

Indexical Predicates

Daniel Rothschild and Gabriel Segal

Truth-conditional semantics is the project of determining a way of assigning truth-conditions to sentences based on A) the extension of their constituents and B) their syntactic mode of composition. Truth-conditional semantics is the major research project of linguistic semantics and the project and its prospects are a central concern in contemporary philosophy of language.¹

Some linguists and philosophers argue that the fact that the extension of certain predicates appears to change dramatically across different contexts indicates that there is a fundamental problem with truth-conditional semantics.² We will state one version of this problem and outline an approach to it. We hope to advance the discussion of the issue by A) giving the explicit semantic theory and B) discussing some empirical considerations that motivate our approach (or at least fail to disconfirm it!).

Two vignettes³ can illustrate the phenomenon of predicates appearing to have context-dependent extensions:

The Greengrocer

The greengrocer stocks two types of watermelons. Both types are green on the outside, one has red flesh and the other has yellow flesh. A customer asks for a red watermelon. The greengrocer points to one and says, ‘How about this one? It's red.’

¹ See Harman (1972), Lycan (1984), Higginbotham (1985), Larson and Segal (1995), and Heim and Kratzer (1998).

² For discussion of this particular issue see Lahav (1989) Travis (1994) Szabo (2001), and Reimer (2002). For more general challenges to truth-conditional compositionality see Chomsky (1977), Sperber and Wilson (1986) Carston (2002), Recanati (2003) Cappelen and LePore (2005).

³ Borrowed from Charles Travis, pc.

The Artist's Studio

An artist is painting a still-life. On his desk is a red-skinned apple and that same watermelon, still green on the outside and red-fleshed. The artist points to the apple and says 'It's red.' He then points to the watermelon and says, 'It's not red.'

It is plausible to assume both of the following facts about these situations.

1. In the greengrocer's statement 'It's red', and the artist's statement 'It's not red', 'it' referred to a red-fleshed, green-skinned watermelon. (That is, the uses of 'it' did not refer to some part of the watermelon but rather the whole thing.)
2. The greengrocer's statement 'It's red', and the artist's statement 'It's not red' were both true.

If these assumptions are right, then it seems that the word 'red' can change its extension across contexts.⁴ Moreover the effect is dramatic enough not to be attributable to ordinary vagueness. Many have regarded examples of this sort as raising a significant challenge for the project of truth-conditional semantics.

We will be focusing on the particular kind of context-sensitivity exhibited in the vignettes. There are, of course, lots of different kinds of context-sensitive expressions, including indexicals, modals, light verbs, quantifiers, relational nouns, adverbs of quantification, and temporal expressions. We will give a semantics for a certain range of context-dependent predicates including 'red' and 'hexagonal'. We think that it's *possible* that this treatment could be extended to treat gradable adjectives such as 'tall', but we're not committed to that claim.⁵

⁴ We only say 'seems' since theoretically the extension of 'red' could simply include the watermelon at one time and exclude it at another time, even if it hasn't changed color in any respect: think 'grue'. This possibility should not be taken seriously for two reasons: the first is that this extension would be bizarre, and the second is that we could rewrite our situations so that the two utterances occur at the same time (i.e. put the artist's studio inside the greengrocer's).

⁵ See DeRose (forthcoming) and Hawthorne (forthcoming) for a discussions of gradable adjectives that indicates that they might fit well into the semantics we

Our semantics is not designed directly to account for all types of context dependence. And we are not committed to any particular views about the extent to which the ideas developed here might usefully be extended to other types of context dependence.

1 Contextual variation and compositionality

The kind of contextual variation of extension exhibited in the vignettes has been taken to present a prima-facie challenge to traditional truth-conditional semantics. Keble College, Oxford, is mainly built of reddish bricks. The examples we gave above make it plausible to think that in one context an utterance of ‘Keble College is red’ is true whereas in another context, an utterance of that same sentence might be false--this without Keble College having changed color. This possibility, along with three other assumptions, leads us to contradiction. One assumption, 2, is a weak compositionality principle. Another, 3, expresses commitment to the basic apparatus of traditional truth-conditional semantic theory. And the final assumption, 4, is that words like ‘red’ do not carry with them an indexical element that varies its extension across contexts. The inconsistent quartet is as follows:

1. ‘Keble College is red’ has different truth values in different contexts.
2. The semantic value of a complex expression is determined by its syntax together with the semantic values of its constituents.
3. The semantic value of an expression is its extension. The extension of a sentence is a truth value.
4. The extension of the components of ‘Keble College is red’ do not vary across contexts of utterance.

outline for ‘red’, which is a gradable adjective in some sense but not a canonical one like ‘tall’. These papers respond to a number of arguments that Jason Stanley has given against the indexical views of gradable adjectives like ‘tall’, see in particular Stanley (2002).

There are a variety of possible responses to the problem raised by the inconsistency of 1 through 4. Probably the most popular response is what can be called the *pragmatist* response. On this sort of line one denies assumption 2. This response would probably be offered by a certain type of fan of the later work of Ludwig Wittgenstein, a fan who thinks that the whole project of formal semantics for natural language is misguided. It is also offered by some theorists who are not at all skeptical about the semantic project, but who think that context supplies determinants of truth conditions that are not shadowed anywhere in the syntax of the object language, for instance any discourse representation theorists.⁶

Skeptics about formal semantics for natural language might or might not also wish to reject assumption 3. Non-skeptics might do so as well. They would accept 2, but deny that *truth-conditional* semantics is the way to go. We believe that Travis, Chomsky (1977), and Paul Pietroski would adopt this position.

Another line of response, typically called ‘semantic minimalism’, is to deny assumption 1. On this account, the truth-conditional semantic project is preserved, at the price of denying that sentences like ‘Keble College is red’ have different truth-conditions in different contexts. Supporters of this view include Cappelen and Lepore (2005) and Borg (2004). One problem with this view is that it rejects the idea that intuitions about the truth values of utterances of sentences provide good evidence about those sentences’ truth conditions. This makes it very hard to understand what the data for a semantic theory are supposed to be.⁷ In other words, it makes it very hard to see how one is supposed to tell whether a semantic account of some fragment of natural language is correct. If it really is so hard to judge whether a semantic theory for a piece of natural language is correct, then the flourishing field of linguistic semantics is in real trouble.

⁶ See, for example, Barbara Partee’s (1989) proposal for treating relational nouns.

⁷ Of course, we don’t regard speakers’ intuitions as sacrosanct. But it is very reasonable to assume, as linguists generally do, that by and large, speakers’ judgements are a reasonably good guide to truth conditions.

There is a less radical view that might amount to a rejection of 1. This view would have it that the strict and literal meanings of words, like ‘red’, determine a single context-independent extension for the words. Calling the watermelon ‘red’ for instance is simply a non-literal use. Red things (literally speaking) are things that are red on most of their visible surface. We don’t know of any philosophers who explicitly hold this view in print (since Cappelen and Lepore (2005) are silent about the extension of any term, their view appears to be compatible with this one). This view seems to be a simple empirical bet. If empirical studies (deficits, brain-imaging, patterns of acquisition or whatever) manage to isolate a core set of literal uses of words like ‘red’ and show that those uses are ones where there is no contextual variation, then this view may turn out to be correct. As far as we know, there is, as of yet, no empirical evidence to suggest that the division between literal and non-literal uses would vindicate a denial of assumption 1.

Another set of views including our own view here and that of Szabó (2001) involves denying assumption 4. On these views, there are one or more indexical elements in ‘Keble College is red’ the extensions of which vary across contexts. Szabó’s view is that at logical form ‘red’ associates with two variables, which have their values contextually determined. One of these picks out a comparison class, the other picks out the part of the object that has to be colored: ‘red(c, p)’.

Our analysis, by contrast, treats ‘red’ itself as a simple indexical, like ‘I’ or ‘that’. There are no variables associated with ‘red’. It’s just a word that happens to change extension across contexts.

2 Our Analysis

According to the standard picture of context-sensitivity, only a small set of lexical items vary their extension across contexts of utterance. These are the classic indexicals such as ‘I’, ‘here’, ‘now’, ‘that’, ‘he’, and ‘this’. These indexicals do not threaten the thesis of compositionality. It is just that semantic value varies

from context to context (in other words, for these words, one denies an equivalent of assumption 4.)

Our task here is to defend a view according to which the class of indexicals is expanded to include certain predicates, like ‘red’. In this section, we give a formal proposal for handling this contextual variation. Our proposal is formulated as a T-theory. This allows us to make clear exactly how one derives the truth conditions for the sentences we treat. We could equally have defined a function that goes from sentences and contexts to truth conditions. It should be easy to see (in principle!) how one would expand the T-theory presented here to cover more extensive fragments of natural language.⁸

Any adequate truth-conditional semantics requires a treatment of indexical elements. One standard treatment is the conditionalized T-theory approach developed by Burge (1974), Weinstein (1974), and Larson and Segal (1995). The basic idea is that the context-independent semantics provides the means to prove a T-theorem, given information about the extensions of expressions in specific contexts. Very roughly, the idea for, say, ‘that is remarkable’ is that the context-independent semantics allows one to derive a conditional along the lines of **D**:

(D) If u is an utterance ‘that is remarkable’ and the speaker uses ‘that’ in u to refer to x , the u is true iff x satisfies ‘is remarkable’.

Suppose that in a particular context, the speaker uses ‘that’ to refer to the Taj Majal. One can then go on to derive **(T)**:⁹

⁸ For outlines of compositional truth-conditional semantics for larger chunks of natural language see Larson and Segal (1995) or Heim and Kratzer (1998).

⁹ As David Kaplan (1977/1989) pointed out, the context-independent semantics of indexical expression is a function from contexts to extensions, a ‘character’. The Burge and Weinstein treatments in effect show how to assign characters to indexicals in a T-theory, without actually talking about functions.

(T) u is true iff the Taj Majal is remarkable.

We have adapted the basic idea behind the Burge/Weinstein sort of theory, so that it can apply to predicates. Our theory requires a rather complicated metaphysics of language. We group together all the tokens of an expression that occur within a given context: so, for example, if the customer says to the greengrocer ‘is that a red one?’ and the greengrocer says ‘yes, it is a red one’, we group together those two tokens of ‘red’. We treat the tokens of an indexical predicate that occur within the same context as tokens of a single syntactic type. No token of that type can occur in another context. We indicate the syntactic type by indexing. Thus all tokens of ‘red’ that occur in a given context receive the same index, as in: ‘red_j’.¹⁰ We will use numbers to keep contexts and these context-bound expressions in line: thus all the ‘red_j’s occur in the j -th context.¹¹ The context-bound expression types are subtypes of larger types, such as the one that includes all the ‘red_j’s, ‘red_k’s etc.. We can think of this larger type as context-independent and possessed of a context-independent semantics. Its semantics, intuitively speaking, is given by a function from contexts to extensions. The semantics of each context-bound ‘red_j’ is the extension it receives relative to its context. This extension is determined by the conversational standards of the context: an object satisfies a token of ‘red_j’ in a context, if it counts as red by the standards of that context. These context-bound indexical predicates occur in sentences, and the sentences have truth conditions relative to the contexts in which they occur. Thus ‘It is red_j’ uttered in the greengrocer’s context is true relative to that context iff the demonstrated melon is red by the standards of the context. We assume further that

¹⁰ ‘Context’ is here a slightly technical term. If it is possible for someone truly to say of a given object ‘Well, it is red, but it’s not red’, using ‘red’ in different senses, then we would represent this by something like this: ‘Well, it is red_j, but it’s not red_k’. Contexts are therefore fine-grained. We have no theory of contexts, but the key thing is that they must include the referential intentions of speakers. If someone insists on asking us what a context is, then we will say that it is an ordered pair of an utterance and the rest of the universe.

¹¹ Thanks to Jon Barton for this idea. If you are fussy, then notice that numbers number contexts and indices, which latter are pieces of syntax.

utterances of sentences containing context-bound indexicals have truth values. These are simply inherited from the truth value that the sentence uttered receives relative to its context. So the grocer's utterance of 'It is red_j' is true, absolutely, iff the sentence uttered is true relative to the context.

The melon that we have been discussing is red by the standards of the greengrocer, but not by those of the artist's studio. The context in which we have been discussing the melon is academia. We are writing a theoretical paper addressed primarily to other academics. It is the 'red' so-used that we need for the meta-language. Our conversational context is academia and we are discussing red in general, and not using 'red' to describe the colour of anything in particular. So we can appropriately label our context 'g', for 'general', and subscript 'red' appropriately: 'red_g'.¹²

Here is the T-theory:

Axioms

- (1) $(x)(n)(x \text{ satisfies 'Keble College', } c_n \text{ iff } x = \text{Keble College})$
- (2) $(x)(n)(x \text{ satisfies 'is red'}^n, c_n \text{ iff } x \text{ is red}_g, c_n)$
- (3) $(S)(NP)(VP) \text{ (If } S = NP \wedge VP, \text{ then } ((n)(S \text{ is true, } c_n \text{ iff } (\exists x)(x \text{ satisfies } NP, c_n \text{ and } x \text{ satisfies } VP, c_n)))$
- (4) $(u)(n)(S) \text{ (if } u \text{ is an utterance of } S \text{ in } c_n, \text{ then } (u \text{ is true iff } S \text{ is true, } c_n))$

¹² We have no very strong views about how 'red_g' works in constructions like 'red_g by the standards of context c'. But we do have a suggestion. We suggest that 'red_g' has a wide extension: things that are red by some standard or other. But we also suggest that the precise extension of 'red' (plus index) is irrelevant in this particular construction. In the greengrocer, someone might say 'Yes, the melon is red_j ... although of course not by the standards of the artist's studio'. Here it looks as though the implicit 'red' in the second sentence is 'red_j'. But the sentence means just the same as 'not red_g by the standards of the artist's studio'. So it looks as though 'by the standards of context c' in effect works like a functor that maps any 'red_n' extension onto things that are red by the standards of c.

Assume information about a particular context c_j :

(5) u_j is an utterance of $S_j = \text{'Keble College is red}_j\text{'}$ in c_j

(6) $(x)(x \text{ is red}_g, c_j \text{ iff } x \text{ is red}_j)$

(To understand the occurrence of the meta-linguistic 'is red_j' on the right-hand side of 6 it is necessary to be in c_j . All that is required for this is that one knows enough about the intentions of the speaker and so on to be able to use 'red' in the same way as the other participants.) We can now plug the contextual information into the context-independent T-theory, and derive a T-theorem, as follows:

(7) $(x)(x \text{ satisfies 'is red}_j\text{'}, c_j \text{ iff } x \text{ is red}_g, c_j)$

[(2)]

(8) If $S_j = \text{'Keble College is red}_j\text{'}$, then $((n)(S_j \text{ is true, } c_n \text{ iff}$

$(\exists x)(x \text{ satisfies 'Keble College', } c_n \text{ and } x \text{ satisfies 'red}_j\text{'}, c_n)$

[(3)]

(9) $S_j = \text{'Keble College is red}_j\text{'}$

[(5)]

(10) $(n)(S_j \text{ is true, } c_n \text{ iff}$

$(\exists x)(x \text{ satisfies 'Keble College', } c_n \text{ and } x \text{ satisfies 'is red}_j\text{'}, c_n))$

[(8) (9)]

(11) $S_j \text{ is true, } c_j \text{ iff}$

$(\exists x)(x \text{ satisfies 'Keble College', } c_j \text{ and } x \text{ satisfies 'is red}_j\text{'}, c_j)$

[(10)]

(12) $S_j \text{ is true, } c_j \text{ iff } (\exists x)(x \text{ satisfies 'Keble College', } c_j \text{ and } x \text{ is red}_g, c_j)$

[(11), (2)]

(13) $S_j \text{ is true, } c_j \text{ iff } (\exists x)(x \text{ satisfies 'Keble College', } c_j \text{ and } x \text{ is red}_j)$

[(12), (6)]

(14) $S_j \text{ is true } c_j \text{ iff } (\exists x)(x = \text{Keble College and } x \text{ is red}_j)$

[(13), (1),]

(15) If u_j is an utterance of S_j in c_j , then $(u_j \text{ is true iff } S_j \text{ is true, } c_j)$

- [(4)]
- (16) u_j is true iff S_j is true, c_j
 [(5), (15)]
- (17) u_j is true iff $(\exists x)(x = \text{Keble College and } x \text{ is red}_j)$
 [(14), (16)]
- (18) u_j is true iff Keble College is red_j
 [(17)]

It is worth noting that our analysis is similar to another, more pragmatic treatment, of ‘red’. On this treatment, ‘red’ has a constant, context-independent semantic value, which value gets enriched in a given context of utterance (see for instance, Recanati, 2003). The newly enriched value then combines compositionally with the rest of elements in the sentence in the usual way. This theory, unlike ours, gives a genuine role to pragmatic enrichment (rather than indexical resolution) in the determination of meaning. On the other hand, the pragmatic enrichment plays the same formal role as the indexical resolution does in our T-theory, so there is little difference in the basic structure of the accounts. We think that this account embodies the same basic idea as ours embedded in a different theoretical framework.

3 Prima-facie Virtues of Analysis

One prima-facie virtue of our analysis is that it preserves both the truth-conditional compositionality of natural language and the idea that our intuitions about the truth conditions of utterances provide reliable data for a semantic theory.¹³ Pragmatism and minimalist semantics sacrifice one or both of these desiderata. We think we should only resort to such measures if forced to.

¹³ At least that such intuitions are reliable with respect to these sorts of predicate; there are types of construction where truth conditions that normal speakers associate with sentences must be explained in part by non-semantic processes, but it is good methodology keep these areas circumscribed and properly accounted for.

Ultimately we may need to relax these assumptions, but as long as good work is based on them, we should try to preserve them if possible.¹⁴

Another virtue of our semantics is its combination of its simplicity and strength. First of all, we do not posit intermediate levels of semantic representation between the output of semantic processing and the extension of an utterance. In other words, we do not posit mechanisms that take sentence meaning along with the broad context as an input and give the truth-conditions of what was said as an output.¹⁵

Secondly, we posit a relatively simple syntax, with only a little more structure than meets the eye. In this respect our proposal is more economical than Szabó's, who posits two hidden variables. Our proposal is as simple as a theory that posits just a single hidden variable referring to the context of utterance and serving as a catch-all covering a very wide range of contextual variations. Our proposal extends to many varieties of contextual variation in which language leaves it up to speakers to pin down an extension. There are many different kinds of case of this ilk. We will mention are just five forms of context-sensitivity that our theory might plausibly handle.

Our proposal can handle the context-sensitivity of vague predicates with varying standards of precision such as 'hexagonal'. It can handle the context-sensitivity of 'knows'. It can handle an element of the context-sensitivity of 'tall' that is not captured by a variable for comparison class: there is variability of how tall one has to be relative to a comparison class to satisfy an utterance of 'tall' (this point is from Fara 2000). It can handle Chomsky's example of 'water'. Water infused

¹⁴ Some philosophers might doubt that truth-conditional linguistic semantics is a flourishing field. Since most people engaged in the project are linguists not philosophers, such doubt is understandable. However, a look at the increasing number of articles, NSF grants, PhD theses, academic jobs, and so on in linguistic semantics might be taken as some evidence of the health of the field.

¹⁵ Relevance theory (Sperber and Wilson, 1986) takes such a mechanism to be a critical part of our linguistic capacity separate from our semantic processing.

with tea leaves does not fall in the extension of most utterances of ‘water’. But if tea leaves were added to tap water at source as a purifier, then the tea-infused water coming from our taps would fall in the extension of ‘water’. Finally, our proposal can handle the ‘forward’ that appears in ‘the meeting was moved forward an hour’, which can mean either that it took place an hour earlier than originally planned, or an hour later, depending on the speaker’s perspective.¹⁶ Our simple proposal thus covers a very wide range of data.

We do not pretend that our semantics explains things it does not. The semantics does not explain how, in a given context, a person determines what the extension of ‘red’ is, i.e. what it is to be red by the standards of the context. On our account, this is an example of indexical resolution. ‘Red’ is rather like ‘that’, in that it is a non-trivial task to explain how the term acquires an extension in a context.

Ultimately how we succeed in communicating with indexical expressions may not be a question which formal semantics itself has much to say about. This is not to say we think that an account of this is either unimportant or easy to give. But it is not part of semantic theory. Our semantic proposal (like many other treatments of demonstratives and indexicals) isolates this complex area from formal semantics.

In sections 4 and 5 we consider how our proposals compares to Szabo's in relation to some specific pieces of data. In section 6 we consider some conceptual issues arising. In section 7 we offer a tentative sketch of a taxonomy of different kinds of context-sensitive expressions.

4 Binding Considerations

The recent literature in philosophy of language is strewn with controversy over the issue of binding and its relation to the semantics/pragmatics distinction.¹⁷ As

¹⁶ For an interesting discussion of this phenomenon, see Pinker (2007), 191.

¹⁷ See for instance Stanley and Szabó (2000), Stanley (2002), Neale (forthcoming).

in first-order logic, a bound variable in natural language is a variable bound by another expression. Linguistics provides many examples where certain expressions, most notably pronouns, behave like bound variables:

(1) Every man likes his mother.

On one reading of (1), it is clear here that ‘his’ functions as if it were a variable bound by the quantificational expression ‘every man’. So a pre-theoretic way of writing out the logical form of (1) would be (2):

(2). For every man x , x likes x ’s mother.

In this case, there are two bound variables in (1).¹⁸ In addition to bound uses, third person pronouns also exhibit free uses:

(3) He was a good man.

In examples like (3) we generally interpret ‘he’ as referring to some contextually salient male individual. This use of third-person pronouns is sometimes called a ‘deictic’ or ‘indexical’ use since it resembles the use of indexicals like ‘that’ and ‘you’.

The exact treatment of bound and unbound pronouns and their relation to each other is a question of considerable controversy within syntax and semantics.

Whether a unified account can be given or not, work needs to be done.¹⁹

There is one important point to note here: a simple indexical semantics for ‘he’ cannot account for its bound use. It is easy to see this: on the simple indexical semantics ‘he’ simply refers to some contextually salient individual. But there is no contextually salient individual which, if taken as the value for ‘his’, in (1) would give us a bound reading. The point is simple and has been made before

¹⁸ Of course, quantifiers may be treated in a variety of possible ways in different semantic theories, our treatment here is just mimicking first-order logic for ease of understanding.

¹⁹ The first steps towards a unified account were made in the early eighties by Heim (1982) and Kamp (1981). There is still considerable debate over the success of these ‘dynamic’ approaches to meaning.

many times: the standard semantics for indexicals does not allow indexicals to act like bound variables (Evans, 1977, Heim and Kratzer, 1995, ch. 9).

In this paper we give a simple indexical semantics for predicates like ‘red’ and ‘hexagonal’. If these predicates exhibit something analogous to the bound usages we find in (1) then our semantics will not be able to account for that.²⁰

This leaves us with the following questions: what sort of ‘bound’ behavior might our predicates exhibit that we would be unable to account for? There are really two different possibilities here. One is that, like ‘he’, our predicate itself acts like a bound variable relative to some quantifier (presumably a quantifier of colors). The other is that our predicate acts like it contains a bound variable, even if it itself is not one. An example of the latter sort of behavior is provided by the noun phrase ‘enemy’:²¹

(4) Every man faced an enemy.

This sentence has a reading under which it means that for every man x , x faced x ’s enemy, not just an enemy of someone or other. In that respect, the sentence behaves as if it has a bound reading. We will call this kind of binding *implicit argument binding* to reflect the fact that what appears to be bound is not the expression itself but some implicit argument.²² In these accounts, what is bound is not so much an implicit argument but either a situation variable or a different pronoun that is pragmatically supplied to fill out a definite description which replaces the original pronoun.

²⁰ What conclusion we should draw from the presence of bound uses is hard to say: the point is one needs to account for them somehow. There are really two options, bring them into the semantics or describe a pragmatic mechanism that can deliver the right readings. Many have read Stanley (2002) as claiming that the latter option is not possible. Such a claim cannot be right: it might for some reason be undesirable to posit pragmatics mechanisms that deliver bound readings where the syntax has no binding, but it is not theoretically impossible. For discussion see Recanati (2003) and Neale (forthcoming).

²¹ This example is from Partee (1989).

²² Note, as pointed out by Neale (2004), that a version of implicit argument binding is used by almost all non-dynamic approaches to donkey anaphora, e.g. Heim (1990) and Neale (1990).

Let us address the possibility of direct binding of ‘red’ first. On our view, ‘red’ is an indexical expression that picks out predicate extensions. Naturally these extensions tend to include just red things. By way of contrast, in its unbound use, ‘he’ is an indexical expression that tends to refer to male individuals. But the pronoun can also be bound by a quantifier quantifying over male individuals.

We will argue that there is no reason to think ‘red’ exhibits the same type of direct binding that ‘he’ does. Here are two examples with ‘red’ in a position where it might be bound by quantifiers over shades of red, each example is paired with a structurally similar example of real binding with pronouns. (The index i shows where the binding is meant to take place.)

- (5) a. Every shade of red _{i} is such that there is a house that is red _{i} .
b. Every man _{i} is such that someone hit him _{i} .
- (6) a. Every shade of red _{i} was used to paint a house red _{i} .
b. Every man _{i} was his _{i} own best friend.

In (5a) the second use of ‘red’ is clearly in a position to be bound by the quantifier ‘every shade of red’. However, (5a) does not exhibit the bound reading according to which for each shade x of red, there is a house that is shade x .²³ On the other hand in (5b) the bound reading is clearly available for the pronoun ‘him’ in a similar syntactic configuration. Likewise (6a) does not seem to have a reading on which the second use of ‘red’ is bound by the quantifier, under which it would mean that every shade of red x was used to paint a house shade x . (Compare (6b))

We should not be surprised that ‘red’ does not exhibit directly bound readings.

There are very few kinds of expressions that can be directly bound by quantifiers.

²³ Perhaps ‘every shade of red’ is somehow not the right quantifier phrase, however, we could not find any that work. We do not discount the possibility that the problem is that a nominal quantifier cannot bind something that’s grammatically an adjective. The problem is that any non-nominal quantifiers (like ‘everywhere’) are also modifiers, and using modifiers to give evidence of binding faces serious problems, as we argue below.

It would have been a surprising syntactic/semantic discovery if ‘red’ could be *directly* bound.

This leaves us with the question of whether ‘red’ exhibits implicit argument binding. Recall that implicit argument binding occurs where an expression contains an unpronounced argument that is bound by a quantifier (or anyway shows behavior that would naturally be analyzed in this manner). Szabó gives us somewhere to search. Recall that he proposed that two variables associate with ‘red’: a location variable (p) and a comparison class variable (c). So ‘red’ looks like ‘red(c)(p)’. First, let us concentrate on the location variable, ‘ p ’. Szabó himself says next to nothing about the nature of these variables, so we cannot be sure whether or not he would expect these variables to be bindable.

All we can do is look at the facts. It is hard to get an example where a bound reading of the hypothetical location variable would appear likely. Imagine, for instance, that one has a cube that has a small sentence written on each face describing the face’s colour. Please consider:

(7) Each face _{i} says the cube is red(c)(p_i).

(8) Each man _{i} says that John hit him _{i} .

On its bound reading, (7) would mean that each face says that the cube is red on that face. We cannot get this reading, nor can we find any other clear examples where it looks like a location variable is bound. As (8) indicates ‘red’ seems to be in the right syntactic position to be bound in (7).²⁴

There is another possible method of binding that needs to be considered. This is binding that occurs when a modifier is introduced, i.e. some sort of adjunct that includes a quantifier. Here is an example:

(9) The field is brown in many places.

²⁴ This is far from certain, however, as direct objects and predicates might not really be in the same syntactic position.

Now it is certainly true that, in some sense, (9) is only saying that the field is brown in the places being quantified over by ‘many places.’ However, this does not show that there is a bindable element in the predicate ‘is brown.’ For we can understand ‘in many places’ as a modifier which takes something of semantic type $\langle e,t \rangle$ (i.e. ‘is brown’) and returns something of semantic type $\langle e,t \rangle$. In doing so it may add into the logical form the element which is bound by the quantifier, so there is no need to suppose that a bindable element was already there. In other words, there is no reason to suppose that the logical form (9) is structurally like (10) rather than (11):

(10) many places (x), the field is brown(x).

(11) many places (x), the field is (brown in x).

Indeed, there is a good reason to understand (9) as having an LF more like (10) than (11). Even when we have quantified over locations there is still a possibility of contextual variation in terms of location of coloring. For instance, (9) could mean in some contexts that *the ground* is brown in many places. In other contexts it could mean that *the plants* are brown in many places. But if we treated the locational variation as a case of actual binding, this extra variation would be left unaccounted for.

It’s worth noting that exactly similar considerations invalidate a naive version of the well-known ‘binding argument’ presented by Jason Stanley (2002) with respect to the sentence ‘it’s raining.’ This naive reconstruction of the argument (which Stanley did not put forward) notes first that the locational variation in ‘it’s raining’ suggests the presence of a bound variable:

(12) Everywhere I go, it rains.

Now the argument goes on to conclude that because a variable was bound in (12), there must be a locational variable all occurrences ‘it rains’. The joke is, of course, that ‘everywhere’ is not simply a plain quantifier over locations, it’s also, grammatically speaking a modifier. Moreover whatever ‘binding’ it accomplishes is not sufficient to capture all the contextual variation over location. For we could

imagine that the locations it quantifies over are too coarse-grained to capture the actual locations of rain:

(13) In every country I go to, it rains.

It is entirely possible that (13) could be false even though on a trip during which I went to Spain, Germany and Italy, it rained *somewhere* in each of those countries while I was there (but that it didn't rain where I was). What this indicates is that whatever binding is happening, it is not simply a matter of the noun phrase 'every country' directly binding the sole location variable in 'it rains'.²⁵

So much for location-variables being implicated in implicit argument binding.

The second variable that Szabó posits is one for implicit comparison classes (following Ludlow, 1989). Implicit comparison classes are classes of items that provide standards for what shades are red enough to counts as 'red'. So being red for an apple is different from being red for a sunburned face, though in both cases it is just the surface that is colored. An apple that is as red as a typical sunburned face might be called 'pink' instead of 'red'.

Of course, one can use a modifier to 'bind' a comparison class variable. One can say, for instance, that every boy bought an object that was red *by comparison to the other objects in his house or for its kind*. However, in these cases the binding is introduced by a modifier, and this is compatible with our semantic proposal for 'red'. What about cases without modifiers? It does not suffice merely to cite sentences like this:

(14) Every kind of animal in the zoo has a member that is red.

One can imagine a situation in which being red for a squirrel, say, is different from being red for a goldfish. However, one extension of 'red' may still do the

²⁵ Stanley's actual example was more complex:

(i) Every time John lights a cigarette it rains.

One puzzle here is that what seems like a temporal modifier acts *also* like a locational modifier---suggesting that tense might also encode locations. Perhaps this is because "'every time'" quantifies over situations rather than just times.

work here: it just needs to pick out canonically red squirrels and goldfish.²⁶ One can try to construct an example for which one extension of ‘red’ will not be adequate. For instance:

(15) Only one collection of fruit was next to exactly one red color tile.

It would be interesting if this could mean something like this:

(16) Only one collection of fruit_i was next to exactly one red(c_i)(p) color tile.

Evaluate (15) in the following scenario. There are two collections of fruit between which is a color tile that counts as red with respect to only one of them. No other color tiles are nearby. If (15) was interpreted like (16), then it would be true in the scenario. But it seems impossible to get this true reading. So, at least on a preliminary examination, there is no compelling evidence that there is a bindable comparison class variable in ‘red’.

Another possibility, departing from Szabó’s proposal, is that ‘red’ has a situational variable and so can be dependent on quantification over situations or contexts. The idea here is that each context or situation provides standards for what counts as red (perhaps by making salient some group of individuals). So, on this proposal ‘red’ itself simply has a bindable situation variable. Here the data are more equivocal.²⁷

²⁶ Zoltan Szabó (p.c.) notes that our proposal might make bad predictions if we assume something like the following, plausible principle:

Color Constancy If two objects are the same shade of color in some context C, then if one of them is red by the standards of C, so is the other.

The problem is brought out by a case where there are no canonically red goldfish but there are still some goldfish that are the same shade of red as, say, a red squirrel is. It seems like in this circumstance (14) might be false. Our response is to reject **Color Constancy**. In a given context, one extension of ‘red’ will include certain sunburns as red while not including identically colored color-tiles. See DeRose (forthcoming) and Hawthorn (forthcoming) for a similar discussion of ‘tall’.

²⁷ In small sample email surveys we conducted on these sorts of examples we found considerable variation and there were differences of opinion on many of the examples even between the two of us.

Let us consider a variation of our earlier example with the artist and the greengrocer.

The Artist and the Greengrocer

Suppose that are two situations: the greengrocer and the artist's studio. In the artist's studio there is a beautiful red apple and a red watermelon (red on the inside). In the greengrocer the situation is different. The greengrocer has six watermelons, five of which are yellow (on the inside) and one of which is red (on the inside). A customer comes in and asks for a red watermelon and the greengrocer points to the requisite item.

Assume this particularly under-stocked greengrocer has no other fruit.

Here is a hypothetical sentence for describing these situations.

(17) In each situation there was exactly one red piece of fruit.

It would seem that (17) could only be true if each situation has different standards of redness. The key point is that one loose sense of 'red' will not do since on this sense a red-fleshed grapefruit is either red or not red. But if 'red' only has one extension in (17) then (17) will have to be false. So to make (17) true we need the grapefruit to count as red in the greengrocer but not to count as red in the artist's studio. This suggests that the logical form of (17) might be something like (18):

(18) In each situation s there was exactly one red(s) piece of fruit.

In a small email survey we found that some people found it natural to interpret sentences like (19) in such a way that they come out true in the scenario described. Others did not.

For a similar example with 'hexagonal' we found that even fewer people found the equivalent of (19) acceptable. The example goes as follows:

The Hexagon Scenario

Many agree that when we talk about geography, a map of France counts as hexagonal. When we talk about architecture or geometry, only something that's closer to a perfect hexagon counts. Suppose that in the architecture classroom there is a map of France and an accurate drawing of a hexagon hanging on the wall, and in the geography classroom there's a map of the

UK and a map of France hanging on the wall, and in the geometry classroom there is just a drawing of an as-near-as-possibly perfect hexagon hanging on the wall.

The sentence we used to test whether ‘hexagonal’ could have a bound context variable was this:

(19) In each classroom, there was exactly one hexagonal shape hanging on the wall.

We found that most people could not get a true reading of (19) in the situation described.

Turning back to ‘red’, even though some people found (17) acceptable, this might not mean that ‘red’ really exhibits bound uses. For one possibility is that those individuals were using a meta-linguistic sense of ‘red’. In other words they interpreted (17) as meaning that in each situation exactly one piece of fruit was ‘red’ (in the sense that it could be called ‘red’ in that situation).²⁸

Our purpose here is not to establish definitely that ‘red’ does not exhibit bound uses---that would take more careful empirical study. All we hope to have done is to establish the plausibility of the claim. The point is that the evidence of bound readings for ‘red’ is weak compared to the evidence of bound readings for words like ‘enemy’ or for quantifier domain restrictions (Stanley and Szabó, 2000).

The only plausible candidate for binding we found was in ‘situation’ type examples. But many people found even these unacceptable. And it is reasonable to account for the judgements of those that did find them acceptable by appeal to meta-linguistic readings. Or, to suggest another possibility, we might hypothesize that they were using an expanded ‘philosopher’s English’ that over-extends the principle of charity---since many of the people surveyed were philosophers. So we will assume that it is plausible that ‘red’ does not introduce a bindable element. Even if this is wrong, we can still propose our semantics for other

²⁸ See Horn (1989) for arguments for the presence of unobvious meta-linguistic readings in ordinary speech.

predicates (such as, possibly, ‘hexagonal’) which don’t give any indication of binding.

If there is no binding phenomenon with ‘red’ or ‘hexagonal’ this does not itself show that our indexical semantics for these terms is correct, or that Szabó’s is incorrect. Nonetheless if it is true, it adds to the tenability of our proposal by showing that we do not need a more complex semantics for ‘red’ in order to handle any bound readings (as is needed for ‘he’, ‘she,’ and ‘it’).

4 Loose Extensions

Recall that our semantics differs from Szabó’s in that rather than positing indexical variables in the syntax introduced along with the predicate ‘red’ we treat ‘red’ itself as an indexical. We argued that our account is both more parsimonious and more powerful than Szabó’s and thus, all else equal, should be preferred.

However, we also think that there are empirical considerations that bear against Szabó’s account. We argued in the previous section that sometimes ‘red’ in a given context can have as its meaning a loose extension that includes a number of things which are red in different ways (e.g. by being red in different locations, etc). We will make two arguments here: 1) that there is direct evidence for such loose extensions and 2) that Szabó’s account cannot handle them. We will also relate these considerations to recent discussions of comparative adjectives, such as ‘tall’.

First of all, evidence for loose extensions. Consider this situation:

The Orange and the Apple

A man is at the greengrocer’s where various fruit, including blood oranges with red flesh and normal, red apples, are displayed uncut. He says to the

greengrocer, ‘I want red fruit for my dinner party, I think red fruits are tastier, and they go better with my tablecloth.’

In this situation, the use of red in ‘red fruits’ can clearly stretch to include both fruits that are red on the inside and fruits that are red on the outside, in this context both the blood orange and the apples. If there is doubt on whether these loose readings really occur we note that one can also say, in many contexts, (20). (20) The apple is more red than the blood orange.

There is clearly a reading of (20) that means that the apple is more red on its surface than the blood orange is in its flesh. But to have this reading the relevant sense of ‘red’ involved needs to be one that covers two different locations.²⁹

Can Szabó’s account handle these loose extensions of ‘red’? That is actually not an easy question to answer. The basic issue is whether there is a single value for the location variable associated with ‘red’ which can cover redness on the surface of some fruit but on the inside of others. Clearly the variable cannot have as a value something like “an arbitrary part”, since that is not a possible value for a variable. Moreover, merely existentially quantifying the variable will not give the desired result, since the whole point is that it is not clear that there is one “location” that covers both areas.

An anonymous reviewer suggested (in a slightly different context) that the variable might function like an ‘e-type description’ (See Evans 1977, Neale 1990). Pronouns, after all, sometimes have e-type semantics, and since hidden variables are taken to be analogous to pronouns we might think that Szabó’s variable could work that way. Consider this pair of sentences:

(21) Susan has a son. Mary talked to him.

Many theories of pronouns take “him” in the second sentence to spell out at some level of analysis as a definite description of the form “the son Susan has”.

Likewise one might think that the variable in the use of ‘red’ in **The Orange and**

²⁹ Thanks to Chris Kennedy for this suggestion. Similar data can be found with ellipsis.

The Apple scenario is really a description of some sort. In this case we would need to find a definite description that covers both the inside of the blood orange and the outside of the apple. The problem with this idea is that e-type uses of descriptions invariable require explicit linguistic antecedents (as in “a son” in (21)). But in our example there is no such antecedent.

So it is simply not clear what the value of the location variable is in **The Orange and the Apple** scenario. This does not prove that Szabo’s account cannot handle these cases, but it shows that there is a significant problem in doing so. Indeed it shows that there is a general lacuna in Szabo’s theory in that it does not make clear what range of values the location variable can have. On our account, by contrast there is no problem in handling this case. We just assume that in this case the particular utterance of ‘red’ in the context expressed a loose extension that covers both the apples and the oranges, an extension similar to that of the expression ‘red on some prominent part’.

It’s worth noting that a similar observation has been made with respect to contextual variance in scalar adjectives by Klein (1980) (see also Ludlow (1989)). The relevant data are very clear. Suppose there are two children, John a four-year-old and Jenny a fourteen-year-old. One can say of them:

(22) John is tall and Jenny is too.

It is important to note that when there is syntactic ellipsis of this sort, it is generally accepted that the elided material must be semantically identical to the unelided material, so that the elided use of ‘tall’ in the second conjunct of (22) must have exactly the same semantic value as the use of ‘tall’ in the first conjunct.

If there were simply a free comparison-class variable in the use of ‘tall’ in (22) then we would expect (22) to mean that John is tall for a four-year-old and Jenny is tall for a four-year-old also (or that they are both tall for fourteen-year-olds). However, the most natural meaning of (22) is not this one but rather one which

says that each is tall for their own age group. Our indexical semantics for ‘red’ if applied to ‘tall’ can capture this reading: it just requires picking an extension of tall that includes all members who are tall for their age group (and, perhaps, gender).

Note, however, that ellipsis is not necessary to show this point. In appropriate contexts we can find uses of ‘tall’ that pick out tall members of different kinds. (23) Children who are tall are generally well nourished.

The variable-heavy view of tall (Ludlow 1989, Stanley 2002, 2005) assumes that tall contains a comparison class variable, ‘c’, and that the meaning of ‘tall(c)’ is essentially tall for things of class c, which yields (vagueness aside) an extension consisting of a set of things higher than a standard determined relative norms for class c. If this is the semantics of ‘tall’ however, then (23) can only have a reading on which children above some set height are well nourished. This is clearly not the intended reading of (23), the intended reading is that children who are tall for their age are well-nourished.

Both Ludlow and Stanley³⁰ have gone to considerable lengths to defend their theories against the problem Klein posed with ellipsis.³¹ However, it should be clear that no simple modification of their accounts can handle cases of loose extensions *without ellipsis*, such as that involved in (23). This is because on their accounts any given use of ‘tall’ picks out all and only entities above some contextually set height (where the contextually set height is determined by which

³⁰ Stanley gives an account of ‘tall’ that is similar to Ludlow’s except that instead of just having a comparison class variable, there is both a functional variable and an objectual variable as its argument which together yield a comparison class.

³¹ The basic idea of both their responses to the problem of ellipsis is to propose that a variable in ‘tall’ is bound by something outside the ellipsis. It is known that bound variables can be elided and thus have different values in ellipsis as in this example:

(ii) John talked to his shrink and Bill did too.

The bound, or ‘sloppy’ reading is the one on which John and Bill each talked to their respective shrinks.

group is in question).³² But sentences like (23) show that this cannot be the case—for on the standard reading ‘tall’ picks out all 10-year olds above some height but not all 15-year olds above the same height (see DeRose forthcoming, for a similar discussion).³³

This indicates that our indexical semantics has an empirical virtue over variable-based accounts of context sensitivity for both ‘tall’ and red’.³⁴

5 Conceptual Issues

The empirical considerations we just went through about binding and ellipsis do not themselves make the case for our analysis. Their purpose was simply to show

³²We have some hesitation here. Stanley posits variables associated with ‘tall’ without giving an explicit semantics for ‘tall’ (not even a toy semantics). In any event, we assume based on his writing that this is the type of semantics he has in mind. In Ludlow’s case this is what we understand him to mean by his meta-language expression “x is tall for a y” (1989, p. 532), but he does not go into detail on this point.

³³ Stanley (2005) gives an interesting objection to the indexical semantics for ‘tall’ (his actual example uses ‘old’ but the point works for ‘tall’ also). He argues that one extension for ‘tall’ is not always sufficient to capture bound readings such as: (iii) Every sports team has exactly one tall member.

The problem he suggests is that some people can play two different sports and be tall for one of them but not the other (e.g. Bo Jackson). Perhaps such examples show that a purely indexical semantics is undesirable for ‘tall.’ However, we have some doubts. It may be rather that it is not truly individuals in the usual sense that satisfy the extension of ‘tall’ but rather (something like) individuals under guises. So for instance, the following sentence might be true:

(iv) Bill Clinton, the president, was great, Bill Clinton, the husband, wasn’t.

The interesting thing about this example is that the elided adjective has no obvious possible differing objects to bind it but nonetheless changes interpretation. This suggests that the proper satisfiers of some adjectives are not individuals in the usual sense. If this is correct, then an indexical semantics can survive Stanley’s objection. We want to stay neutral, however, on the question of whether an indexical semantics is appropriate for ‘tall’ as well as ‘red’. As Chris Kennedy pointed out to us, the same point can be made using comparatives:

(v) Bill Clinton, the president was better, than Bill Clinton, the husband.

For some supporting observations about this topic see Hawthorne (forthcoming).

³⁴ See DeRose (forthcoming) and Segal (forthcoming) for further objections to hidden-variable accounts.

that our analysis was tenable and empirically preferable to Szabó's similar analysis. The main argument in favor of our analysis is its simplicity and its compatibility with the aims and methodology of truth-conditional, compositional linguistic semantic research---aims and methodology which we think should only be abandoned for good reasons.

A typical response to our sort of proposal is to argue that our treatment of 'red' as an indexical is undesirable on, roughly speaking, conceptual grounds. Since basically conceptual grounds motivate our approach we need to be able to respond to this charge.

The most prevalent argument against our kind of proposal is the claim that 'red' does not feel like an indexical. Perhaps the mildest and most cautious complaint along these lines comes from Travis (1997) who writes:

There are several respects in which the present phenomena are unlike central cases where the parameter approach seems promising. One difference is this. In the central cases, such as 'I' and 'now', pointing to given parameters seems to be part of the terms meaning what they do.... By contrast, it is not part of what 'green' means, so far as we can tell, that speakings of it speak of, or refer to, such-and-such parameters. If its contribution, on a speaking, to what is said is a function of some parameters – say, implausibly, the speaker's intentions – saying so is not part of what 'green' means. The parameter approach does not automatically suggest itself as it did with 'I'.

One might say same thing about treating 'the' as a quantifier:

There are several respects in which the present phenomena are unlike central cases where the quantifier approach seems promising. One difference is this. In the central cases, such as 'all' and 'some', expressing generality seems to be part of the terms meaning what they do.... By contrast, it is not part of what 'the' means, so far as we can tell, that speakings of it speak generally. If its contribution, on a speaking, to what

is said is to express generality saying so is not part of what ‘the’ means.

The quantifier approach does not automatically suggest itself as it did with ‘all’.

We concede that neither account automatically suggests itself. We concede, too, that maybe the claims of semantic theorists about ‘green’ and ‘the’ do not conform to prima facie intuitions about what ‘seems to be part of the terms meaning what they do’. But such shallow seemings provide no real evidence against independently-motivated semantic theories.³⁵

Neale (forthcoming) is more outspoken. He claims that a key feature of indexicals is their perspectival nature:

I am strongly inclined to think that indexical words are *essentially* perspectival, that perspective is the *hallmark* of indexicality. (p. 336)
... We should be skeptical about any claim to the effect that an expression (phonic or aphonic) is indexical if the expression is not perspectival in some way. This is one reason I am deeply skeptical about “contextualist” accounts of the meaning of ‘know’. The idea that this verb is indexical in some way makes a mockery of the idea of indexical expressions. (p. 337)

Neale argues at length that every syntactically real context-sensitive term gives perspectival information. For instance, words such as ‘this’, ‘I’, ‘you’, ‘here’, ‘now’ all contain some sort of implicit reference to the perspective of the given speech act that they appear in.

We think that as an empirical matter Neale’s claim is dubious. The fact is that many expressions whose semantics is obviously not constant such as ‘he’ and ‘it’ do not seem to have anything perspectival about them (perhaps explaining why they can be bound by quantifiers). Anticipating this thought, Neale argues that such third-person pronouns indicate that the referent is neither the speaker or the

³⁵ This is not to say that the quantificational approach to definite descriptions is correct. Only that our pre-theoretical intuitions do not provide very good evidence for or against it.

hearer, which thereby gives them perspectival content. However, this kind of perspectival nature is presumably not essential to the semantics of the pronouns but probably just arises from their competition with other forms that do have a perspectival nature. That this is the case is indicated by the fact that one can, if awkwardly, speak truly about oneself or a conversational partner in the third person, however one cannot speak truly about someone else using the first person. For example, one person, say Tom, can use (24) to make a (potentially) true predication of himself, whereas Tom cannot use (25) to speak about another person, Jerry:

(24) Tom, he's a Jet.

(25) Jerry, I'm a Jet.

Further evidence for this (suggested by an anonymous reviewer), comes from the fact that when the identity of a the referent a pronoun is unknown the third person is always used:

(26) (talking about a boy in a family photograph) He must be either me or my brother.³⁶

Indeed it is a rather standard idea in the linguistics literature that, with regard to both person and gender, "he" is an unmarked pronoun (e.g., Sauerland, forthcoming).

Even if we accept that standard indexical forms encode perspective that would by no means show that words like 'red' cannot be indexicals. For it is not clear why positing that a form is indexical needs to be backed up by an explanation of their indexicality. It is rather a piece of accepted data that words like 'red' are used to do different things in different contexts. Giving these words an indexical semantics is a way of handling the data within a truth-conditional semantic theory. It does not seem that any further motivation or explanation for this move

³⁶ Another possible example of unmarked forms in natural language occurs with number. Plural marking is often taken to be the unmarked form of number. This explains why (vi) does not mean the same as (vii).

(vi) There are no baseballs.

(vii) There aren't more than one baseballs.

is necessary. There is no reason to think a constant (non-indexical) semantics should be the default--choosing a constant semantics is as much of a theoretical move as choosing an indexical semantics is. Rather, like every other semantic proposal, our account must compete against other proposals. There is no obvious sense in which our account is deficient because the use of an indexical treatment was unmotivated by any perspectival features.

Szabó (2001) suggests a rather different argument for rejecting a proposal on which color-adjectives are treated as simple indexicals. He writes:

If ‘green’ were an indexical like ‘I’, we would expect the dictionary to contain an informative clause which tells us how to select its content in a given occasion of its use. If ‘green’ is context-dependent, its context-dependence is of a different kind.

One way of understanding this passage is to take it as meaning that any indexical word (i.e. word whose meaning is context-dependent) requires a recipe for understanding how its extension varies.³⁷ Now it is obviously true that for ‘I’ a very simple rule that generally holds: any use of ‘I’ refers to the producer of the utterance it appears in. However, the rule is less simple with ‘that’ and ‘he’. A word like ‘he’ does not even need to pick out demonstrated individuals, or even always the most salient individual. Likewise one cannot easily phrase a simple rule for what ‘that’ picks out. If one tries, the suggested rule has little explanatory power. Here is an example:

[‘That’-rule] ‘that’ picks out the object demonstrated by the speaker.

A rule of this ilk is clearly going to be very weak. For demonstration can happen in countless different ways. Obviously formulating a rule with any predictive power will be very hard.

³⁷ Szabó (p.c.) explained to us that in this passage he actually only meant to prove that ‘red’ is not an indexical in the narrow sense, like ‘I’, ‘you’, and ‘here’, not that ‘red’ is not a context sensitive expression (or indexical in the broad sense). We thought it worth keeping this discussion since the way we interpret his argument seems to be similar to thoughts others have on this matter.

Different context-sensitive expressions work in different ways. These differences do not show, however, that a large class of these expressions are not, in their semantic hearts, all indexicals. Clearly different indexicals must encode different, looser or tighter, lexical constraints on what they can pick out. But there is no motivation for saying that the constraints must be perspectival (as Neale does) or that they must be easily codifiable in a rule of some sort.

Szabó and Neale at least try to identify features that some indexicals share and argue that lexical items not containing those features are not indexicals. A more common attitude is simply to produce a short list of recognized indexical terms, ‘that’, ‘I’, ‘you’..., and say that anything not on that list is not an indexical. The argument, if there is one, tends to be that since speakers only intuitively recognize the small list of indexicals to be context-sensitive then we should only treat that list as indexicals. Our reply is that speakers’ meta-linguistic intuitions (even as reflected by their uses of indirect discourse, as discussed in Cappelen and LePore, 2005) are simply not an adequate guide to what the shape of natural language semantics really is.

The case for this point is not hard to make. There are many abstract syntactic/semantic properties that clearly are relevant to syntactic/semantic processing that ordinary speakers have no conscious grasp of. An obvious example is downward-entailingness (Fauconnier, 1975, Ladusaw, 1979). Downward-entailingness was ‘discovered’ through semantic research to be relevant for the licensing of *negative polarity items*, expressions like ‘any’ and ‘ever’. But even though it took researchers some time to identify the property, it is clear that implicitly we must be sensitive to it, as native speakers are easily able to judge whether a use of negative polarity item is felicitous or not. Anyone working in syntax and semantics must recognize that folk theories of semantics do not constrain real theories. The fact that ‘red’ does not appear in the canonical list of indexicals is no reason not to treat it as one.

6 A Taxonomy of Context Sensitivity

We'll close with a tentative consideration of the taxonomy of context-sensitive expressions. There are, at least, two types of context-sensitive expressions: those like pronouns and tense which can be bound by other operators and those such as 'I' and 'now' and 'red' (on the current proposal) which always have their extension determined by the context of the speech act. This simple, dichotomous picture is essentially the one that Stanley (2002) gives.

My own view of the truth-conditional role of context is very conservative. First, there are expressions which are obviously indexicals in the narrow sense of the term, words such as 'I', 'here', 'you', 'now', and their brethren. Secondly, there are expressions which are obviously demonstratives, such as 'this' and 'that'. Third, there are expressions that are obviously pronouns, such as 'he' and 'she'. Overt expressions that are in none of these classes are not context-dependent. If the truth-conditions of constructions containing them are affected by extra-linguistic context, this context dependence must be traced to the presence of an obvious indexical, demonstrative, or pronominal expression at logical form, or to a structural position in logical form that is occupied by a covert variable.

(400)

Stanley here actually gives four different categories of context-dependent items: indexicals, demonstratives, pronouns and covert variables. However, the difference between demonstratives and indexicals does not seem particularly important for our taxonomic purposes (and Stanley makes nothing of it in the paper). Nor does the distinction between pronouns and covert variables: Stanley treats covert variables as unpronounced pronouns, they can either have free uses determined by context or they can be bound by higher-up operators. The interesting hypothesis then, is that variables in context-sensitive expressions are just silent pronouns. Assuming a unified treatment of indexicals and demonstratives, this would yield an elegant dichotomy. Since we treat 'red' as falling into the first category, of demonstratives and indexicals, this picture is

attractive for us also. However, we will argue in this section that things are not actually so simple.

The dichotomous picture is problematic in a number of ways. First, consider Stanley's characterization of standard indexicals:

The three central features of such words is, first, that they are primitive lexical items, second, that they are not bindable by operators, and, third, that their interpretation shifts from context to context.

In fact, standard indexicals, unlike 'red', do exhibit certain bound readings. First, in some languages such as Amharic, indexicals can be bound by the subject of a propositional attitude attribution. Schlenker (2003) attributes this to the fact that standard indexicals actually get their value determined by reference to a context variable which propositional attitudes operators quantify over. Moreover, Heim (in unpublished work reported in, e.g., Schlenker, 2003) gave some examples where indexicals seem to need to be treated like bound variables:

(27) Only I did my homework.

(28) I did my homework and John did too.

In both cases a bound variable analysis of the first person pronoun seems necessary to get one of the readings.

So, perhaps 'real' indexicals such as 'I' do not have a simple indexical semantics as Stanley suggests. If this is the case we would still, absent further evidence, suggest a plain indexical analysis of 'red', making at least two categories of indexicals, words like 'red' that exhibit mere unbindable context-sensitivity, and words like 'I' that have a more complex semantics (to be discovered by further, cross-linguistic research). So we suggest that a) Stanley might mischaracterize indexicals, and b) words like 'red' should be considered as part of the semantic class of unbindable indexicals/demonstratives.

What about the other half of Stanley's taxonomy, pronouns and variables? We have already reduced the work that the class of variables needs to do, since we do

not think it needs to cover all context-sensitive expressions that are not canonical indexicals/demonstratives. But what about words like ‘enemy’ and ‘local’ that do exhibit bindable context-sensitivity? The interesting proposal here, advocated also by Marti (2003), is that these forms of context-sensitivity are due to variables that essentially act like unpronounced pronouns.

Stanley offers some support for such a claim by arguing that, like pronouns, variables are subject to syntactic constraints such as weak cross-over. We tested this hypothesis by checking some examples on ourselves and others in small email survey. Our survey indicated that in fact the context-sensitivity of expressions such as ‘local’ and ‘enemy’ did not exhibit the same weak-crossover constraints as regular pronouns do.³⁸ First of all, we found Stanley’s own data unconvincing.³⁹

(29) Her_i local bar sponsored every reporter_i.

(30) A local bar sponsored every reporter (where the bar is local to each reporter)
Stanley labels both (29) and (30) as ungrammatical. However, almost everyone we surveyed found (29) to be awful, and (30), if flawed at all, merely inelegant or slightly awkward. It is easy to find other examples where (non-pronominal) contextually-sensitive expressions don’t show weak cross-over effects. For example, weak cross-over effects are extremely weak if present at all in implicitly bound uses of ‘enemy’:

(31) A political enemy sabotaged each congressman’s campaign (understood as an enemy of the congressman)

(32) His political enemy sabotaged each congressman’s campaign.

While (32) is clearly awful, (31) sounds quite good. This leads us to believe that positing covert pronouns to associate with bindable context-sensitive expressions is not syntactically well-motivated. Without further study it is not even clear that

³⁸ For a similar verdict with different examples see Carlson and Storto (forthcoming).

³⁹ These examples are from Stanley (2002), p. 423.

the non-indexical context dependent expressions such as ‘enemy’ and ‘local’ form a unified class.

This does not leave us with a very simple picture of context-sensitivity. Here is a start of a taxonomy, though: First, there are the classical indexicals, which seem to exhibit some limiting binding phenomenon (even in English), second, there are the truly unbindable indexicals (such as ‘here’ and the simple demonstrative ‘that’) to which category we add words like ‘red’, third, there are the normal pronouns, and fourth there are other expressions, such as ‘enemy’ and ‘local’, that seem neither to be straightforward indexicals (owing to their binding behavior) nor expressions with silent pronouns. It would be desirable to give principled explanations of the differences (and/or deeper connections) between these categories, but that is beyond the scope of this paper.⁴⁰

References

- Borg, Emma (2004). *Minimal Semantics*. Clarendon Press.
- Burge, Tyler (1974). Demonstrative constructions, reference, and truth. *Journal of Philosophy*, 71:205–223.
- Cappelen, Herman and LePore, Ernest (2005). *Insensitive Semantics*. Blackwell.
- Carlson, Greg and Storto, Gianluca (forthcoming). Sherlock Holmes was in no danger.
- In Betty Birner and Gregory Ward (eds.), *Drawing the Boundaries of Meaning: Neo-Gricean Studies in Pragmatics and Semantics in Honor of Laurence R. Horn*. John Benjamins.
- Carston, Robyn (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. OUP.
- Chomsky, Noam (1977). *Essays on Form and Interpretation*. Elsevier.

⁴⁰ We are grateful to Chris Barker, Jon Barton, Eliza Block, Richard Breheny, Noam Chomsky, Mike Gabbay, Jen Hornsby, Keith Hossack, Ruth Kempson, Chris Kennedy, Shalom Lappin, Guy Longworth, David Papineau, Richard Samuels, Brett Sherman, Jason Stanley, Zoltán Szabó, Mark Textor, Charles Travis and two anonymous reviewers for many helpful comments on previous drafts. We would also like to acknowledge audiences at the CSMN/Arché Adjectives Workshop and the UCL Pragmatics Reading Group for their useful questions and comments.

- (2003). Reply to Ludlow. In L. Antony and N. Hornstein (eds.), *Chomsky and his Critics*. Blackwell.
- DeRose, Keith (forthcoming). Gradable adjectives: A defence of pluralism. *Australasian Journal of Philosophy*.
- Evans, Gareth (1977). Pronouns quantifiers and relative clauses (i). *Canadian Journal Of Philosophy*, 7(3):467–536.
- Fara, Delia Graff (2000) Shifting Sands: An Interest-Relative Theory of Vagueness, *Philosophical Topics* 28: 45-81.
- Fauconnier, Giles (1975). Polarity and the scale principle. *Chicago Linguistics Society*, 11:188–199.
- Harman, Gilbert (1972). Logical form. *Foundations of Language*, 9:38–65.
- Hawthorne, John (forthcoming). Context dependency and comparative adjectives. *Analysis*.
- Heim, Irene (1982). *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst.
- (1990). E-type pronouns and donkey anaphora. *Linguistics and Philosophy*, 13:137–177.
- Heim, Irene and Kratzer, Angelika (1998). *Semantics in Generative Grammar*. Blackwell.
- Higginbotham, Jim (1985). On semantics. *Linguistic Inquiry*, 16:547–593.
- Horn, Larry (1989). *A Natural History of Negation*. CSLI.
- Kamp, Hans (1981). A theory of truth and semantic representation. In J. Groenendijk et al. (ed.), *Formal Methods in the Study of Language*, pages 277–322. Mathematish Centrum.
- Kaplan, David (1977/1989). Demonstratives. In J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan*. Oxford University Press.
- Klein, Ewan (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4(1):1–45.
- Ladusaw, William (1979). *Polarity Sensitivity as Inherent Scope Relations*. Ph.D. thesis, University of Texas Austin.
- Lahav, Ran (1989). Against compositionality: the case of adjectives. *Philosophical Studies*, 55:111–129.
- Larson, Richard and Segal, Gabriel (1995). *Knowledge of Meaning*. MIT Press, Cambridge, MA.
- Ludlow, Peter (1989). Implicit comparison classes. *Linguistics and Philosophy*, 12:519-533.
- Lycan, William (1984). *Logical Form in Natural Language*. MIT Press.
- Marti, Luisa (2003). *Contextual Variables*. Ph.D. thesis, University of Connecticut.
- Neale, Stephen (1990). *Descriptions*. MIT Press, Cambridge, MA.
- (2004). This, that, and the other. In Marga Reimer and Anne Bezuidenhout (eds.), *Descriptions and Beyond: An Interdisciplinary Collection of Essays on Definite and Indefinite Descriptions*. Oxford University Press.
- (forthcoming). On location. In M O'Rourke and C. Washington (eds.), *Situating Semantics: Essays In Honor of John Perry*. MIT Press.

- Partee, Barbara (1989). Binding implicit variables in quantifier contexts. In C. Wiltshire, B. Music, and R Graczyk (eds.), *Papers from CSL 25*. Chicago Linguistics Society.
- Pinker, Steven (2007) *The Stuff of Thought*. Allen Lane: an imprint of Penguin Books.
- Recanati, Francois (2003). *Literal Meaning*. Cambridge University Press.
- Reimer, Marga (2002). Do adjectives conform to compositionality. *Philosophical Perspectives*, 16:183–198.
- Sauerland, Uli (forthcoming) On the semantic markedness of phi-features. Proceedings of the Phi-Workshop, McGill University, Montreal, Canada.
- Schlenker, Philippe (2003). A plea for monsters. *Linguistics and Philosophy*, 26:29–120.
- Segal, Gabriel (forthcoming). Five flies in the ointment: some challenges for traditional semantic theory. In *Prospects for Meaning*, Richard Schantz (ed), Oxford University Press, Oxford.
- Sperber, Dan and Wilson, Deidre (1986). *Relevance: Communication and Cognition*. Blackwell.
- Stanley, Jason (2002). Context and logical form. *Linguistics and Philosophy*, 25:391–434.
- (2005). Semantics in context. In G. Peters and G. Preyer (eds.), *Contextualism*, pages 221–253. Oxford University Press.
- Stanley, Jason and Szabó, Zoltan (2000). On quantifier domain restriction. *Mind and Language*, 15:219–261.
- Szabó, Zoltan Gendler (2001). Adjectives in context. In I. Kenesei and R.M. Harnish (eds.), *Perspectives on Semantics, Pragmatics, and Discourse*. John Benjamins.
- Travis, Charles (1994). On constraints of generality. *Proceedings of the Aristotelian Society*, 44:165–88.
- (1997). Pragmatics. In Bob Hale and Crispin Wright (eds.), *Blackwell Companion to the Philosophy of Language*, pages 87–107. Blackwell.
- Weinstein, Scott (1974). Truth and demonstratives. *Nous*, 8:179–184.