Just What Is It That Makes Travis's Examples So Different, So Appealing?*

Nat Hansen

Forthcoming in J. Collins, A. Davies, and T. Dobler (eds.), *Themes from Charles Travis: On Language, Thought and Perception*, Oxford: Oxford University Press.

Abstract

Odd and memorable examples are a distinctive feature of Charles Travis's work: cases involving squash balls, soot-covered kettles, walls that emit poison gas, faces turning puce, ties made of freshly cooked linguine, and people grunting when punched in the solar plexus all figure in his arguments. One of Travis's examples, involving a pair of situations in which the leaves of a Japanese maple tree are painted green, has even spawned its own literature consisting of attempts to explain the context sensitivity of color adjectives ("green", e.g.). For Travis, these examples play a central role in his arguments for occasion-sensitivity, which he takes to be a pervasive feature of how we understand natural language. But how, exactly, do these examples work? My aims in this paper are to put Travis's examples under the microscope, using recent experimental studies of Travis-style cases to raise worries about aspects of the way Travis's cases are informally presented, but then show how his examples can be redesigned to respond to these doubts.

1 Introduction

Since the 1970s, Charles Travis has been waging a guerrilla war against mainstream truth conditional theories of meaning. In the last decade or so, Travis has been caught up in skirmishes between minimalists, moderate and radical contextualists, indexical and non-indexical contextualists, relativists, and other factions fighting over the proper way to understand the interaction of meaning and context. Some of Travis's examples have played a prominent role in these debates. For example, a squadron of theories have formed in

*Thanks to Zed Adams, Alex Davies, Chauncey Maher and Eliot Michaelson for comments.

response to Travis's examples that involve color adjectives ("green", "red", etc.) (Clapp, 2012; Hansen, 2011; Kennedy and McNally, 2010; Predelli, 2005; Rothschild and Segal, 2009; Sainsbury, 2001; Szabó, 2000; Vicente, 2012).

Travis's examples are, to borrow a Wittgensteinian phrase (and to switch metaphors), the "decisive movement in the conjuring trick". Because they generate the empirical foundation on which wide-ranging debates about context and meaning rest, they should be closely scrutinized. My aims in this paper are to put Travis's examples under the microscope, using recent experimental studies of Travis-style cases to raise worries about aspects of the way Travis's cases are informally presented, but then show how his examples can be redesigned to respond to these doubts. .

2 Occasion-Sensitivity, Reasonable Judges, and Evidence of Semantic Facts

Travis refers to his view of meaning and context using different names in various places: "occasion-sensitivity" (Travis, 2008c, p. 4), "S-use sensitivity" (Travis, 1989, p. 18), "the speaking-sensitive view of words" (Travis, 2008a, p. 70), or "the pragmatic view" (Travis, 2008f, p. 109). In outline, Travis's view of how meaning and context interact is the following:[1]

- English (e.g.) words have occasion-invariant meaning.
- The meaning of words does not fix their extension; it merely "constrains" what those words can be used to talk about (or what the truth condition of a sentence is, or what a sentence says).
- What a sentence says or what an expression refers to on a particular occasion is fixed by a correct understanding of the use of that expression on an occasion.
- A correct understanding of the use of an expression on an occasion is how a "reasonable judge" would understand it. Reasonableness, and hence also correct understand-

[1] Gustafsson (2002) gives a useful introduction to Travis's conception of meaning and context.

ing, is not equivalent to an appreciation of some finite set of contextual parameters.

A representative example illustrates how the parts of Travis's view fit together:

> Hugo and Odile are discussing their new apartment. Hugo says, 'The walls are beige'. The walls are, indeed, painted beige, though made of white plaster. Are Hugo's words true? There are various ways of counting walls as beige or not. On some of these ways, the walls do so count, and on others they do not. Correspondingly, it is sometimes true to say that the walls are beige, and sometimes true to say that they are not. So there are different things to be said, on different occasions, in saying so... Nothing said so far decides which of these things Hugo said, nor, hence, what is required for his words to be true—or at least enough about this to determine whether they are true (Travis, 1989, p. 122).

At this preliminary stage, given a meaningful sentence ("The walls are beige") and a situation that the sentence can be used to describe (walls that are painted beige but made of white plaster), it is not yet fixed (i) what Hugo says (if anything) by using the sentence, or (ii) whether what he says (if anything) is true or false of the walls. To get to a stage where what Hugo said by saying "The walls are beige" is fixed and whether what he said is true, aspects of an occasion of use—the "overall activities of which Hugo's words are a part"—must be added:

> But now suppose we consider the overall activities of which Hugo's words were a part. Suppose that Hugo and Odile are choosing a rug. For that purpose, it would not be reasonable to take 'beige', used in the way Hugo did, to be subject to standards on which it would not be true of those walls, or to use it so that it was subject to such standards... So it is most reasonable to take it to

> have been subject to standards on which it is true of the walls. So Hugo said what is true…

If the "overall activities of which Hugo's words are a part" change, it may become reasonable to take what Hugo said to be *false* of the walls. So on Travis's view the meaning of a sentence underdetermines (does not "fix") content, truth conditions, or "what is said". Only a sentence used on some occasion and understood as a reasonable judge would understand it says something that can be true or false.

Occasion-sensitivity involves both a view of "descriptive semantics" and a view of "foundational semantics" (or meta-semantics). Descriptive semantics "says what the semantics for the language is, without saying what it is about the practice of using that language that explains why that semantics is the right one" (Stalnaker, 1997, p. 535). The descriptive component of occasion-sensitivity is its rejection of certain views about the way that context and linguistic meaning interact to yield the content of what is said: namely, that ambiguity or ellipsis or hidden indexicality can adequately capture all of the types of variation in truth conditional content that words are subject to. What given words, used on an occasion, say is what a reasonable judge would take them to say on that occasion.

The foundational or meta-semantic component of occasion-sensitivity is the view that what constitutes or grounds the semantics that given words have are the ways we (when we are acting as reasonable judges) react to uses of those words:

> [W]e and our reactions—in the first instance, to particular candidate semantic items—are at the core of the present story of the making of semantic fact. By 'we', I mean you and I and most of our friends, in the first instance, and any one else with a suitable competence in the language in question (where one is) and an appropriate background in the activities in which the semantic item in question is being or is to be employed…Not all English speakers are fit for understanding remarks about dace and bream, for example (Travis, 1989,

pp. 45–46).

The remark about dace and bream indicates that whether one counts as a reasonable judge may vary depending on the topic of conversation. Since I don't know much about fish, I am not a reasonable judge regarding the meaning of "dace", and whether or not I call a Chinese mud carp a "dace" does not play a role in determining whether or not what I say is true. But the judgments of icythyologists and well-informed fishermen would play that role. And some semantic items are such that no *special* expertise (beyond being a reasonably worldly English speaker) is required to be a reasonable judge with respect to them: "knows", "beige", "weighs 80 kg", for example, would be such semantic items.[2]

Given that the reactions of reasonable judges constitute semantic facts, if we know who the reasonable judges are, then we can observe their reactions as a way of figuring out what the semantic facts are. And "often enough", we *do* know who the reasonable judges are:

> Semantic facts are detectable, the present thought goes, if there are reliable detectors of them, and if these detectors are often enough identifiable. But there are. We detect or discern the semantic facts often enough by relying on ourselves as detectors of them, in the first instance; or more generally, by relying on those who are recognizably reacting to the relevant semantic items (on relevant confrontations with them) as a reasonable judge would then react to them. . . Often enough, there is no doubt as to who these judges are, and often enough no doubt that we are among them (Travis, 1989, pp. 62–63).

In short, when we know how reasonable judges would assess some semantic fact (e.g. the truth of what is said by someone in using a sentence on an occasion), then we thereby

[2]Travis's view about the role of reasonable judges in the constitution of semantic facts might usefully be compared to the "judgment-dependent" account of meaning proposed in Wright (2001). A central difference between the two views is that Travis does not try, or seem to think it possible, to give an informative, non-circular account of what Wright calls "C-conditions"—those conditions that have to obtain for a judgment to be extension-determining. In Travis's terms, there is no informative, non-circular specification of the conditions that make someone a reasonable judge. For some remarks comparing Travis's occasion-sensitive view and Wright's judgment-dependent view, see Miller (2009).

know those semantic facts. That is the background to Travis's examples: assuming that we are reasonable judges, then our reactions to his examples tell us what the relevant semantic facts are. And, the thought goes, once we have observed our reactions we learn that whether a given use of a sentence says something true or false depends on all kinds of features of the occasion of use that aren't obviously explicable in terms of ambiguity or ellipsis or hidden indexicality—that is, that the sentence's semantics is occasion-sensitive.

3 Reasonable Judges and Experiments

How should we use reasonable judges (that is, usually, ourselves) as detectors of semantic facts? Given a concept C (*being beige*, e.g.), "[o]ne might distinguish cases of an item's fitting C from others by… appealing directly to intuitions as to what does fit C (or so counts on an occasion)" (Travis, 1989, p. 62).[3] In other words, we can discover facts about the concept *being beige* or the expression "is beige" by asking ourselves whether various items are beige or not on various occasions.

Travis allows that we might discover that certain expressions have occasion-*insensitive* semantics, if we can't find any occasions on which we react differently to the expression. This seems to be the case with mathematical expressions: Travis (p. 28) asks, "could the number 3 be in a position to exhibit S-use [occasion] sensitivity in the property of being odd?", and Szabó (2001, p. 124) offers what I think is the right answer to a question in the neighborhood of Travis's question: "If one uses the sentence 'The number is even', talking about the number four, the sentence expresses a truth. One cannot construct some special scenario where… 'The number is even' says something false".[4] But for expressions

[3] I understand "intuition" as Travis uses it here to just be a *judgment*, with no commitments to whether it is spontaneous or "reflective". For discussion of different ways of understanding "intuition", see Nagel (2012).

[4] In the context wars, Travis is usually classified as a "radical contextualist". Radical contextualism is sometimes distinguished from "moderate" contextualism in terms of the *scope* of context sensitivity: moderate contextualists hold that *some* natural language expressions are context sensitive, while radical contextualists are said to think that *all* natural language expressions are context sensitive (Cappelen and Lepore, 2005, pp. 5–6), (Szabó, 2006, p. 32). And Travis frequently says that he is committed to the generality of occasion-

that concern "sublunary" (non-mathematical) affairs, Travis thinks there is good evidence, supplied by our reactions, indicating that they will have occasion-sensitive semantics.

Whether they indicate that the semantic facts vary on different occasions or that they remain the same, our reactions are *data* that are generated by Travis's examples, which are informal *experiments*. They function by varying an independent variable, namely the occasion on which some sentence is used, and observing the resulting variation (or lack thereof) in a dependent variable, namely our reactions. A schematic representation makes this structure—shared by many of his examples—clear:

Table 1: The Structure of Travis's Examples

	Occasion 1	Occasion 2
SENTENCE	REACTION	REACTION

Travis's example of the beige walls fits the structure laid out in Table 1 nicely. The sentence that is used on both occasions is "The walls are beige", as spoken by Hugo, and the two occasions on which we're supposed to evaluate uses of the sentence both involve walls that are made of white plaster but have been painted beige. On occasion 1, Hugo and Odile are choosing a rug, and on occasion 2 the superintendent of Hugo and Odile's building is trying to determine whether the walls in their apartment need to be demolished. They will be demolished if they emit poison gas, and only walls made of beige plaster emit the gas. The superintendent asks whether they have beige walls, and Hugo, wanting new walls, says "The walls are beige" (Travis, 1989, p. 122).[5]

sensitivity. For example, he says "if I am right, occasion-sensitivity is everywhere" (Travis, 2008c, p. 1) and that "any English sentence" (Travis, 2008b, p. 235), "any English predicate" (Travis, 2008d, p. 276), and "predicates in general" (Travis, 2008g, p. 207) can be used to illustrate occasion-sensitivity. The fact that Travis allows that certain (mathematical) expressions could be occasion-insensitive might seem to be in tension with this characterization of radical contextualism. But the distinction between the meta-semantic and descriptive semantic commitments of occasion-sensitivity help resolve the tension; Travis can hold that the semantic facts are constituted by the reactions of reasonable judges (meta-semantic occasion-sensitivity), even when reactions don't vary (yielding semantic occasion-insensitivity at the descriptive level).

[5]In Travis's original, Hugo actually says "They're beige", referring to the walls on the second occasion. But, as I will discuss in more detail later, varying the linguistic material in addition to features of the sur-

Travis, in the guise of a reasonable judge, offers his reactions to this example as indicated in Table 2 (p. 122). Assuming Travis's reactions are those of a reasonable judge, then

Table 2: Beige Walls

	Choosing a rug	*Finding walls that emit poison gas*
"The walls are beige"	HUGO SAID WHAT IS TRUE	WHAT HUGO SAID [IS] FALSE

on Travis's picture of how we have knowledge of semantic facts we know that what Hugo said varies its truth value across the two occasions.

Many of the partisans in the debates that have developed in response to Travis's examples take it for granted that the reactions that Travis describes to his examples are indeed tracking semantic facts but offer an explanation of those reactions using the resources not of occasion-sensitive semantics, but of traditional truth conditional semantic theory (Hansen, 2011; Kennedy and McNally, 2010; Predelli, 2005; Szabó, 2001).

A different way of challenging the descriptive component of occasion sensitivity is to target Travis's claim that we *know* the relevant semantic (truth-involving) facts, by raising doubts about whether our reactions are those of a reasonable judge or not. Here is one example of how such a challenge can be made.[6] J.L. Austin (1966, p. 432) aims to demonstrate the subtly different meanings of the words "deliberately" and "intentionally" by describing a situation in which a go-cart is run over and then asking whether we should say that the go-cart was run over intentionally or deliberately. Here's the situation he describes:

> I am summoned to quell a riot in India. Speed is imperative. My mind runs on the action to be taken five miles down the road at the Residency. As I set off

rounding occasion makes it more difficult to identify what is responsible for the variation in reactions, so I have used the same sentence on both occasions.

[6]This discussion is closely based on Hansen and Chemla (2013a).

> down the drive, my cookboy's child's new gocart, the apple of her eye, is right across the road. I realize I could stop, get out, and move it, but to hell with that: I must push on. It's too bad, that's all: I drive right over it and am on my way. In this case, a snap decision is taken on what is essentially an *incidental* matter.

Austin then immediately follows up the description of the situation with an unambiguous statement of what we should say about it—a statement that is endorsed by Ferguson (2003, p. 93), Searle (2001, p. 223), and Williams (2009, p. 24). And I think many more philosophers would be inclined to endorse Austin's judgment about the case, given his continuing reputation for having an acute ear for making subtle distinctions in meaning that other philosophers miss. Now consider the following two ways of glossing the situation Austin describes. One is Austin's gloss (slightly abbreviated to make a tighter parallel possible), the other is simply the reverse of his gloss. Which one is more reasonable?

Gloss A: In this case, a snap decision is taken on what is essentially an incidental matter. I did drive over the gocart intentionally, but not deliberately... At no time did I deliberate whether to drive over it. It was incidental to anything I deliberated about doing, which was simply to get to the scene of the riot in order to quell it. However 'odd' it may sound, I feel little doubt that we should say here that we did run over the go-cart intentionally *and* that we should not care to say we ran over it deliberately. We never deliberated whether to run over it.

Gloss B: In this case, a snap decision is taken on what is essentially an incidental matter. I did drive over the gocart deliberately, but not intentionally... At no time did I intend to drive over it. It was incidental to anything I intended to do, which was simply to get to the scene of the riot in order to quell it. However 'odd' it may sound, I feel little doubt that we should say here that we did run over the go-cart deliberately *and* that we should not care to say we ran over it intentionally. We never intended to run over it.

When both possible glosses are available, I submit that it appears a lot less obvious what the "reasonable" way of understanding the situation is. (Often I have to look back at the Austin even to remember which way he "felt little doubt" was the right way to describe the situation—the reader may want to look at the relevant passage in "A Plea for Excuses"

to be reminded which description is Austin's.)

In Hansen and Chemla (2013a) we presented versions of Austin's "go-cart" example to subjects who saw either Gloss A, Gloss B, or just the description of the situation with no gloss.[7] We found (unsurprisingly) that subjects' reactions could be reversed by reversing the gloss, and (more surprisingly) that in the neutral, gloss-free condition, their reactions did not align with Austin's gloss of the situation.

Where does this leave us with regard to what the reasonable reaction to the go-cart scenario is? I think it means that we don't know what the reasonable response to it is (or even that there is one unique reasonable response). Given Travis's view of our access to the semantic facts, that means that we don't know what the fact of the matter is regarding the correct (or, choosing between "intentionally" and "deliberately" as descriptions of how we ran over the go-cart, even the *better*) description of the go-cart story. That doesn't mean that there *isn't* a fact of the matter about which words correctly describe the story, or even that it is more reasonable to use one over the other. It only means that we're not currently in a position to know which is more reasonable, and therefore, not in a position to know what the semantic fact of the matter is regarding whether "intentionally" or "deliberately" applies to the go-cart story.

In the following section, I will argue that several features of the design of Travis's examples give us reason to doubt that we know what the reasonable reaction to them would be. Four features in particular might spark such doubts about our reactions to Travis's examples:

1. experimenter effects

2. order of presentation effects

[7]The versions of Gloss A and Gloss B that appeared in our experiment were slightly different than those described here; they were switched from the first person to the third person, and the language was slightly altered but kept parallel as much as possible.

3. contrast effects

4. effects of the "rule of accommodation"

In addition to these features, there are miscellaneous linguistic asymmetries in some of Travis's examples that pose problems for what Travis calls the "method of contrasting pairs", a method which his examples are intended to be demonstrations of, even if we assume that our reactions are the reactions of reasonable judges.

4 Sources of Skepticism about Reactions to Travis's Examples

4.1 Experimenter effects

Suppose that our judgments about Travis's examples could be *reversed* if the substance of the example (the sentence used and the details of the occasion of use) were left exactly the same but the accompanying 'gloss' on how to respond to the example were changed. If that kind of reversal is possible, it wouldn't show that our reactions when prompted by Travis's gloss are *not* reasonable. But it should give us reason to *wonder* what the reasonable reaction to the example is. Is it the reaction we have when it is glossed one way, but not the other? Is the way an example is glossed something that the reasonable judge should take into consideration? I myself am inclined to think that the optimal condition for reacting to a case is when it is free of any accompanying gloss, or presented with counterbalanced glosses (see Hansen 2013 for such a proposal). Ideally, we should compare how we react to the examples when they are presented without any gloss and when they are glossed in the way Travis describes. If it turned out that our reactions don't vary across the two conditions (glossed and gloss-free) that would address the worry that the glosses are distorting our reactions to the examples. But if our reactions do vary across the two conditions, then without some convincing reasons to prefer the glossed condition, we don't know what the reasonable reaction to the examples would be.

Consider Travis's "milk in the refrigerator" example, first stripped of any accompanying gloss, then with two different glosses:

> Suppose that the refrigerator is devoid of milk except for a puddle of milk at the bottom of it. Now consider two possible speakings, by Odile, of the words, 'There's milk in the refrigerator'. For the first, Hugo is seated at the breakfast table, reading the paper, and from time to time looking dejectedly (but meaningfully) at his cup of black coffee, which he is idly stirring with a spoon. Odile volunteers, 'There's milk in the refrigerator'. For the second, Hugo has been given the task of cleaning the refrigerator. He has just changed out of his house-cleaning garb, and is settling with satisfaction into his arm-chair, book and beverage in hand. Odile opens the refrigerator, looks in, closes it and sternly utters the above words ['There's milk in the refrigerator'] (Travis, 1989, pp. 18–19).

Here are two ways of glossing that example:

Reaction A: "There is milk in the refrigerator", as used on both occasions, says what is true.

Reaction B: "There is milk in the refrigerator", as used in the first case, says what is false, but as used in the second case, says what is true.

Reaction B is Travis's reaction.[8] Here is the worry about the effect of glosses on our reaction to Travis's examples: When the milk in the refrigerator example is accompanied *only* by Travis's reaction, it might seem much more plausible than when it is juxtaposed with Reaction A, or when the example is presented by itself, without any accompanying

[8] When Travis introduces the milk in the refrigerator example, he says "I will not argue here that the above description of those speakings is correct. The present point, after all, is not to prove the dominant picture [of semantics] wrong, but to identify something which would be at odds with it" (Travis, 1989, p. 19). But elsewhere, Travis does take this kind of reaction to be evidence that the semantics of English, or of "predicates as such" (Travis, 1978, p. 424), is occasion-sensitive.

gloss. This worry gets a further grip when combined with demonstrations of how experimenters can—even unconsciously—affect the reactions of their subjects by indicating what the expected reaction to a given task is.

For example, Doyen et al. (2012) challenge a widely-cited finding in social psychology that seems to indicate that priming subjects with words related to the concept of old age slows down the speed at which they walk down a corridor during the experiment (relative to those who haven't been so primed) (Bargh et al., 1996). After they failed to replicate the original priming effect, Doyen et al. manipulated the experimenters' expectations about whether participants would walk slowly or quickly, and found that only when experimenters expected primed participants to walk slowly did the participants walk significantly slower than those who had not been primed with words associated with old age. When the experimenters expected participants to walk faster when primed with words associated with old age, there was no significant difference between those who had been primed and those who had not been primed. That is evidence that experimenters' expectations play a key role in generating the supposed "priming" effect.

Intons-Peterson (1983) showed that subjects' responses to tasks involving perception and imagination were affected by manipulating experimenters' beliefs about what outcomes to expect. In one experiment, for example, participants were instructed to "indicate whether an outline of a hand was that of a right or a left hand when viewed from the back. . . a priming hand, in an upright position, was either imagined or seen by the experimental groups" (p. 404). Experimenters were "led to believe that decisions about the hands would be faster with imaginal primes than with perceptual primes or they were led to believe the converse". Decision times were found to mirror experimenters' beliefs, resulting in faster decisions for imaginal primes than for perceptual primes when that's what the experimenter expected to happen, and vice versa.

Finally, the well-known case of "Clever Hans" is the classic demonstration of an experi-

menter effect. "Sober and sedate Germany was thrown into a turmoil of newspaper debate" (Pfungst, 1911, p. v) by the fact that it seemed that Clever Hans, a horse, was capable of answering complex arithmetical problems, indicating the correct answer to problems set by questioners by tapping his hoof. But it was eventually discovered that Hans was actually responding to unconscious cues from his questioners, who were inadvertently communicating the answers to Hans. Thus, "All wonderful feats of counting and computation which were accomplished while thus experimenting with the horse are to be accredited, not to the horse, but to the questioner" (p. 141).

Given those striking demonstrations of how subtle forms of experimenter expectations can affect responses in a variety of situations, we should question whether Travis's account of the truth values of what is said in his examples might be influencing our own reactions. It is of course also possible that Travis's descriptions point us in the direction of the reasonable response to the examples. But it seems plausible that we should be able to react reasonably to the examples without any accompanying gloss, because whether it is reasonable to judge that what a speaker says is true or not should not depend on whether Travis (or anyone else) *says* that such a judgment is reasonable. If our reactions to Travis's examples when not prompted by his gloss differ substantially from our reactions when prompted by his gloss, then it's not obvious whether it is our reactions to his examples as standardly presented or our reactions to his examples when not prompted by any gloss that are the reactions of a reasonable judge. For that reason, it's important to compare reactions to Travis's examples in both the original, glossed form and also when the examples are presented gloss-free. In §5 I discuss the results of an experiment that presented the examples without any gloss, which we can compare with our reactions to Travis's glossed examples.

4.2 Order of presentation effects

Philosophers might dismiss the idea that they are susceptible to the same kinds of effects as experimental subjects, but Schwitzgebel and Cushman (2012) found evidence that not only were professional philosophers' reactions to examples in moral philosophy affected by the *order* in which the examples were presented, but that philosophers were subject to an additional form of order of presentation bias that ordinary subjects were not: their responses to questions about general moral principles were also significantly affected by the order in which they viewed certain scenarios. For example, whether philosophers endorsed a statement of the doctrine of double effect was significantly affected by whether they saw an example ("push") that "involved killing one person through direct physical contact as a means of saving five people" first or an example ("switch") that "involved one person's dying, without direct physical contact from the agent, as a side effect of an action to save five" first (p. 138).[9]

Table 3: An Example of Order of Presentation Effects on Philosophers

	% Endorsing Doctrine of the Double Effect
"Push" first	46%
"Switch" first	62%

These findings raise an analogous concern about potential order of presentation effects in Travis's examples. For example, it may be the case that once one encounters an occasion that has relatively demanding standards for the application of some concept (in which, say, being painted green isn't sufficient for something to count as *green*—it has to be naturally green), it is more difficult to feel the appeal of the relatively less demanding standards that

[9]The general moral principles were phrased in the form of questions, like the following: "Sometimes you can save several people by actively and purposefully killing one person whom you could have let live. Other times you can save several people by purposefully allowing one person to die whom you could have saved. Is the first action morally better, worse, or the same as the second action?" (p.139). Responding with "worse" constituted an endorsement of the doctrine of double effect for the purposes of the experiment.

are operative on a second occasion (in which either being painted green or being naturally green would count as being green) than if the order in which the occasions were encountered were reversed. If our reactions to Travis's examples are affected simply by changing the order in which the contrasting cases are presented, we would have an additional reason to wonder whether our reactions are reasonable, or whether we are being pushed around by extraneous factors that a reasonable judge should not be affected by.[10]

4.3 Contrast Effects

Travis sometimes refers to his examples as involving "contrasting pairs" (see, e.g. Travis 1989, p. 204). He spells out the "method of contrasting pairs" as follows:

> Beginning with words W, look for some words W^*, which share as many already recognized semantic features with W as possible, but still differ intuitively from W in their proper understanding. Try to see how like W words can be while still recognizably different in their semantics. Success may lead to recognition of new aspects of what is said in words, new semantic features or properties which may constitute, in part, a proper understanding. (Travis, 2008a, p. 91)

The method of contrasting pairs is an experimental method, which varies selected features of the occasion of use as a method of uncovering previously overlooked semantic features of words. There are two different approaches to applying the method of contrasting pairs. One approach introduces a potentially significant experiment-specific feature that is not present (or not present to the same degree) in the use of words on ordinary (non-experimental) occasions.

[10] One might argue that order is something that a reasonable judge should be sensitive to in making truth value judgments about Travis's examples, but I can't imagine what reasons one could offer in favor of that idea.

Experimenters distinguish *between–* and *within–subject* experimental designs. In a between-subjects design, different groups of subjects see and respond to different experimental conditions, and are not allowed to see and respond to contrasting conditions. So, for example, a between-subjects experiment employing one of Travis's examples would ask different groups of subjects to respond to one or the other of the two contrasting occasions of use for the term "beige" (but not both), and then compare the reactions of the two groups (see Table 4).

Table 4: Beige Walls: Between-subjects design

	Occasion 1: Rug	*Occasion 2: Gas*
"The walls are beige"	REACTION FROM GROUP A	REACTION FROM GROUP B

In a within-subjects design, a single group of subjects sees and responds to contrasting conditions. This is the standard set-up when we consider our reactions to Travis's examples in the informal context of reading his papers and books: we, the experimental subjects, get to compare both occasions side-by side while gauging our reactions to his examples (see Table 5).

Table 5: Beige Walls: Within-subjects design

	Occasion 1: Rug	*Occasion 2: Gas*
"The walls are beige"	REACTION FROM GROUP A	REACTION FROM GROUP A

Which design is a better way of implementing the method of contrasting pairs?[11] It has been observed that experiments that employ a between-subjects design more closely approximate conditions in which we would react in ordinary circumstances, because, as Daniel Kahneman puts it, "We normally experience life in the between-subjects mode, in which contrasting alternatives are absent" (Kahneman, 2011, p. 354).[12] That is, nor-

[11] The discussion of the importance of contrast that follows is drawn from Hansen (2014).

[12] For other examples of the claim that everyday life resembles a between-subjects experiment, see Kahneman (2000, p. 682) and Shafir (1998, p. 72).

mally we don't get to compare different occasions side-by-side when reacting to uses of language. And certain effects of context on our reactions only arise in experiments employing a within-subjects design (and disappear in experiments with a between-subjects design) (see, e.g., Phelan 2013). So one might worry that the fact that there are intuitive contrasts in our reactions to Travis's examples when we read them with a de facto within-subjects format is an artifact of that particular experimental design. Our reactions might differ substantially if we only ever encountered individual occasions without the benefit of a contrasting case to consider.

On the other hand, there are reasons to think that rationality is served by considering more than a single case at a time. Again, Kahneman (2011, p. 361) gives a succinct statement of the idea (where "joint evaluation" refers to situations akin to experiments involving a within-subjects design and "separate evaluation" to experiments involving a between-subjects design):

> . . . rationality is generally served by broader and more comprehensive frames, and joint evaluation is obviously broader than single evaluation.

An illustration of how the consideration of more options can produce improved judgments is given in Kahneman and Tversky (1996) in reference to the "conjunction fallacy". The "conjunction fallacy" is the tendency of subjects, in certain conditions, to judge that $p\&q$ is more probable than p alone. For example, consider the following example:

> Linda is in her early thirties. She is single, outspoken, and very bright. As a student she majored in philosophy and was deeply concerned with issues of discrimination and social justice. Suppose there are 1,000 women who fit this description. How many of them are
>
> (a) high school teachers?
>
> (b) bank tellers? or

(c) bank tellers and active feminists? (Kahneman and Tversky, 1996, p. 587)

Kahneman and Tversky report that when participants were allowed to see all three options (a), (b) and (c), 64% conformed to the conjunction rule, which holds that that conjunctions cannot be more probable than either conjunct. But in an experiment when participants saw only either options (a) and (b) or (a) and (c), "the estimates for feminist bank tellers . . . were significantly higher than the estimates for bank tellers" (p. 587).

So the method of contrasting cases might be a way of making our reactions more reasonable than if we only ever considered one occasion of use at a time. And there is no reason to stop at considering only *pairs* of occasions of use. Travis himself says that "there is nothing special about the number two" in the method of contrasting pairs, and proposes that we can speak of "contrasting n-tuples" (Travis, 2008e, p. 28). But that raises a question about the effects of considering more contrasting occasions of use: if there is reason to think that considering pairs produces more reasonable reactions than considering single cases, does that mean that there is also reason to think that further increases to the number of contrasting cases lead to further improvements in our reactions? If so, then we should be constructing examples that go beyond the standard design of Travis's examples by including more than two contrasting occasions of use. In §5 I will discuss a way of redesigning Travis's examples that allows us to assess this possibility by comparing our initial reactions to Travis's examples with our reactions to them after we have considered several versions of the same example.

4.4 The "Rule of Accommodation"

Almost all of Travis's examples involve evaluating a single sentence as used on two different occasions (see Table 1). Keith DeRose (2005, 2011), in contrast, has argued in favor of using a different design, in which a positive sentence is evaluated on one occasion and a

negative sentence (the positive sentence's contradictory) is evaluated on another occasion (see Table 6).

Table 6: The Structure of DeRose's Examples

	Occasion 1 ("low standards")	Occasion 2 ("high standards")
POSITIVE SENTENCE	REACTION	
NEGATIVE SENTENCE		REACTION

DeRose justifies his alternative design, which he uses to evaluate the context sensitivity of the word "knows", by invoking Lewis's (1979) "rule of accommodation", which according to DeRose is a general principle of interpretation that puts pressure on us as interpreters to find what a speaker asserts true. Assuming that the rule of accommodation is guiding our interpretation of what the speaker in DeRose's well-known "bank" case asserts on both the "low" and "high" standards occasions, then we should find both assertions true (see Table 7).

Table 7: DeRose's Reactions to the "Bank" case

	"low standards"	"high standards"
"I know the bank will be open on Saturday"	TRUE	
"I don't know the bank will be open on Saturday"		TRUE

The fact that we are inclined to judge that contradictory sentences are true in high and low standards contexts motivates DeRose's theory about the context sensitivity of "know".

In contrast, if the standard design (usually employed by Travis) was used to evaluate the context sensitivity of "know" using DeRose's "low" and "high" standards contexts, then DeRose predicts that it would fail to uncover evidence of context sensitivity, since the rule of accommodation would put pressure on us to find both assertions of "I know the bank is open on Saturday" *true* (see Table 8).

Table 8: The Rule of Accommodation and a Modified "Bank" Case

	"low standards"	"high standards"
"I know the bank will be open on Saturday"	TRUE	TRUE

DeRose is providing a reason to think that speakers will be inclined to judge two uses of the same sentence in different contexts both *true*, thereby obscuring any potential occasion-sensitive effects. And several recent experimental investigations of the context sensitivity of "know" that used the traditional design (evaluating a single positive sentence in two different contexts) have found roughly what DeRose predicted: subjects are inclined to find what the speaker asserts when he says "I know the bank will be open on Saturday" *true* in both the low and high standards contexts.[13] That is, these studies failed to find evidence of the effects of context on reactions to uses of the sentence. DeRose criticized these studies for neglecting his recommended design and argued that their failure to find evidence of context sensitivity does not constitute a challenge to the context sensitivity of "know".[14] If DeRose is right, then his design will provide better evidence of the effects of the occasion of use on what is said by the use of a sentence than Travis's standard design.

4.5 Linguistic asymmetries

The "method of contrasting pairs", as Travis spells it out, aims to uncover previously unrecognized semantic properties of some words W by finding some words W^* that "share as many already recognized semantic features with W as possible", while still differing in how they should be "properly understood". The method is most interesting when words W and W^* are *exactly the same*, and yet varying the surrounding occasion of use produces surprising changes in how we understand those words. The more that W and W^* differ, the

[13]For DeRose's discussion and response to these experimental studies, see DeRose (2011). I discuss DeRose's reasons for preferring his design in Hansen (2012).

[14]In Hansen and Chemla (2013b), we criticize these studies for a different reason: they take the absence of evidence of variation to be evidence of absence of variation.

less reason there is to think that differences in our reactions to uses of those words is due to anything new or surprising, rather than due to differences in the conventional meaning of the words involved.

Some of Travis's best examples elicit different reactions to the use of the very same sentence on two different occasions. The beige wall and milk in the refrigerator examples have this structure, for example. Other examples diverge from this ideal in important ways. Consider his most-discussed example, the case of the painted leaves (presented in Table 9):

> A story. Pia's Japanese maple is full of russet leaves. Believing that green is the color of leaves, she paints them. Returning, she reports, 'That's better. The leaves are green now'. She speaks truth. A botanist friend then phones, seeking green leaves for a study of green leaf chemistry. 'The leaves (on my tree) are green', Pia says. 'You can have those'. But now Pia speaks falsehood (Travis, 2008f, p. 111).

Table 9: The Painted Leaves Case

	Occasion 1	Occasion 2
"The leaves are green now"	PIA SPEAKS TRUTH	
"The leaves (on my tree) are green"		PIA SPEAKS FALSEHOOD

This example diverges from the ideal of using the same words on both occasions in one important way. Pia says "The leaves are green *now*" on the first occasion, whereas on the second occasion the temporal indexical "now" does not appear in what she says. That poses a problem for Travis's argument in the paper where the example appears, in which he is trying to establish the viability of "The Pragmatic View", which is the idea that "semantics [the meaning of words] has little or nothing to do with truth conditions" (Travis, 2008f, p. 109). Travis's strategy for making the Pragmatic View plausible involves showing that a sentence with a fixed semantics [linguistic meaning] can have different truth conditions

on different occasions of use. But because he varies the wording of the sentence being used on the two occasions, it isn't the case that the example (even prima facie) shows that a sentence with fixed linguistic meaning can vary its truth conditions.

A similar problem arises in Travis's discussion of knowledge ascriptions, when he describes two contrasting cases in which the topic of discussion is whether Hugo knows where the milk is (Travis, 1989, p. 156):

> Here is one simple thing that one might know: where the milk is, or that it is in the refrigerator. Here is one case. Hugo, engrossed in the paper, says, 'I need some milk for my coffee'. Odile replies, 'You know where the milk is'. Suddenly defensive, Hugo replies: 'Well, I don't really *know* that, do I? Perhaps the cat broke into the refrigerator, or there was just now a very stealthy milk thief, or it evaporated or suddenly congealed'...
>
> A contrasting case. Hugo and Odile are just leaving for a weekend in the country. As they are stepping into the car, Odile asks, 'Are you sure everything is put away?' Hugo replies, 'Yes'. She then asks, 'Where's the milk?' He answers, 'In the refrigerator'. She asks, 'Are you *sure* it is?' 'I know it', replies Hugo 'after all, that's where the milk is kept'.

However one reacts to these contrasting cases, there's a straightforward problem with drawing any semantic conclusions about "know" on the basis of one's reactions.[15] That's because again there are some important asymmetries in the sentences we are reacting to (see Table 10).

The first asymmetry is that the two sentences are of different *polarities*.[16] It is well known that sentences containing negations are harder to process, and create environments in which certain specific polarity-sensitive words can (in the case of negative polarity items)

[15]Travis's reaction (p. 156) to these cases is that both things that Hugo says are "absurd".

[16]Note that I have turned Hugo's rhetorical question into a declarative sentence.

Table 10: Knowing that the milk is in the refrigerator

	Occasion 1	Occasion 2
"I don't really *know* that [the milk is in the refrigerator]"	REACTION	
"I know [the milk is in the refrigerator]"		REACTION

or cannot (in the case of positive polarity items) occur. For example, the negative polarity item "anymore" can combine with the first sentence, but not the second:

(1) I don't really know that the milk is in the refrigerator anymore.

(2) * I know the milk is in the refrigerator anymore.[17]

Switching the polarity of the sentence between occasions of use therefore introduces a semantically significant change that makes it difficult to determine what factor (the changing polarity of the sentence or the changing occasion of use) is responsible for our changing reactions.[18] (There *is* a way that manipulating the polarity of the sentence can contribute to an improved experimental design, but it requires a more substantial revision to the standard design of Travis's contrasting cases; I will discuss the revision in §5)

The second important asymmetry is the presence of the adverb "really" in the first sentence, modifying "know". Cullen (2010, p. 287) observes that "the intensifying degree adverb *really* implicates a higher standard for knowledge than simply 'knows'". It would be easy to mistake an apparent contrast generated by the presence of of "really" in the first sentence (and its absence from the second) for an effect generated by features of the occasion of use. In order to effectively deploy Travis's method of contrasting pairs to uncover surprising features of the semantics of words, it is essential to eliminate these linguistic asymmetries from the contrasting cases under consideration.

[17]Labov (1975, p. 106) reports the existence of a Philadelphia dialect that employs positive "anymore", e.g., "John is smoking anymore"—meaning that whereas John previously did not smoke, he is now smoking.

[18]This is equally a problem for DeRose's preferred design, which is discussed in Hansen (2014) and Hansen and Chemla (2013b).

4.6 Summary

We need a way of addressing the questions raised above about *experimenter effects*, *order of presentation effects*, *contrast effects*, and effects of the *rule of accommodation*, which might make us wonder whether our reactions to Travis's examples as they are standardly presented are reasonable, or whether our reactions are being pushed around by inessential features of the design of his examples. And even if we are confident that we are reacting to the elements of his examples that are supposed to matter (namely the meaning of the words being used and their occasion of use), the presence of subtle linguistic asymmetries can undercut the effectiveness of the method of contrasting pairs.

In the following section, I'll describe how these worries can be addressed by redesigning Travis's examples, and describe evidence that Emmanuel Chemla and I collected using suitably redesigned versions of Travis's examples that supports the reactions that Travis describes.

5 Experimenting with Travis's Examples

Emmanuel Chemla and I (Hansen and Chemla, 2013b) experimented with a variety of Travis-style examples and confirmed that participants' reactions aligned with Travis's reactions in examples featuring color adjectives and in the "milk in the refrigerator" and "weighs 80 kg" examples, even after we controlled for experimenter, order, and contrast effects, and eliminated some problematic linguistic asymmetries.

To control for potential experimenter effects, we presented Travis's examples to our experimental subjects purged of his own statements of what the intuitive or "reasonable" response to them would be, and subjects read and responded to the examples remotely, so they would be shielded from our own expectations about how they would respond.[19] Order

[19]Subjects were recruited and reimbursed through Amazon's Mechanical Turk. For a discussion of the reliability of "MTurk" as a method of conducting linguistic experiments, see Sprouse (2011).

effects were controlled for by randomizing the order in which subjects saw the examples.

Controlling for contrast effects required a more sophisticated modification to Travis's original designs. Instead of simply switching to a between-subjects design, we developed a multiple "block" design that made it possible to collect reactions from subjects to a large number of Travis-style examples while isolating reactions to the *first* examples of each of several types that they saw (which varied between participants so that each type of example was seen first by some participants), thereby minimizing contrast effects in their reactions to those examples.[20] This design has the additional advantage that it allows for a comparison of how subjects respond to examples when they first encounter them with how they respond to similar examples later on in the experiment.

Addressing DeRose's argument regarding the role that the "rule of accommodation" plays in influencing responses to contextualist examples also required a substantial modification to existing experimental methodology. Instead of asking only for reactions to a single positive sentence as used on different occasions (as in Travis's more traditional design) or asking only for reactions to positive and negative sentences in "low" and "high" standards contexts (as in DeRose's modified design), we embedded both designs within a more comprehensive design that asked for reactions to *all four* possible combinations of sentence polarity (negative vs. positive) and context type ("high" and "low" standards in the case of knowledge ascriptions, "choosing a rug" and "searching for poison gas" in the case of the beige walls example, and so on) (see Table 11).

Table 11: Four-Cell Design

	Occasion 1	Occasion 2
Positive Polarity Sentence	REACTION	REACTION
Negative Polarity Sentence	REACTION	REACTION

In addition to allowing us to simultaneously evaluate both Travis-style and DeRose-

[20] See Hansen and Chemla 2013, p. 301 for a diagram of the block design.

style claims about how people react to uses of sentences on different occasions, the four-cell design enables two further improvements over existing ways of collecting data in support of contextualist or occasion-sensitive theories. First, by asking for reactions to both positive and negative sentences on both occasions of use, we can assess the relative contribution made by sentence polarity vs. the occasion of use to our reactions to Travis-style examples. The second, related advantage of the four-cell design is that we can evaluate the effect of the rule of accommodation. If we find that a sentence and its negation are both judged true (on the same occasion), that would be evidence in favor of DeRose's rule of accommodation.

We also eliminated whatever linguistic asymmetries we could when revising Travis's examples for use in our experiment. We broke his examples down into three component parts: (i) a description of the relevant situation that remains fixed between the two occasions; (ii) two contrasting occasions of use, which we tried to make as close in length to each other as possible; and (iii) a positive and a negative polarity sentence for evaluation on the two occasions of use. The following bracketed method of presentation highlights the three components of the painted leaves case (the example comes from Hansen and Chemla (2013b, p. 317)):

i. Pia has a Japanese maple tree in her backyard that has russet (reddish brown) leaves. She paints the leaves of the tree green.

ii.
- *Decorator*: A friend of Pia's who is making decorations for a play asks if Pia has any green leaves she can use in her stage set.
- *Chemistry*: A friend of Pia's who is conducting a study of green-leaf chemistry asks if Pia has any green leaves she can use in her study.

iii.
- *Positive*: 'Yes, the leaves on my tree are green,' Pia says.
- *Negative*: 'No, the leaves on my tree aren't green,' Pia says.

The four "cells" in our design are the result of combining the two elements in part ii with

the two elements in part iii. So the "positive–decorator" cell would look like this:

> Pia has a Japanese maple tree in her backyard that has russet (reddish brown) leaves. She paints the leaves of the tree green. A friend of Pia's who is making decorations for a play asks if Pia has any green leaves she can use in her stage set. 'Yes, the leaves on my tree are green,' Pia says.

And the "negative–chemistry" cell would look like this:

> Pia has a Japanese maple tree in her backyard that has russet (reddish brown) leaves. She paints the leaves of the tree green. A friend of Pia's who is conducting a study of green-leaf chemistry asks if Pia has any green leaves she can use in her study. 'No, the leaves on my tree aren't green,' Pia says.

5.1 Formal Experimental Evidence in Support of Travis's Reactions to His Examples

Our experiment with Travis's examples yielded several interesting results. First, we found that when we conducted a formal experiment designed to control for all of the sources of bias discussed above, there was a clear tendency for experimental subjects to align with Travis's reactions. Subjects responded to four different examples involving color adjectives (three of which were closely based on Travis's "beige walls" (Travis, 1989, p. 122), "painted leaves" (Travis, 2008f, p. 111) and "black kettle" (Travis, 2008e, p. 26) examples), Travis's "milk in the refrigerator" (Travis, 1989, p. 18–19) example, his "weighs 80 kg" (Travis, 2008g, p. 207) example, and four examples of knowledge attribution adapted from other philosophers.[21] In each type of example, changes in the occasion of use had statistically significant effects on subjects' reactions to the use of target sentences.

[21]The examples we tested were based on DeRose's (1992, 2009) bank example, Feltz and Zarpentine's (2010) truck example, Fantl and McGrath's (2002) train example, and Pinillos's (2012) spelling example.

Second, we found significant differences in the *strength* of the effects on subjects' reactions between different *types* of examples (p. 305). The color adjective, milk and weight examples revealed stronger effects of occasion on reactions than the knowledge scenarios did. This is one advantage of employing formal measures of subjects' reactions—it allows the detection of differences in strength that aren't usually apparent when considering informal truth value judgments.

Third, when we only considered reactions from the "first block" of examples, when subjects were seeing a particular example for the first time (without contrast), subjects' judgments still aligned with Travis's reactions across all of the color adjective examples, as well as the "milk in the refrigerator" and "weighs 80 kg" examples. But when we only considered reactions to the first block of examples, contextual effects on reactions to examples involving knowledge ascriptions (which were not based on Travis examples) disappeared. That is, it seems that contrast (or acclimation) is necessary for contextual effects on reactions to knowledge ascriptions to show up. And as discussed in the previous paragraph, when those effects do show up, they are weaker than the effects on the color, milk, and weight examples.

Fourth, we found evidence of a general bias in favor of finding positive sentences "more true" than negative sentences, but we did not find evidence of DeRose's "rule of accommodation", which should affect all sentences equally (p. 305).

Overall, our findings don't establish that Travis's reactions to his examples are indeed reasonable. But they indicate that the reactions to his examples that Travis describes can be reproduced without the effects of experimenter bias, order or contrast effects, or the "rule of accommodation". The *possibility* that those aspects of Travis's examples (as they are standardly presented) could be shaping our reactions was enough to raise doubts about whether our reactions were the reactions of reasonable judges, and hence doubts about what the semantic facts are. By showing that reactions to Travis's examples aligned with

his reactions even in conditions designed to control for those sources of bias, those specific worries about our status as reasonable judges can be dismissed.

Our findings also raise novel questions about the significance of reactions to the examples. The most important questions prompted by our experimental results concern the differences we observed between reactions to the knowledge cases and the examples from Travis we experimented with: *Why does changing the occasion of use only affect knowledge ascriptions when subjects have seen more than one knowledge example?* And *why is there a difference in strength between subjects' reactions to the knowledge cases vs. the color, milk and weight cases?* One possible explanation of these observations is simply that Travis can write short, memorable scenarios that provoke stronger reactions than cases imagined by other contextualists. (This possibility is discussed in Hansen and Chemla 2013b, p. 308.) The complexity, length and blandness of the knowledge cases in the contextualist literature may be a barrier to finding an effect of changing contexts on subjects' reactions. As Jennifer Nagel writes, "Participants who are not interested in a particular story may be more inclined to respond to it randomly" (Nagel, 2012, p. 516).

In this respect, following submaxim 6 of Travis's ("purely hypothetical") maxims of conversation ("Supermaxim: Always talk so as to leave the best impression of yourself possible"!) may make the difference between finding an effect of the occasion of use on our reactions to an example and failing to find such an effect:

> 6. Be as interesting as possible, and as inventive and imaginative as (but no more than) would be appreciated. (Subsubmaxim: Be bizarre, if you think your audience would go for that.) (Travis, 2008e, p. 61)

6 Conclusion

Travis's examples play a central role in his argument for occasion-sensitivity. According to Travis's meta-semantic view, the reactions of reasonable judges are constitutive of the

semantic facts, and we can use our reactions as indicators of what the semantic facts are (insofar as our reactions are reasonable). In §4, I presented several reasons to worry that features of the informal design and presentation of Travis's examples might make us question the reasonableness of our reactions. But those worries can be addressed by a more meticulous approach to the design of the examples: eliminating experimenter effects, order of presentation effects, evaluating the effects of contrast and the "rule of accommodation". And balancing the language involved in the contrasting pairs as much as possible improves the chances of uncovering previously unnoticed ways that the occasion of use affects what is said by the use of a sentence.

After producing two pairs of examples of knowledge ascriptions meant to illustrate occasion-sensitivity, Travis (1989, p. 157) says "Let us take it that we can now construct pairs of cases on the above model *ad lib*". But a great deal of care has to be taken in constructing the cases if they are to work as intended. Given their importance for his account of occasion sensitivity, it is essential that Travis's examples be designed to exclude features that might prevent us from reacting to them the way a reasonable judge would. The more sophisticated design described here should make Travis's examples even more appealing demonstrations of the complex ways that context and meaning interact.

References

Austin, J.L. (1966). "Three Ways of Spilling Ink". *The Philosophical Review 75*(4), 427–440.

Bargh, J.A., M. Chen, and L. Burrows (1996). "Automaticity of Social Behavior: Direct Effects of Trait Construct and Stereotype-Activation on Action". *Journal of Personality and Social Psychology 71*(2), 230–244.

Cappelen, Herman and E. Lepore (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.

Clapp, Lenny (2012). "Indexical Color-Predicates: Truth-Conditional Semantics vs. Truth-Conditional Pragmatics". *Canadian Journal of Philosophy 42*(2), 71–100.

Cullen, Simon (2010). "Survey-Driven Romanticism". *Review of Philosophy and Psychology 1*(2), 275–296.

DeRose, Keith (1992). "Contextualism and Knowledge Attributions". *Philosophy and Phenomenological Research LII*(4), 913–929.

DeRose, Keith (2005). "The Ordinary Language Basis for Contextualism, and the New Invariantism". *The Philosophical Quarterly 55*(219), 172–198.

DeRose, Keith (2009). *The Case for Contextualism*. Oxford: Oxford University Press.

DeRose, Keith (2011). "Contextualism, Contrastivism, and X-Phi Surveys". *Philosophical Studies 156*(1), 81–110.

Doyen, Stéphane, O. Klein, C.-L. Pichon, and A. Cleeremans (2012). "Behavioral Priming: It's all in the Mind, but Whose Mind?". *PLoS ONE 7*(1), 1–7.

Fantl, Jeremy and M. McGrath (2002). "Evidence, Pragmatics, and Justification". *The Philosophical Review 111*(1), 67–94.

Feltz, Adam and C. Zarpentine (2010). "Do You Know More When It Matters Less?". *Philosophical Psychology 23*(5), 683–706.

Ferguson, Kennan (2003). "Three Ways of Spilling Blood". In F. Debrix (Ed.), *Language, Agency, and Politics in a Constructed World*, pp. 87–100. Armonk, New York: M.E. Sharpe.

Gustafsson, Martin (2002). "Meaning, Saying, Truth". In M. Gustafsson and L. Hertzberg (Eds.), *The Practice of Language*, pp. 177–197. Dordrecht: Kluwer Academic Publishers.

Hansen, Nat (2011). "Color Adjectives and Radical Contextualism". *Linguistics and Philosophy 34*(3), 201–221.

Hansen, Nat (2012). "On an Alleged Truth/Falsity Asymmetry in Context Shifting Experiments". *The Philosophical Quarterly 62*(248), 530–545.

Hansen, Nat (2013). "A Slugfest of Intuitions: Contextualism and Experimental Design". *Synthese 190*(10), 1771–1792.

Hansen, Nat (2014). "Contrasting Cases". In J. Beebe (Ed.), *Advances in Experimental Epistemology*. New York: Continuum.

Hansen, Nat and E. Chemla (2013a). "Beating a Dead Donkey: Linguistic Experiments and Ordinary Language Philosophy". Unpublished ms. University of Reading and LSCP.

Hansen, Nat and E. Chemla (2013b). "Experimenting on Contextualism". *Mind & Language 28*(3), 286–321.

Intons-Peterson, Margaret Jean (1983). "Imagery Paradigms: How Vulnerable Are They to Experimenters' Expectations". *Journal of Experimental Psychology: Human Perception and Performance 9*(3), 394–412.

Kahneman, Daniel (2000). "A Psychological Point of View: Violations of Rational Rules as a Diagnostic of Mental Processes". *Behavioral and Brain Sciences 23*(5), 681–683.

Kahneman, Daniel (2011). *Thinking, Fast and Slow*. New York: Farrar Straus & Giroux.

Kahneman, Daniel and A. Tversky (1996). "On the Reality of Cognitive Illusions". *Psychological Review 103*(3), 582–591.

Kennedy, Christopher and L. McNally (2010). "Color, Context, and Compositionality". *Synthese 174*(1), 79–98.

Labov, William (1975). "Empirical Foundations of Linguistic Theory". In R. Austerlitz (Ed.), *The Scope of American Linguistics: Papers of the First Golden Anniversary Symposium of the Linguistic Society of America*. Lisse: The Peter de Ridder Press.

Lewis, David (1979). "Scorekeeping in a Language Game". *Journal of Philosophical Logic 8*(1), 339–359.

Miller, Alexander (2009). "Review of *Thought's Footing: A Theme in Wittgenstein's Philosophical Investigations*". *Mind 118*(469), 211–215.

Nagel, Jennifer (2012). "Intuitions and Experiments: A Defense of the Case Method in Epistemology". *Philosophy and Phenomenological Research LXXXV*(3), 495–527.

Pfungst, Oskar (1911). *Clever Hans*. New York: Holt, Reinhart and Winston.

Phelan, Mark (2013). "Evidence That Stakes Don't Matter for Evidence". *Philosophical Psychology* (iFirst), 1–25.

Pinillos, Angel (2012). "Knowledge, Experiments, and Practical Interests". In J. Brown and M. Gerken (Eds.), *New Essays on Knowledge Ascriptions*, pp. 192–219. Oxford: Oxford University Press.

Predelli, Stefano (2005). "Painted Leaves, Context, and Semantic Analysis". *Linguistics and Philosophy 28*(3), 351–374.

Rothschild, Daniel and G. Segal (2009). "Indexical Predicates". *Mind & Language 24*(4), 467–493.

Sainsbury, R.M. (2001). "Two Ways to Smoke a Cigarette". *Ratio XIV*(4), 386–406.

Schwitzgebel, Eric and F. Cushman (2012). "Expertise in Moral Reasoning? Order Effects on Moral Judgments in Professional Philosophers and Non-Philosophers". *Mind & Language 27*(2), 135–153.

Searle, John (2001). "J.L. Austin". In A. Martinich and D. Sosa (Eds.), *A Companion to Analytic Philosophy*, pp. 218–230. Oxford: Blackwell.

Shafir, Eldar (1998). "Philosophical Intuitions and Cognitive Mechanisms". In M. R. DePaul and W. Ramsey (Eds.), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, pp. 59–83. Lanham, MD: Rowman and Littlefield.

Sprouse, Jon (2011). "A validation of Amazon Mechanical Turk for the collection of acceptability judgments in linguistic theory". *Behavior Research Methods 43*(1), 155–167.

Stalnaker, Robert (1997). "Reference and Necessity". In B. Hale and C. Wright (Eds.), *A Companion to the Philosophy of Language*, pp. 534–554. Oxford: Blackwell.

Szabó, Zoltan Gendler (2000). *Problems of Compositionality*. Studies in Philosophy. New York: Garland Publishing, Inc.

Szabó, Zoltan Gendler (2001). "Adjectives in Context". In I. Kenesei and R. M. Harnish (Eds.), *Perspectives on Semantics, Pragmatics, and Discourse: A Festschrift for Ferenc Kiefer*, pp. 119–146. Amsterdam: John Benjamins Publishing Company.

Szabó, Zoltan Gendler (2006). "Sensitivity Training". *Mind and Language 21*(1), 31–38.

Travis, Charles (1978). "Meaning versus Truth". *Dialogue 17*(3), 401–430.

Travis, Charles (1989). *The Uses of Sense: Wittgenstein's Philosophy of Language*. Oxford: Oxford University Press.

Travis, Charles (2008a). "Annals of Analysis". In *Occasion-Sensitivity: Selected Essays*, pp. 65–93. Oxford: Oxford University Press.

Travis, Charles (2008b). "Attitudes as States". In *Occasion-Sensitivity: Selected Essays*, pp. 227–252. Oxford: Oxford University Press.

Travis, Charles (2008c). "Introduction". In *Occasion-Sensitivity: Selected Essays*, pp. 1–16. Oxford: Oxford University Press.

Travis, Charles (2008d). "On Constraints of Generality". In *Occasion Sensitivity*, pp. 271–289. Oxford: Oxford University Press.

Travis, Charles (2008e). "On What Is Strictly Speaking True". In *Occasion-Sensitivity: Selected Essays*, pp. 19–64. Oxford: Oxford University Press.

Travis, Charles (2008f). "Pragmatics". In *Occasion-Sensitivity: Selected Essays*, pp. 109–129. Oxford: Oxford University Press.

Travis, Charles (2008g). "Vagueness, Obseration, and Sorites". In *Occasion-Sensitivity: Selected Essays*, pp. 206–226. Oxford: Oxford University Press.

Vicente, Augustin (2012). "On Travis Cases". *Linguistics and Philosophy 35*(1), 3–19.

Williams, Merle A. (2009). *Henry James and the Philosophical Novel: Believing and Seeing*. Cambridge: Cambridge University Press.

Wright, Crispin (2001). "Wittgenstein's Rule-Following Considerations and the Central Project of Theoretical Linguistics". In *Rails To Infinity: Essays On Themes From Wittgenstein's Philosophical Investigations*, pp. 170–213. Harvard University Press.