences *say*, and about how what a sentence says may depend upon features of context in a linguistically systematic manner. We can use the framework to implement some influential recent proposals for explaining Travis cases as instances of linguistically rule-governed context dependence. These proposals exemplify what I call grammatical contextualism.

3.2 Grammatical contextualism

Perhaps the most sophisticated versions of grammatical contextualism have been formulated for adjectives of various kinds: for example, scalar adjectives like ‘tall’, color adjectives like ‘green’, and evaluative adjectives like ‘good’. Kennedy (1999) has argued that a scalar adjective denotes a “measure function” from objects to degrees on an interval scale associated with that adjective. Roughly, an object x satisfies ‘tall’ as used in a given context if the degree on the scale of height associated with x is at least as great as the degree contextually fixed as the “standard” for tallness. Accounts of this kind treat the lexical meaning of a scalar adjective as incorporating a contextually variable parameter. Szabó (2001) has argued that color adjectives and evaluative adjectives, too, have parametric meanings. The semantic representation of ‘green’, on his view, can be written as $(green(C, P))(x)$, where ‘ C ’ ranges over comparison classes and ‘ P ’ over possibly relevant *parts* of an object, each variable to be assigned a value contextually. The semantic representation of ‘clever’ can be written as $(clever(R))(x)$, where ‘ R ’ is a “role-variable standing for a way of being clever” (2001, p. 137).

On Szabó’s account of the lexical meaning of ‘clever’, then, we might expect our semantic theory to derive something like the following conditionalized T-sentence for ‘She is clever’:

- (9) If U is an utterance of ‘She is clever’, and x is referred to by ‘she’ in the course of uttering U , and x is female, and R is the cleverness-role referred to in the course of uttering U , then U is true iff x is clever in role R .

This T-sentence makes vivid a consequence of Szabó’s proposal: for any given context, one does not grasp the truth-conditional contribution of a use of ‘clever’ unless one grasps, of the cleverness-role that is salient in that context, that the relevant object is being said to be clever in *that* role. This consequence holds for any account, like Kennedy’s and Szabó’s, that assigns contextually variable semantic parameters to the meanings of predicates: one does not grasp these predicates’ truth-conditional contributions in context unless one assigns the correct values to those parameters. (I will return to this point in Sect. 4.)

3.3 Extra-grammatical contextualism

The grammatical and extra-grammatical approaches agree that a Travis case involving a predicate P illustrates that P makes different contributions to what is strictly and literally said by different utterances of sentences containing it. The approaches disagree on whether these contextual variations are linguistically rule governed. Extra-grammatical contextualists deny that grammatical rules, learned in the course of learning context-dependent predicates, regulate the dependence of what is said upon context. Unsurprisingly, they typically deny that the meanings of context-dependent predicates can be explicated in terms of notions generally employed to explicate the concept of what is said—such as proposition, reference, and truth. Some, including such advocates of “Relevance Theory” as Dan Sperber, Deirdre Wilson, Robyn Carston, and Anne Bezuidenhout, regard the semantic component of a speaker’s grammar as generating schematic, non-propositional representations that serve as inputs to an inferential procedure designed to maximize “relevance,” that is, to generate the informationally richest propositional representations for the least amount of cognitive processing. The relevance-theoretic explanation of utterance comprehension appeals to such an inferential procedure, one whose domain is not limited to the interpretation of specifically *linguistic* stimuli. Within this framework the computation of what is said is assigned, not to the semantic component of the speaker’s grammar, but to a different cognitive system altogether.

Other extra-grammatical contextualists, like Charles Travis, go further than to deny that what is said is recovered through the application of *linguistic* rules. They deny that such contents are recovered by applying rules at all. For Travis, what is said is *doubly* context dependent: not only can a typical predicate express, compatibly with a unitary linguistic meaning, any of several properties, depending upon the context, but we have no context-independent, explicit way of specifying either this range of properties or the ways that a given context pins one down. For we could manage to do so only if the language in which we described linguistic content contained predicates that did not manifest inexorable context-dependence:

[Such a predicate] would be one which did not admit of different possible understandings of what it would be for some item to fit the description which that

predicate provides (or for the description to fit some item). The right understanding of it would foresee every eventuality in or to which the description might be applied. There is reason to think that no such predicate is available to human beings, at least given the way we in fact cognitively conduct our affairs. (1997, 99)

Travis's view is that where there is no possibility of giving an explicit, context-free characterization of how the property a predicate expresses is fixed in context, there is *a fortiori* no possibility of showing how that information is recovered by the application of (context-free) rules.¹⁴

I should emphasize, once again, that contextualist views of all types typically appeal to a conception of what is said: what is strictly and literally expressed in or by uttering a sentence, as distinct from what may thereby, but in addition, be communicated.¹⁵ What motivates a contextualist theory for a given expression is the intuition that what it contributes to what is literally expressed varies with the context of use. The difficulties for contextualist approaches to explaining Travis cases derive principally from this intuition. Let us now consider some of these difficulties.

4 Difficulties for the standard approaches

I do not pretend, by the criticisms I shall advance in this section, to refute any of the standard approaches to Travis cases that I consider, contextualist or anti-contextualist. The criticisms' purpose is to raise *prima facie* doubts about these standard approaches, and thereby to motivate the key thesis on which the specificational approach differs from the others: that the content of a predicate is to be analytically distinguished from the type of classification that a speaker might use the predicate to effect. I begin with contextualism, turning from there to the most familiar form of anti-contextualism.

4.1 Difficulties for contextualist approaches

A contextualist treatment of a Travis case locates a contextual difference in the contribution of the predicate to what is literally said. It follows, then, that a hearer in a given context must grasp this contribution on pain of failing to grasp what was said. If in a particular context a speaker uses 'yellow' to say of the pane that it is made of yellow glass, then one cannot understand the word's locutionary content as used in that context unless one recognizes that it makes that contribution to what is said—unless, that is, one understands the utterance, in that context, as saying of the pane that it is made of yellow glass.

¹⁴ See Travis (1990) for a criticism, along the lines just sketched, of the relevance theorists' attempt to characterize utterance interpretation within a computational framework.

¹⁵ Extra-grammatical contextualists employ different notions of what is said; some maintain that what is said is fully propositional, while others do not. Despite this difference, they agree that grammatical meaning fails to include a rule determining a function from contextual factors to the contents of what is literally expressed, or implicit therein.

A grammatical contextualist treatment of ‘yellow’ would assimilate it to uncontroversial pieces of context-dependent language. Consider evident indexicals like ‘I’ and ‘that’. If we find ourselves in a situation in which the referents of these words are not clearly salient and non-trivially identifiable (i.e., identifiable other than as “the utterer of ‘I,’” or “the thing meant by ‘that’”), then we will attain at best a partial grasp of the utterance’s content. If we hear a stranger around the corner and obscured from view utter ‘I like that one’, then we can deduce that someone has said, of some salient thing, that he or she likes that thing. But if this counts as grasping what was literally said, it is something like a degenerate case. It seems more like a *schematic* understanding of the utterance than a full understanding.

But does our understanding of what was said with ‘yellow’ seem schematic, or partial, in *this* sense? There seems nothing intrinsic to the meaning of the word that would suggest that some kind of contextual supplementation is necessary to grasp its propositional contribution. One may partially understand a word, of course (technical jargon, for example); but this is obviously a different matter. When someone utters ‘It’s yellow’, referring to a pane made of colorless glass but fitted with a yellow film, there is no particular contextual parameter to which we must direct our attention in order to fix the meaning of ‘yellow’ and grasp what the speaker is saying. Specifically, it is not necessary for us to understand the point and purpose of the speaker’s ascription for us to understand that she is describing some pane of glass as yellow. For it seems clearly possible for a hearer to understand ‘yellow’ as meaning *yellow* even though she *lacks* information relevant to determining what sort of classification the speaker aims to mark out by using ‘yellow’ (e.g., a pane that is made of yellow glass, or that is merely fitted with a yellow film). It is hard to understand how a hearer’s ignorance of this aim could prevent her from grasping an utterance of ‘yellow’ as meaning what it appears to mean, namely, *yellow*.

But would a contextualist object to this last point? After all, a precisely analogous point holds about the pronoun ‘I’: namely, that a hearer can understand ‘I’ as referring to the speaker even though she lacks information relevant to determining who (or which thing) that is. But then, if this point about ‘I’ is compatible with the contextualist’s claim that ‘I’s meaning contains a parameter that must be filled in the semantic interpretation of its utterances, why not regard my point about ‘yellow’ as similarly compatible with a contextualist treatment of that predicate? Suppose, for ease of illustration, that the contextualist’s treatment of ‘yellow’ takes the linguistic meaning of that word to incorporate a semantic parameter to be filled by the contextually determined weak specification of *yellow*. Then why not say that a speaker who hears ‘It’s yellow’, but who is unable to make a non-trivial identification of the weak specification associated with ‘yellow’ in that utterance, can nonetheless deduce that the speaker has said, of some particular thing and of some particular specification, that *that* thing bears *that* specification of *yellow*?

I suggest that the key difference between ‘I’ or ‘it’, on the one hand, and ‘yellow’, on the other, lies in the token-reflexivity attaching to the meanings of the former. We can claim that we know what is being said, in respect of the contribution of a given token of ‘I’, provided that we have a way of identifying that token’s utterer. I can say that the person who just uttered ‘I like that one’ said that he or she likes some particular salient thing simply by being able to identify the utterer *as* the utterer, and as

the thing demonstrated *as* the thing demonstrated—whoever and whatever they might be.¹⁶ I am able to do this because I know the token-reflexive rules that determine (or at least constrain) the reference of ‘I’ and ‘that’. But it is precisely because of the possibility of such *trivial* identifications of the referents of ‘I’ and ‘that’ that we might regard some ways for a hearer to represent the content of an utterance containing ‘I’ or ‘that’ as schematic; indeed, it might be acceptable to say in such cases (even if only “in a manner of speaking”) that the hearer does not know what has been said. But even if we suppose that a hearer who can make only a trivial identification of the referent of a token of a token-reflexive expression can still count as knowing the utterance’s content, it is important to note that the trivial mode of identification yields a metalinguistic representation of that content. With ‘I like that one’, the trivial modes of identification of the referent of ‘I’ include “the speaker (whoever that is),” “the utterer of ‘I’ (whoever that is),” “the person saying this (whoever that is),” and the like; and they yield such metalinguistic representations of the utterance content as “The speaker (whoever that is) is saying that he or she likes whatever it is he or she is calling ‘that one’.”

It is doubtful that anything analogous is true of ‘yellow’. Even if we are ignorant of which specification attaches to the speaker’s use of ‘yellow’, we require no metalinguistic representation of what she is saying in using that predicate. We may admit that we do not know *what she meant* by calling it yellow; we may even allow that, for this reason, we cannot say whether she is *right* in so calling it. But ignorance of what was meant need not, of course, entail ignorance of what was said. If the meaning of ‘I’ is token-reflexive, but the meaning of ‘yellow’ is not, we can explain the disanalogy. To identify the utterer of a token of ‘I’ is to identify that token’s referent, and therefore any difficulties in making the identification constitute difficulties in grasping its semantic content. But if there is nothing in the meaning of ‘yellow’ that requires assigning a contextual value to a semantic parameter, then we would expect that the difficulties that might arise in understanding a speaker’s use of ‘yellow’ pertain to understanding how she is using it, rather than with understanding what she is expressing in uttering it.

Of course Szabó and others have made specific proposals about the contextual parameters that color adjectives like ‘yellow’ lexically determine. But first, the character of such parameters in our linguistic knowledge would be quite different from those of words like ‘I’, with evidently token-reflexive meanings, since there is no reason to expect that theoretically unsophisticated speakers could report which parameters these might be and how, in any given case, they determine their values; and second, there might be no clear limit to the dimensions and varieties of contextual variation a Travis case for ‘yellow’ could be constructed to display.¹⁷

¹⁶ These descriptions of utterer and demonstratum are to be construed attributively, since the referential identifications made by their means are trivial.

¹⁷ Here I find myself in apparent disagreement with Szabó. See Szabó (2001, p. 138): “An object is green if some contextually specifiable (and presumably sufficiently large) *part* of it is green.” Thus Szabó includes a “parthood” variable in ‘green’s lexical representation. But if parthood is a lexically determined parameter, why isn’t, say, the quality of illumination? Sometimes—but not always—one calls an object green when it is illuminated in green light, even if it would present a rather different appearance under other, more typical,

The points I have made apply to extra-grammatical contextualism too, since it is a common assumption of all versions of contextualist explanations of Travis cases that the contribution made by predicates figuring in such cases to the propositional content literally expressed—to “what is said”—can vary with features of the context of utterance. I have argued that reflection on what is involved in comprehending speech involving ordinary predicates reveals the *prima facie* implausibility of this basic contextualist assumption.

Travis cases, then, do not *in themselves* reveal context dependence in the conditions on understanding predicates. Rather, the contextual variation evinced by Travis cases reveals something about the kinds of *classifications* that a speaker may use a predicate to mark out: the utterer of ‘yellow’ in (2a) classifies the pane differently from the utterer in (2b). That variation is irrelevant to the conditions a hearer must satisfy to understand what ‘yellow’ expresses in each situation.¹⁸

Suppose we grant that one’s ability to grasp what is said by a use of a predicate does not involve taking into account the contextual differences on which Travis cases depend. To grant this is not to deny that in a wider sense, the speaker’s point and purpose in uttering the sentence cannot be achieved unless the hearer is in a position to appreciate the kind of classification the speaker is making by using the predicate. Misunderstandings in this wider sense are possible, and they can undermine our aims in interlocution.

Imagine, for instance, a variation upon Case 2. Suppose that someone enters my shop and says, “I’d like a pane of yellow glass.” I show her a pane of clear glass covered in a transparent yellow film; she replies that this isn’t the sort of thing she wants. Perhaps I have misunderstood *her*; perhaps she has failed to make her intention clear. We have a case of miscommunication. But to explain what has gone wrong, we need not attribute to anyone a *semantic* misapprehension—a misrepresentation (on the part either of speaker or of hearer) of what is said, in that context, in uttering ‘I’d like a pane of yellow glass’. I haven’t failed to grasp what she said in uttering that sentence; and she hasn’t failed to say what she intended *if she intended to say that she would like a pane of yellow glass* (among the many other things she may have intended to do, or meanings she may have intended to get across, in and by uttering those words). If she did *not* so intend, then she indeed misspoke. But it is coherent to imagine that she did so intend, although she miscalculated in assuming that *she* would be understood by saying to me that she would like a pane of yellow glass. So we will have imagined a communicational misfire that arises because the hearer has failed to recognize the relevant standard for classifying a pane of glass as yellow, a failure for which (perhaps) the speaker is responsible. The misfire is explained by reference to a factor that operates to generate a Travis case for ‘yellow’; but the misfire is not explained by reference to anyone’s having misrepresented the semantic content that, in context, is determined by ‘yellow’s lexical meaning. We can explain what goes

Footnote 17 continued

lighting conditions. (Though I do not wish to argue the point here, I find it implausible that such uses of color terms are, though pragmatically acceptable, literally false.)

¹⁸ I think that further corroboration of this point comes from considering our practices of indirect speech reporting. For more on this topic, see Cappelen and Lepore (2003, 2004).

wrong in a case of this kind without assuming that the predicate at issue is context dependent.

A contextualist treatment would explain the misfire by assuming that my interlocutor used the sentence ‘I’d like a pane of yellow glass’ to *say* that she would like a pane of glass that is yellow “through and through,” or something of the sort, whereas I understood her to be saying something different, such as that she would like a pane of glass that filters yellow light. My point, then, is that first, there is an at least equally plausible explanation of such misfires that does not assume that speaker and hearer understand the *utterance* differently; and second, that there are independent reasons to think that the contextual factors on which this different in utterance content is alleged to depend are irrelevant to understanding the utterance.

I conclude, then, that the specificational approach offers a way to explain the intuitions elicited in Travis cases that is substantially different from contextualist approaches, and that the specificational approach is consistent with further considerations about these cases that contextualism cannot accommodate. There is, however, an anti-contextualist, or *invariantist* approach to Travis cases that I have not discussed, by comparison with which I can bring out more clearly the distinctive commitments of the specificational approach. Let us turn now, therefore, to a brief critical review of invariantism.

4.2 Difficulties for invariantist approaches

The invariantist approach is utterly straightforward: it shares the specificational approach’s assumption that a Travis case illustrates no context dependence in what a speaker says, but it maintains that the predicate marks out the *same* classification in each case. Returning to Case 2, for instance, invariantism maintains that the pane of glass covered in transparent yellow film is classified as satisfying ‘yellow’ or as not satisfying ‘yellow’ in *whatever* context it is used. This implies that in at least one of the conversational interchanges, the answer given to the question ‘Is the glass yellow?’ is strictly speaking *wrong*. If the right answer is ‘Yes’ in *any* context, then the right answer is ‘Yes’ in *all* contexts. I interpret Cappelen and Lepore (2005) as endorsing the invariantist approach.

I claim that the invariantist approach, simple and intuitive though it may be, carries substantial and revisionary implications about the norms that govern ordinary linguistic practice. If ‘yellow’, as used in both interchanges (i.e., (2a) and (2b)), correctly classifies the pane, then it is incorrect for the speaker in (2b) to withhold the predicate. If it misclassifies the pane in both interchanges, then it is incorrect for the speaker in (2a) to ascribe the predicate to it. Whatever assumption we make, it would seem to follow that the use of ‘yellow’ in at least one of the two interchanges results in an incorrect ascription.

There are two problems here, one pragmatic and one semantic. The pragmatic problem falls into two parts. First: if Travis cases involve the misapplication of the relevant predicates, we require a satisfactory explanation of why such misapplications occur. Either the speaker misapplies the predicate knowingly or unknowingly. Why would she knowingly misapply the predicate? It would seem that a misapplication

would be easily avoided, simply by choosing a predicate phrase that semantically encodes the relevant specification. (The speaker in (2a) could have used a phrase like ‘tinted yellow’; the speaker in (2b) could have used a phrase like ‘contains colorless glass’.) Why would the speaker knowingly misuse her semantic resources? Suppose, on the other hand, that she is unaware that she has chosen a semantically unsuitable predicate to effect her intended classification. Is it plausible to ascribe semantic ignorance to the speakers in our imagined interchanges? The invariantist approach needs to explain why the semantic content of ‘yellow’ might be epistemically opaque to a competent speaker.

A second part of the pragmatic problem is to explain how an incorrect ascription nevertheless manages to serve the speaker’s communicational purposes (as it evidently does). This, of course, is a general challenge facing a Gricean “error-theoretic” treatment of any area of discourse. The challenge seems less pressing in areas that involve nonliteral or loose speech (such as irony and hyperbole). Travis cases, however, are not plausibly explained in terms of nonliterality. Case 2, certainly, seems to resist such treatment. In which of the two contexts—(2a) and (2b)—is one plausibly speaking nonliterally? It is not obvious that ‘yellow’ as applied to the pane *strictly* classifies it in one way, as opposed to the other.

The pragmatic problem of explaining which use in a Travis case is, as it were, “faultlessly incorrect,” carries with it the semantic problem of singling out the classification the predicate marks out in every context of its use. The contextualist is likely to complain that there can be no such classification. The invariantist approach faces the difficulty that we lack any robust intuition as to which of the two utterances—‘Yes, it’s yellow’; ‘No, it isn’t yellow’—is literally false. One might be tempted to choose the genuinely “literal” reading of ‘yellow’ on the basis of its greater naturalness. One might claim that in calling a glass pane yellow, it is much more natural to mean that the glass itself is yellow, regardless of the tint of the film laid upon it. Perhaps, indeed, this is the “unmarked” reading of ‘yellow’ in this context. But we should be wary of reading semantic significance into facts about the typicality of usage. It might turn out that hearers, in default of appropriate background information, will be “primed” to interpret a speaker’s ascription of ‘yellow’ in accordance with the most typical use. Such facts are of great interest for the psychology and epistemology of communication; but it requires justification, to say the least, to draw any conclusions about the semantic properties of ‘yellow’ on their basis.¹⁹

In summary, the specificational approach I have sketched proposes that for predicates capable of figuring in Travis cases, we should distinguish the predicate’s semantic

¹⁹ At this juncture a comment on the label ‘invariantism’ is in order. The invariantist approach, as I construe it, differs from the specificational approach in how it conceives the relation between the literally expressed and linguistically determined *content* of an utterance containing a predicate like ‘yellow’, and the correctness conditions of the *classification* that the user of the predicate is making in producing that utterance. On the specificational approach, but not on the invariantist approach, content remains invariant while classification-conditions vary. I am thus using the label ‘invariantism’ in a more specific sense than the broader, but also useful, sense in which any view counts as “invariantist” that explains Travis cases without supposing a contextual variation in the literally expressed and linguistically determined content of utterances of the predicate. The specificational approach, of course, is invariantist in this broader sense: *content*-invariantist (i.e., anti-contextualist).

content from the correctness conditions of the various classifications that a speaker can use the predicate to effect. What varies from context to context are the latter; the predicate's grammatically determined semantic content remains invariant. From the standpoint of the specificational approach, both contextualism and invariantism catch glimpses of different parts of the truth about the semantics and pragmatics of ordinary predicates: contextualism rightly acknowledges the thought that speakers can use predicates to effect different classifications, but invariantism rightly denies that thought the kind of semantic significance the contextualist takes it to have.²⁰

5 Predicates, properties, extensions, and perspectives

My argument to this point has focused on the question of what Travis cases reveal about what a speaker uses a predicate to say. I have drawn a distinction between saying and classifying, and I have claimed that while facts about what predicates contribute to what is said are semantic (and grammatically determined) facts, cross-contextual differences in how a predicate is used to classify objects are non-semantic differences. These cross-contextual differences, I have claimed, are differences among the standards for the correctness of particular uses of a predicate to classify objects. It might fairly be remarked, though, that talk of such standards seems to belong nowhere but to the conceptual territory of referential *semantics*. For it seems evident that a standard for correct classification, insofar as it attaches to a predicate, determines that predicate's satisfaction condition, hence the condition that defines the predicate's extension. According to the specificational approach, the classificatory standards associated with a predicate are contextually variable; but then so, surely, is the predicate's extension. The concepts of extension and satisfaction conditions, however, are central to referential semantics. It therefore seems that on my view, the linguistically determined seman-

²⁰ So-called "truth-relativist" accounts of the semantics various types of expressions (such as epistemic modals in Egan et al. (2005), predicates of personal taste in Laserson (2005), and future contingents in MacFarlane (2003)) deserve acknowledgement as well. Of the frameworks currently under discussion in the literature on philosophical semantics, it is the truth-relativist approach that is perhaps most congenial to the specificational approach; an anonymous reviewer observed that a truth-relativist treatment of Travis cases allows one to accommodate the key claim of the specificational approach that the semantic content of a predicate in a Travis case remains invariant while its extension varies.

I grant that the truth-relativist framework attractively accommodates Travis cases, as John MacFarlane has himself suggested (2009, p. 246). I am reluctant to allow truth relativism the final word on these cases, however, since the types of expression for which the truth-relativist treatment seems most attractive (e.g., epistemic, evaluative, and temporal locutions) are ones in which it is at least arguable that the terms in which we semantically evaluate utterances containing them involve some sort of reference to context—either the speaker's or the evaluator's. ("She said that it's fun. Maybe that's true for her.") This does not seem to be the default case for predicates susceptible of Travis cases. It is by no means obvious that we employ a speaker- or assessment-relative notion of truth when we evaluate the truth of what is said by a speaker who associates a particular weak specification with her use of 'yellow'. The specificational approach maintains that different speakers and hearers/evaluators in relevantly different contexts think of the semantic content of 'yellow' (namely, the property *yellow*) differently; but I think it is possible to explicate these differences in the speakers' and evaluators' "ways of thinking" without attributing to them thoughts *about* context, or about parameters to which semantic evaluation is putatively relativized. Cf. the discussion of perspectives in Sect. 5. I hope that these brief remarks will suffice in the present context; I grant that the relations between the specificational approach and truth-relativist frameworks deserve fuller consideration.

tic content of a predicate, insofar as it is cross-contextually fixed, does not determine its (contextually variable) extension. The specificational approach therefore seems to imply that a predicate that participates in Travis cases lacks a linguistically determined extension, and therewith, a referential semantics.

Worse, it seems that if the foregoing implications hold, then incoherence threatens. For the specificational approach would be committed to the following claims about, for instance, Case 2:

- (10) In interchange (2a), the speaker uses ‘yellow’ in such a way that its extension includes the pane.
- (11) In interchange (2b), the speaker uses ‘yellow’ in such a way that its extension excludes the pane.
- (12) In both interchanges, ‘yellow’ literally expresses the same property.

(10) and (11) follow from the assumption that the weak specification of *yellow* associated with ‘yellow’ as used in (2a) differs in the relevant respect from the one associated with ‘yellow’ as used in (2b). But (12) apparently implies that the *grammar* assigns to ‘yellow’ a uniform extension: that borne by the property *yellow* (since the latter, we have been assuming, individuates the semantic content of ‘yellow’). Contradiction results, unless we draw a distinction between the extension of ‘yellow’ as used on a particular occasion and the extension of ‘yellow’ as determined by the grammar, perhaps by invoking a Kripkean distinction between “speaker’s reference” and “semantic reference.”

I doubt that it will help, in resolving this puzzle, to invoke the distinction between speaker’s reference and semantic reference. For I maintain that the thought the speaker intends to get across to her hearer, in uttering ‘The pane is yellow’ in (2a), need be nothing other than the thought that the pane is yellow, and not (for instance) the thought that the pane filters yellow light. In that case, it would be a mistake to posit a distinction between two levels of communicated content, by reference to which we may distinguish what the predicate picks out and what the speaker uses it to pick out.

I suggest that a satisfactory description of what is at work in a Travis case must include reference to something that the argument to a contradiction from (10), (11), and (12) leaves out: *reference to the speaker’s perspective*. To describe a Travis case, we require a notion of how a speaker thinks of a property—more specifically, of what the speaker is taking it to *come to* for a thing to instantiate a property. She may think of *yellow* in the way reflected by (2a), so that a pane of glass’s being yellow *just is* its filtering yellow light. If she does, then we should say that in thinking that the pane is yellow, *and* in thinking that the pane falls in the extension of ‘yellow’, she is thinking that it filters yellow light. When we consider (10)’s and (11)’s reference to the extension of ‘yellow’ as the speaker uses it, we should be mindful of the distinction between what the speaker *does* with the predicate and what she *believes* (or, if it matters, *knows*) about its semantics. She is doing different things—effecting different classifications—in the interchanges (2a) and (2b). Yet in each interchange, she has the same semantic belief about ‘yellow’: it picks out the yellow things. One may be tempted to ask whether, according to the specificational approach, it is the correctness condition of the classification, or instead the property that figures in the speaker’s semantic belief (i.e., *yellow*), that determines the extension of ‘yellow’ as used in

(2a) and in (2b). My point is that from the speaker's perspective, this is a distinction without a difference. Our own perspective as evaluators of the speaker's classification allows, and perhaps even requires, us to be more specific about what she "means" by calling the pane yellow—to use a more specific form of words to characterize her classification where *she* would simply use the word 'yellow'. But it is *yellow* nonetheless that she is thinking and talking about.

We should interpret (10), (11), and (12) as follows. (10) and (11) reflect facts about the correctness conditions governing the speaker's classifications in (2a) and (2b). (12) reflects a fact about the linguistically determined semantics the speaker believes 'yellow' to have. What unites (10) and (12) is that in interchange (2a), the speaker is thinking of the pane's being yellow in such a way that the correctness of the thought that it *is* yellow depends on the pane's satisfying the relevant specification there in force. The same holds true *mutatis mutandis*, with a different specification in force, for (11) and (12). Since the correctness conditions of these classifications are not aspects of the semantics of 'yellow', it is misleading to speak crypto-semantically, as (10) and (11) do, of the extension(s) of 'yellow' as used on those particular occasions. (10) and (11) are better phrased in a way that refers to how the speaker is classifying the pane, eliminating reference to the *word's* extension.

This does not mean that we should stop speaking of 'yellow' as having an extension at all. On the contrary: as I will explain in Sect. 5.3 below, there is nothing in the classificational approach to jar with the idea that a speaker has beliefs about the extensions of the predicates in her language. But the distinction between classification and meaning requires that we handle our semantic terminology with care, and be chary of conflating the speaker's perspective on her linguistic acts with our own perspective as evaluators.

But what about the extension of *yellow* itself? Even if we grant that classification is one thing, and meaning another, what of the fact that 'yellow' expresses *yellow*, which in turn has a perfectly definite extension? Surely, we must still allow that 'yellow' has a semantics that imposes that uniform extension in whatever context it is used—call it the extension of *yellow*, *tout court*. Further, this fact would seem to undermine the claim I made above, to the effect that speakers can take different perspectives on *yellow* in their *yellow*-thoughts, in which different weak specifications come to the fore. After all, if *yellow* has an extension *tout court*, then any thought about *yellow*—in which that very property serves to individuate that thought's content—will be a thought about, *inter alia*, all and only those things that fall into *yellow's* extension. If a thinker deploys the concept YELLOW, or utters the word 'yellow', in classifying a thing *x* as yellow, then the conceptual or linguistic classificatory act she performs is one whose correctness depends on whether *x* falls in the extension of *yellow*, *tout court*. All these points seem merely truisitic. What conceptual space exists for a distinction between the extension of *yellow*, *tout court*, and its extension as determined by the perspective the thinker takes on it (if this is supposed to differ)?

5.1 Extension ascriptions deflated

We need to distinguish two styles of talk about the extensions of properties: a deflationary style and an inflationary style. Extension-talk in the deflationary style is

metaphysically trivial: the extension of F consists of the F s; the extension of *yellow* consists of the yellow things; and so on, innocently, for any property we care to name. Extension-talk in the inflationary style is metaphysically loaded: the extension of F is, putting it picturesquely, the upshot of the way F sorts all of the elements of reality, actual and otherwise, into F s and non- F s (ignoring issues of no relevance to our present concerns, such as vagueness). Extension-talk in the inflationary style, however, goes along with a way of thinking of properties that allows no conceptual space for weak specifications. On this way of thinking of properties, *yellow* divides reality along a yellow/non-yellow line: *yellow* puts the pane of glass on one side of that divide or the other. This is to say that there can be no such thing as a weak specification of *yellow*: if *yellow* puts the pane on the non-yellow side of the dividing line (let us suppose), then the way the pane is could not be a way of being yellow. The specificational approach, of course, is committed to rejecting such a way of thinking of properties. A weak specification of F is a way for a thing to have F that does not metaphysically necessitate having F itself. Cognitive agents can use weak specifications of F to sort things into F s and non- F s, and different specifications yield jointly incompatible, albeit individually acceptable, sortings. The specificational approach does, then, make a substantive assumption about the nature of properties. It rejects the picture of properties as laying down a dividing line across metaphysical space in relation to which each thing falls on one side or the other. The specificational approach maintains that it is cognitive agents thinking F -involving thoughts, and not F itself, that sort the F s from the non- F s.²¹

The specificational approach prefers, therefore, the deflationary style of talking about properties' extensions. We can speak neutrally and innocently of the yellow things. It is possible to give a hard metaphysical edge to this talk, of course: a skeptical philosopher may ask whether "the pane in itself really is, or really is not, among the yellow things." To raise the question is to talk in the inflationary style, and the friend of weak specifications will not accept such talk uncritically. But talk of "the extension of *yellow*, *tout court*" will be rejected too, since it has a similarly inflationary inflection. When we adopt the specificational approach to describing and explaining thinkers' classificatory practices with predicates, we are better served by eschewing talk of extensions altogether, save the innocent, deflationary variety.

The upshot of this subsection, and of the immediately preceding paragraphs, is that we should block the inference to a contradiction from (10), (11), and (12) by rejecting inflationary talk of the extension of the property *itself* (the dividing line the property intrinsically marks out), as opposed to the *particular* demarcations that thinkers make with particular uses of predicates expressing that property. To say that *yellow* is borne by yellow things is harmlessly trivial. If we want a substantive account of the correctness conditions of speakers' uses of predicates, we must take into account their own sorting activity: we must look to their perspectives on the properties that figure in

²¹ What is the *scope* of the thesis that it is cognitive agents, and not properties themselves, that sort F s from non- F s? Does it apply to all properties, or only to those that figure in Travis cases? If the latter, then can we give an illuminating philosophical account of such properties as a kind? A satisfactory elaboration of the metaphysical assumptions of the specificational approach must address these large questions. It will serve the purposes of this paper well enough if I restrict the thesis to those properties that figure in Travis cases.

their predications. In doing so, we may find ourselves referring explicitly to the weak specifications that figure only implicitly in the speakers' own thoughts. We will then be characterizing perspectives, but not, I suggest, by ascribing concepts that directly encode those specifications. Rather, we will be describing functional or dispositional facts about the speakers' use of those concepts that directly express the properties the specifications of which are our focus. The next subsection clarifies the point.

5.2 Perspectives as cognitive dispositions

It is common in post-Fregean philosophy to place various explanatory burdens on the distinction between the thing that is thought of and *how* that thing is thought of. Philosophers and cognitive scientists invoke such notions as senses and concepts to explicate the distinction. One may, for instance, think of *yellow* under the concept (or sense) YELLOW, or under the concept (or sense) MY DAUGHTER'S FAVORITE COLOR, the difference between them helping to explain certain types of failure, familiar since Frege, to integrate various pieces of information a single thinker might have about *yellow*. The "perspectives," or "ways of thinking," that the specificational approach requires, however, are individuated differently from senses and concepts, at least as these latter are standardly invoked in theorizing about communication and cognition. The specificational approach requires a notion of ways of thinking that ties them to the various weak specifications of a given property that are relevant to explaining Travis cases. The difference between thinking of *yellow* under YELLOW and of thinking of it under MY DAUGHTER'S FAVORITE COLOR cuts across the differences in the weak specifications of *yellow* for which the specificational approach requires a notion of a way of thinking of that property. Regardless of which of these two concepts (/senses) under which the speaker thinks of the color of the pane—and so regardless of whether the speaker utters 'The pane is yellow' or 'The pane is my daughter's favorite color'—the differences among weak specifications of *yellow* to be accounted for in a Travis case still arise. To take a particular perspective on *yellow* is not a matter of deploying a particular concept that expresses something other than *yellow*. It is to deploy a concept that expresses *yellow*, but to deploy it in a particular style.

I suggest, then, that the notion of a "perspective" on a property that we need is a *functional* notion: it is given by the inferential dispositions, or (more broadly, perhaps), the dispositions to making certain cognitive transitions, in which the speaker's concept of the property participates in the stretch of the speaker's cognition of which her use of the predicate is a part. The way in which the utterer of 'The pane is yellow' in (2a) is thinking of *yellow* is given by such functional facts as that she has just heard her interlocutor speaking about the color of the light he wants his skylight to filter. Compare this with (2b), in which the interlocutor has specifically set aside the coloring effect of the film on the pane. The speaker in (2a) has been caused to be disposed or "primed" to draw an inference like 'filters-yellow-light(x) \rightarrow yellow(x)';²² the speaker in (2b)

²² Here and in the next paragraph, I am making free and casual use of the notation by which I am representing the speakers' inferences. For instance, it is probably desirable to interpret the variable ' x '

has been caused to be disposed *not* to draw that inference. Each speaker deploys YELLOW (let us suppose); but in each speaker's cognition, YELLOW adopts a different inferential profile.

Considering Case 1 likewise: the way in which the utterer of 'The house isn't dirty' in (1b) is thinking of *dirty* is given by such facts as that his interlocutor has focused the conversation on the condition and appearance of the house's exterior. Given the information he has received in the conversation, he has been caused to be disposed to draw the inference ' $\text{dirty}(x) \rightarrow \text{not-dirty-exterior}(x)$ '; his counterpart in (1a) has been caused to be disposed precisely not to draw that inference. I suggest that it is in terms of such inferential dispositions as these that we should spell out the kind of talk of "perspectives" or "ways of thinking of a property" required by the specificational approach, and therefore it is in such terms that we should explain how a property becomes associated with a particular weak specification in a Travis case.²³

5.3 Specifications and referential semantics

Given the position I have staked out in Sect. 5, there is a way to reconcile the specificational approach to explaining Travis cases with a referential framework for semantics. Suppose that among the things a speaker believes of the predicate 'yellow' is that it applies to a thing x just in case x is yellow (or that she believes that the extension of 'yellow' is $\{x : x \text{ is yellow}\}$); and suppose that this belief is drawn upon in computing a truth condition for 'The pane of glass is yellow', to wit, that it is true if and only if the pane of glass is yellow.²⁴ These suppositions—provided that the semantic terminology is construed in the appropriately deflationary style—are consistent with the claims I made above about the relation between perspective and extension. For the semantic beliefs I have mentioned are specified by their contents; in attributing them to a speaker, we presuppose nothing about how she is thinking about the property

Footnote 22 continued

as domain-restricted. Here, we might do well to let it range only over pieces of transparent or translucent material. In the inference described in the next paragraph, it might be best to let ' x ' range over houses or similar structures to which an interior-exterior contrast is applicable. My imprecision in these matters is principled. My aim here is to introduce the idea of perspectives as inferential dispositions, and not to discuss how best to delimit and individuate these dispositions. (Note 23 announces a similar explanatory disclaimer.)

²³ A reviewer of an earlier draft pointed out that one can have all sorts of perspectives on a property that do not contribute to isolating a weak specification of it. I may, for instance, be regularly disposed to infer that a thing is beautiful from the assumption that it is yellow. That inferential disposition, however, obviously cannot make it the case that *yellow*, as I think about it, is restricted to beautiful things. This point is quite correct. A theorist of specifications should explain how, for any given property F , we should distinguish three types of inferential/cognitive dispositions: (1) the dispositions (if any) that are constitutive of being able to think of F at all, (2) the dispositions that individuate the perspectives on F that associate particular weak specifications of F , and (3) those various and sundry dispositions that are irrelevant to (1) and (2). Of course, any theory of the cognitive mind that explicates conceptual states, acts, and capacities in dispositional terms—which is to say, any of a large variety of theories of the cognitive mind—bears a similar explanatory burden.

²⁴ I am abstracting away from any context dependence manifested by the constituents of the sentence, such as temporal indexicality in the tense of 'is' or quantifier domain restriction in the noun phrase 'the glass'.

yellow mentioned in those contents. We can ascribe the belief that the ‘The pane of glass is yellow’ is true if and only if the pane of glass is yellow to either of the speakers in Case 2 without making any presupposition about the particular inferential dispositions with the concept YELLOW that individuate the speaker’s perspective on *yellow* (hence without identifying the weak specification associated with the classification the speaker is making with ‘yellow’).

Moreover, grasping a speaker’s perspective is irrelevant to comprehending her utterance. When I hear my interlocutor utter the sentence ‘I’d like a pane of yellow glass’, the grammar that I know for my language determines that I hear her as saying that she would like a pane of yellow glass—that is, saying something that is true if and only if she would like a pane of yellow glass. It may be that the truth of what she said *consists in* her wanting a pane of glass that is yellow “through and through.” That is a fact that is determined by her perspective, i.e., that she is thinking of the yellowness of the glass in such a way that this yellowness consists in instantiating the weak specification of being yellow through and through. This fact may be evident to me the hearer, or it may not. She may succeed in getting that fact across, or she may not. Regardless, there remains a stable piece of information about what my interlocutor strictly and literally said in making the utterance.

I have tried to reconcile the specificational approach with referential semantics for *language* by making a proposal about how we should understand the referential aspects of *thought*. My proposal is that the referential aspects of our thoughts about properties may reflect our perspectives: often, one’s thinking of a property *F* involves a particular conception of what it is to instantiate *F*, so that whether an object *fits* one’s thinking about *F* depends on whether it instantiates *F* in that particular way. If this is so, then as I suggested in Sect. 5.1, the question “Setting aside any and all perspectives, what is really included in the extension of *F*?” may have no answer. The specificational approach requires a conception of the nature of properties that allows us only a deflationary and trivial notion of a property’s extension.

6 Conclusion

Travis cases have attracted attention among philosophers and cognitive scientists because their interpretation seems to bear directly on how to understand a central problem in the study of language: the interface between semantics and pragmatics. The standard terms of the debate divide the competing views in roughly this way. Some attempt to explain away the intuitions that Travis cases appear to elicit, and deny that ordinary predicates are used to effect different classifications of objects. Others maintain that a Travis case exhibits context dependence in what is literally expressed by uses of the predicate (whether this is linguistically controlled, as grammatical contextualists maintain, or triggered by non-linguistic principles of interpretation, as extra-grammatical contextualists contend). These approaches presuppose that what is dependent upon context is an aspect of what is strictly and literally expressed in the use of the predicate. The specificational approach that I have proposed rejects this presupposition. A principal benefit of the approach is that some problems that seemed to force complications into semantic theory are now seen to lie outside its purview. The

problems that Travis cases raise in the philosophical analysis of linguistic (and mental²⁵) representation remain as puzzling as they have ever been; but if my approach is correct, they do not threaten the goal of semantic theory: to delineate a central part of the knowledge that underlies our understanding and use of language.

Acknowledgments Special thanks are due to Jacob Beck and to two anonymous reviewers for numerous invaluable suggestions for improvement. I also thank Richard Heck, Susanna Siegel, Bernhard Nickel, Michael Glanzberg, Michael Rescorla, Charles Parsons, Elisabeth Camp, Catherine Wearing, and audiences at Harvard University, MIT, UC San Diego, and the University of Western Ontario for helpful comments on earlier versions of this paper.

References

- Atlas, J. D. (1989). *Philosophy without ambiguity: A logico-linguistic essay*. Oxford: Clarendon Press.
- Bach, K. (1994). Conversational implicature. *Mind and Language*, 9, 124–162.
- Barwise, J., & Perry, J. (1983). *Situations and attitudes*. Cambridge, MA: MIT Press.
- Bezuidenhout, A. (2002). Truth-conditional pragmatics. *Philosophical Perspectives*, 16, 105–134.
- Cappelen, H., & Lepore, E. (2003). Context shifting arguments. *Philosophical Perspectives*, 17, 25–50.
- Cappelen, H., & Lepore, E. (2004). A tall tale in defense of semantic minimalism and speech act pluralism. In G. Preyer (Ed.), *Contextualism*. Oxford: Oxford University Press.
- Cappelen, H., & Lepore, E. (2005). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*. Oxford: Blackwell Publishers.
- Chomsky, N. (2000). *New horizons in the study of language and mind*. Cambridge: Cambridge University Press.
- Egan, A., Hawthorne, J., & Weatherson, B. (2005). Epistemic modals in context. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning and truth* (pp. 131–168). Oxford: Oxford University Press.
- Higginbotham, J. (1988). Contexts, models, and meanings: A note on the data of semantics. In R. Kempson (Ed.), *Mental representations: The interface between language and reality*. Cambridge: Cambridge University Press.
- Kennedy, C. (1999). *Projecting the adjective: The syntax and semantics of gradability and comparison*. New York: Garland Press.
- Laserson, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28, 643–686.
- MacFarlane, J. (2003). Future contingents and relative truth. *Philosophical Quarterly*, 53, 321–336.
- MacFarlane, J. (2009). Nonindexical contextualism. *Synthese*, 166, 231–250.
- Moravcsik, J. (1998). *Meaning, creativity, and the partial inscrutability of the human mind*. Stanford: CSLI Publications.
- Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
- Reimer, M. (2002). Do adjectives conform to compositionality?. *Philosophical Perspectives*, 16, 183–198.
- Sanford, D. (2006). Determinates vs. determinables. *Stanford Encyclopedia of Philosophy*. Retrieved June 28, 2010, from <http://plato.stanford.edu/entries/determinate-determinables/>
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Cambridge, MA: Harvard University Press.
- Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23, 391–434.
- Szabó, Z. (2001). Adjectives in context. In I. Kenesei & R. M. Harnish, *Perspectives on semantics, pragmatics, and discourse: A festschrift for Ferenc Kiefer*. Amsterdam: John Benjamins.
- Travis, C. (1981). *The true and the false: The domain of the pragmatic*. Amsterdam: John Benjamins.
- Travis, C. (1985). On what is strictly speaking true. *Canadian Journal of Philosophy*, 15, 187–229.
- Travis, C. (1989). *The uses of sense: Wittgenstein's philosophy of language*. Oxford: Clarendon Press.

²⁵ It is implicit in the argument of this paper, and especially in the discussion of perspectives on properties as cognitive dispositions, that the kind of context-dependence arising in Travis cases is equally manifest in mental representation. I discuss the issue further in other work.

- Travis, C. (1990). Critical notice of Sperber and Wilson, *Relevance: Communication and cognition*. *Canadian Journal of Philosophy*, 20, 277–304.
- Travis, C. (1997). Pragmatics. In B. Hale & C. Wright (Eds.), *A Companion to the philosophy of language*. Oxford: Blackwell Publishers.
- Travis, C. (2000). *Unshadowed thought: Representation in thought and language*. Cambridge, MA: Harvard University Press.