

A Plea for Semantic Localism

Agustín Rayo
Massachusetts Institute of Technology

October 17, 2011

The purpose of this paper is to defend a conception of language that does not rely on linguistic meanings, and use it to address the Sorites and Liar paradoxes.

I Two Models

I would like to contrast two different answers to the following question:

What sort of semantic information must a subject associate with the members of a language's basic lexicon in order to master the language?

Both answers will be a little oversimplified. But I hope that the contrast between them will nonetheless shed light on the question of what language-mastery consists in.

Here is the first of the two proposals:

I. THE SPECIALIZED-KNOWLEDGE MODEL

In order to count as mastering a language, the subject must be in possession of a significant amount of specifically linguistic information. In particular, her usage of the language's basic lexicon must be guided by *semantic rules* which determine the range of application of a given expression (relative to a context of utterance) independently of the subject's general-purpose abilities, such as common-sense.

With the lexical item ‘blue’, for instance, the subject might associate a semantic rule to the effect that ‘blue’ should be applied to all and only objects with thus-and-such hues, saturations and luminosities (relative to a given context of utterance).

The alternative is to think that language-mastery does not require that a speaker have access to semantic rules that determine the range of application of basic lexical items independently of the speaker’s general-purpose abilities. By exercising sensitivity to context and common sense, the right kind of subject in the right sort of context might be in a position to come to a sensible decision about what to treat as the range of application of a given expression for the purposes at hand. But the specifically linguistic information that the subject associates with the relevant lexical item will typically not suffice to fix its range of application. (See (Travis, 1985; Recanati, 1989; Carston, 2002; Wilson and Sperber, 2004; Stalnaker, typescript); see also (Chomsky, 2000a) and the essays in Part II of (Preyer and Peter, 2004).)

The second of the two proposals I would like to consider is a particular implementation of this general idea:

2. THE GRAB-BAG MODEL

With each expression of the basic lexicon, the subject associates a ‘grab bag’ of mental items: memories, mental images, pieces of encyclopedic information, pieces of anecdotal information, mental maps, and so forth. With the expression ‘blue’, for example, a subject might associate two or three particular shades of blue, the information that a paradigmatic instance of ‘blue’ is the sky on a clear day, a memory of a blue sweater, and so forth. Different speakers might associate different grab bags with the same lexical item.

A grab-bag will typically not be enough to determine a range of application for the relevant lexical item independently of the subject’s general-purpose abilities. But, by exercising sensitivity to context and common sense, the right kind of subject in the

right kind of context might be in a position to use the grab bag to come to a sensible decision about what to treat as the expression's range of application for the purposes at hand.

In some cases—such as the logical connectives—a grab bag might contain something akin to a semantic rule, which is enough on its own to settle a range of application. And in some cases—such as explicitly defined technical terms—the grab-bag might contain the mental analogue of a piece of paper on which an object-language definition of the relevant lexical item has been written.

The following pair of examples (adapted from (Fara, 2000)) can be used to illustrate the differences between our two models:

CASE 1

You know that the party is at one of the two houses in front of you, but you're not sure which. You ask your companion, and she asserts: 'The party is at the blue house'. The house to the left is gray; the house to the right is blueish gray.

On the basis of your companion's assertion, you come to believe that the party is at the blueish gray house (which is exactly what your companion intended).

CASE 2

As before, except that this time the house to the left is bright blue, and the house to the right is blueish gray.

On the basis of your companion's assertion, you come to believe that the party is at the bright blue house (which is exactly what your companion intended).

What is interesting about these examples is that in one case the assertion is used to pick out the blueish-gray house and in the other the assertion is used to exclude it, even though the same sentence is asserted on both occasions.

Let us see how each of our models would explain the fact that you were able to use your companion's assertion to come to a belief about the location of the party. On the Specialized-Knowledge Model, you start by applying the semantic rule corresponding to 'blue' to the context at hand and get a range of application for 'blue'. (Relative to the context of Case 1, bright blue houses and blueish gray houses count as satisfying 'blue', but solid gray houses do not. Relative to the context of Case 2, bright blue houses count as satisfying 'blue' but blueish gray houses and solid gray houses do not.) Once you know the range of application of 'blue' (and other lexical items) relative to the context of the assertion, you use a grammar for the language to assign *truth-conditions* to the assertion.

There is then a straightforward story to tell about how you come to form a belief about the location of the party. You work on the assumption that you and your companion speak the same language, and therefore use the same semantic rules to associate truth-conditions with assertions. So—by taking your companion to be truthful—you come to the conclusion that she believes that the world is such as to satisfy the truth-conditions that you assigned to the assertion. Finally—by treating your companion as reliable—you come to the conclusion that the world is such as to satisfy these truth-conditions.

On the Grab Bag Model, in contrast, most of the work is done by exercising sensitivity to context and common sense. A detailed version of the story will have to wait until section 2.1, but let me give you a first approximation.

When your companion asserts 'The party is at the blue house' it is a bit as if she handed you a grab bag containing a sample of some particular shade of paradigmatic blue. You know that your companion is using the color sample to send you a message. You also know that the issue under discussion is the location of the party, and that your companion is being fully cooperative. So a good guess is that she is trying to use the color sample to indicate the location of the party.

There are two salient possibilities: (i) that the party is at the house to the left, and (ii) that the party is at the house to the right. So you try to determine whether one of these possibilities is rendered salient by the color sample. In Case 1, the house to the left is gray and the house to the right is blueish gray. Neither of the houses is very close in color to your sample, but the blueish gray house is a significantly better fit than the gray house, and this is enough to render the relevant possibility salient. So you conclude that the party is at the blueish gray house. In case 2, the house to the left is bright blue and the house the right is blueish gray. Both houses have colors that bear some similarity to the color of your sample, but the bright blue house is a significantly better fit than the grayish blue house, and this is enough to render the relevant possibility salient. So you conclude that the party is at the bright blue house.

A slightly better approximation would be to imagine that when your companion asserts ‘The party is at the blue house’ it is as if she handed you grab bags corresponding to each of the lexical items in the asserted sentence, where each of the grab bags contains an assortment of different items. It would be even better to imagine that your companion hands you the grab bags plus a bit of ‘syntax’: an explanation of how the grab bags should be combined to render one of the relevant possibilities salient. Better still: imagine that your companion draws attention to some of *your own* grab bags, and hands you some syntactic information.

Later in the paper I will discuss some advantages of embracing the Grab Bag Model instead of the Specialized-Knowledge Model, but one advantage is immediate. On the Specialized-Knowledge Model, a huge amount of specialized semantic information must be somehow stored in a subject’s cognitive system before she will count as mastering the language. For the subject must know semantic rules corresponding to every basic expression in the lexicon, and each semantic rule determines the full range of application of a given lexical item across possible worlds, relative to a large range of contexts of utterance. In contrast, the Grab Bag Model is relatively frugal when it comes to specialized semantic information.

The subject must associate a set of mental items with each basic expression of the lexicon. But nothing more is required as far as the lexicon is concerned. The rest of the work is done by the subject's grasp of grammar, and by her general-purpose abilities.

2 Localism

The Grab Bag Model is intuitively appealing. But you might worry that there is no real reason to think that the sketchy—and largely metaphorical—ideas that have been presented so far could be developed into a theory detailed enough to allow for adequate assessment. The purpose of this section is to convince you otherwise.

The key idea is *localism*. The best way of understanding what localism amounts to is by contrasting it with globalism. Globalism is the claim that in order for an assertion to be in good order, it must succeed in determining a partition of the entire space of possibilities. On the simplest version of globalism, the partition is two-fold: possibilities are divided into *verifiers* and *falsifiers*. (There are versions of globalism whereby the space of possibilities is divided into more than two groups, but I will ignore them here to simplify the exposition.) Consider a particular assertion of 'the party is at the blue house' as an example. A globalist would say that there must be a definite range of hues, saturations and luminosities such that the assertion is verified at a possibility p if and only if, at p , the party in question takes place at a house whose color falls within the relevant range.

In contrast, a localist would claim all that is required for an assertion to be in good order is for it to succeed in dividing *the possibilities that are relevant for the purposes of the assertion* into verifiers and falsifiers. Consider an assertion of 'the party is at the blue house', and suppose the context is as in Case 1 of the preceding section. A localist would say that all it takes for the assertion to be in good order is for it to be clear whether or not the assertion would be verified at each of the two possibilities that are salient in the case at hand: (*i*) that

the party be at the gray house, and (ii) that the party be at the blueish gray house. There need not be a fact of the matter about whether the sentence would be verified at a possibility in which the party takes place at a blueish-green house.

A proponent of the Specialized-Knowledge Model might be tempted by globalism. She might be tempted to think that, once contextual parameters have been fixed, the language's semantic rules determine truth-conditions that divide the entire space of possibility into verifiers and falsifiers. But set the Specialized-Knowledge Model aside and there is no real reason to think that globalism is true. For language is first and foremost a tool for communication, and successful communication does not normally require discriminating amongst possibilities that are not relevant for the purposes at hand.

The Grab Bag Model gains plausibility when localism is in place. For a syntactically articulated assortment of grab bags might well be enough to enable a contextually sensitive and commonsensical subject to discriminate amongst the possibilities that are relevant for the purposes of the assertion even if it wouldn't enable her to discriminate amongst *arbitrary* possibilities. Go back to Case 1. When the purposes at hand require only that we decide between the possibility that the party is at the gray house and the possibility that it is at the blueish gray house, a grab bag consisting of a particular color sample was enough for our subject to get the job done. But no amount of sensitivity to context and common sense would enable the subject to use that color sample to generate a principled division of the *entire* space of possibility into verifiers and falsifiers.

2.1 The Hermeneutic Procedure

In this section I will suggest a localist implementation of the Grab Bag Model. I will work on the assumption—familiar from (Stalnaker, 1979)—that conversation takes place against the background of a *context set*. A context set is a set of (mutually exclusive) possibilities that a conversational participant treats as *live options* for the purposes of a given stage of the

conversation. For a possibility to be treated as a live option is for it to be compatible with every proposition that the participant treats as a *presupposition* of the relevant stage of the conversation. A proposition is treated as a presupposition by a conversational participant just in case she is disposed to act, for the purposes of the relevant stage of the conversation, as if the proposition was common knowledge. As the conversation evolves, so does each participant's context set.

There may be differences between the context sets corresponding to different conversational participants. Such differences can be tolerated as long as they do not interfere with the goals of the conversation. But if differences become significant enough to threaten conversational goals, participants can be expected to adjust their presuppositions in such a way that discrepancies are reduced. (To simplify the exposition, I will speak of *the* context set of a conversation whenever differences between the context sets corresponding to different conversational participants are irrelevant for the purposes of our discussion.)

From a localist point of view, *the main objective of semantic theorizing is to explain and predict the evolution of the context set throughout a conversation.* The Grab Bag Model is a view about the sort of information that a subject must associate with a language's lexical items in order to count as a competent user of the language. As such, it has nothing explicit to say about the evolution of the context set. But it can nonetheless have a role to play in explaining context set dynamics. For in taking a stand about the sort of information that competent speakers have at their disposal in performing and interpreting speech acts, it imposes constraints on the ways in which speakers are able to use a speech act to modify the context set. What I mean when I say that I will suggest a localist implementation of the Grab Bag Model is that I will suggest an account of the evolution of the context set that relies on a theory of linguistic competence based on the Grab Bag Model.

Before telling you about the story in detail, let me remind you of the basic picture. When an assertion takes place, the hearer will try to use the assertion to come up with a hypothesis

about how the speaker proposes to divide the context set into verifiers and falsifiers. In doing so, she will make use of specifically semantic information. But not much: the lion's share of the work will be done by sensitivity to context and common sense.

Now for the details. When an assertion takes place, the hearer's cognitive system carries out some variant of the following *Hermeneutic Procedure*:

1. Use your overall understanding of the situation at hand to form a hypothesis about which possibilities are being treated as live options by the speaker.

[*Example*: Consider a particular assertion of 'the party is at the blue house': case 1 from section 1, say. A reasonable hypothesis is that the speaker is treating the following two possibilities as live: (i) that the party be at the gray house, and (ii) that the party be at the blueish gray house.]

2. Use your grasp of the language's phonology, morphology and syntax to form a hypothesis about which sentence has been asserted.

[*Example*: You hypothesize that your companion's assertion is built up from the lexical items <'the', 'party', 'is', 'at', 'the', 'blue', 'house'>, and that they are meant to be articulated into a sentence with a certain syntactic structure.]

3. Use your competence with the language's lexicon to associate a grab bag of mental items with each basic lexical item occurring in the sentence identified in step 2.

[*Example*: You associate a grab bag with each of 'party', 'blue', 'house', 'the', 'is' and 'at'. In the case of 'party', 'blue' and 'house', the grab bags can be expected to consist of a rich array of mental items, such as mental images and pieces of encyclopedic information. In the case of 'the', 'is' and 'at', the grab bags can be expected to consist of semantic rules.]

4. Assess whether the grab bags delivered by step 3 succeed in rendering any objects or

properties salient given your general understanding of the situation at hand. In light of your conclusions, use common sense to choose a *reference* for lexical items occurring in the sentence asserted.

[*Example:* Given what you know about the situation at hand, the party you are trying to get to can be expected to be rendered salient by the grab bag corresponding to ‘party’, and each of the two houses in front of you can be expected to rendered salient by the grab bag corresponding to ‘house’. In addition, the color of the blueish-gray house—but *not* the color of the gray house—can be expected to be rendered salient by the grab bag corresponding to ‘blue’.

This makes it reasonable to choose the following referents for ‘party’, ‘house’ and ‘blue’:

- Relative to each of the two possibilities hypothesized as live in step 1, the extension of ‘party’ is the set consisting of the relevant party.
- Relative to each of the two possibilities hypothesized as live in step 1, the extension of ‘house’ is the set consisting of the two relevant houses.
- Relative to each of the two possibilities hypothesized as live in step 1, the extension of ‘blue’ is the set consisting of the blueish-gray house.

Finally, referents for ‘the’, ‘is’ and ‘at’ are selected on the basis of the relevant semantic rules while displaying sensitivity to the situation at hand.]

5. Evaluate the sentence identified in step 2 with respect to the assignment of referents proposed in step 4. Then determine whether, so interpreted, the sentence succeeds in generating a division of the possibilities identified in step 1 that is *complete, non-trivial* and *sensible*:

- (a) For the division to be *complete* is for each of the possibilities in question to be counted as a verifier or a falsifier.
 - (b) For the division to be *non-trivial* is for at least one possibility to be counted as a verifier and one possibility to be counted as a falsifier.
 - (c) For the division to be *sensible* is for it to be plausible in light of your overall understanding of the situation at hand that the possibilities that are counted as falsifiers are possibilities the speaker might propose to rule out in the context at hand and for the possibilities that are counted as verifiers to be possibilities the speaker might propose not to rule out in the context at hand.
6. If step 5 delivers a complete, non-trivial and sensible division of possibilities, go on to step 7. If not, try modifying the assumptions that you made in carrying out steps 1–4.

[*Example:* Suppose that your initial hypothesis was that the speaker asserted ‘the party is at the *new* house’. Since both houses seem pretty new and neither of them appears significantly newer than the other, step 5 delivers the *trivial* division of possibilities whereby both possibilities are counted as verifiers. You go back to step 2, and see what happens when you work on the assumption that the speaker said ‘blue’ instead of ‘new’.]

An important kind of revision you may need to carry out corresponds to what Lewis (1979) calls ‘accommodation’. You *accommodate* when you revise the set of possibilities you treat as open by assuming that the speaker is making a presupposition that has not been made explicit but would deliver a result in a more sensible way of interpreting her assertion.

[*Example:* Suppose you initially had no idea that a party was taking place. In order to interpret the speaker’s assertion of ‘the party is at the blue house’, you hypothesize that the speaker is presupposing that a party will be taking place, and therefore exclude

any non-party possibilities from the set of possibilities you take her to be treating as live.]

Another important kind of revision you may need to carry out corresponds to what Stalnaker (1979) calls ‘diagonalization’. (I omit discussion for the sake of brevity.)

If your best effort to revise earlier hypotheses fails to deliver a complete, non-trivial and sensible division of possibilities, try working on the assumption that the speaker is to be understood non-literally (e.g. ironically), or that she is hoping to create a conversational implicature.

If that doesn’t work, give up: conclude that you are unable interpret the assertion (and, if contextually appropriate, ask for clarification from the speaker.)

7. If you’ve made it this far, you have succeeded in identifying a complete, non-trivial and sensible division of possibilities. Now test for straightforwardness and uniqueness:

- (a) *Straightforwardness*

Verify that there there is no immediate reason to think that the speaker had access to a significantly more straightforward way of proposing the division of possibilities that has already been identified. (If she did, make sure you have a reasonable explanation for why she was being less than straightforward: she might have hoped to be funny, or to come across as highly educated, or to not be understood by the children in the room.)

- (b) *Uniqueness*

Verify that there isn’t an alternate set of assumptions for steps 1–4 which isn’t significantly less plausible than your current assumptions but delivers a rival division of possibilities.

If straightforwardness and uniqueness are satisfied, you’re done: work on the assump-

tion that the speaker is proposing that the context set be modified in accordance with the division of possibilities that has already been identified.

Otherwise, you can try revising your assumptions once again. Or you can give up: you can conclude that you're not sure how to interpret the assertion (and, if contextually appropriate, ask for clarification from the speaker.)

2.2 The Principle of Clarity

The Hermeneutic Procedure is an account of how the *hearer* succeeds in interpreting the assertion. But it can also be used to understand why the *speaker* acts as she does. For in a well-run conversation speaker and hearer cooperate to secure the successful transfer of information. So the speaker can be expected to comply with the following principle:

PRINCIPLE OF CLARITY

Make your assertion in such a way that it is clear to your audience how you propose that it be used to modify the context set.

But the Grab Bag Localist would claim that the speaker is, in effect, presupposing that her audience will attempt to carry out something along the lines the Hermeneutic Procedure. So she can be expected to comply with the Principle of Clarity by choosing her assertion in such a way that sensible application of the Hermeneutic Procedure is likely to deliver the proposal for modifying the context set that she would like to set forth.

2.3 How to Assemble a Grab Bag

The Hermeneutic Procedure supplies a way of understanding how an assignment of grab bags to lexical items can be used to interpret the speaker's assertion. But it also gives us a way of understanding how language users determine which mental items to include in their

grab bags. For it is often clear from context how the speaker proposes to modify the context set by making a particular assertion. In such cases the hearer can work backwards, and use the Hermeneutic Procedure to make an educated guess about the sorts of referents that the speaker might have intended the hearer to associate with lexical items in the asserted sentence for the purposes of the assertion. That in turn can be used to form a hypothesis about the sorts of mental items that it might be useful to include in a grab bag for the relevant lexical item. The subject can then gauge the effectiveness of her hypothesis on the basis of whether future communication is smooth: whether it seems to her that she is able to successfully interpret assertions involving the relevant lexical item, and whether it seems to her that her own assertions have the intended effect on others. And, of course, she can continue to make adjustments as time goes on.

If this story is along the right lines, there is no reason to expect that different speakers will associate the same grab bag of mental items with a given lexical item. But one should expect that the differences won't normally interfere with successful communication: even if different grab bags are used by different speakers, the participants in a conversation will typically end up with a similar set of possibilities divided in a similar way.

2.4 Meager Contexts

Attention so far has been focused on the simplest case: verbal assertions in which speaker and audience have a rich array of contextual cues to help them along the way. But linguistic communication can also take place in more meager contexts. One might, for instance, write a text without knowing much about who one's audience will be: one may not know when or where it will be read, or if readers of the text will know much about its author.

The picture I have been defending predicts that it will be much more difficult to communicate effectively in such contexts. One will often need to 'create a context' by using earlier sentences to generate contextual cues that can be used to interpret later ones. And one may

need to fend off misunderstanding by using several sentences where a richer context might have called for only one.

Of course, that is just what one finds in practice. Small children can often communicate effectively in person, but it usually takes a more sophisticated language user to communicate effectively over the phone. A brief email is often all it takes to communicate effectively with a friend, but it takes skill to write an good letter to someone one has never interacted with. *Pride and Prejudice* is able to communicate effectively across cultures and generations. But that is part of what makes Austen great. She is able to create enough of a context for later assertions to be readily interpretable, and she is able to fend off much potential misunderstanding. Still: the more one knows about Austen, and about the customs of the gentry in pre-Victorian England, the smoother communication becomes.

If the context is meager enough, one may be unable to come to a firm conclusion about what particular division of possibilities is being proposed by the speaker. The most one can do is form hypotheses about what might have been meant on different assumptions about what was being presupposed by the speaker at the time of the assertion. If one needs to know exactly what was meant, one has no choice but to ask for clarification, or try to find out more about the contextual assumptions that the speaker may have been relying on. Fortunately, one doesn't always need to know exactly what was meant in order to satisfy one's goals: it is often enough to know what the speaker might have intended relative to a few reasonable hypotheses about the conversational presuppositions. And it is important to keep in mind that even contexts that are relatively meager contain vast amounts of information that is common ground between speaker and audience—information that can be invisible because it is so obvious. Even the most culturally removed of Austen's readers knows that humans fall in love.

2.5 Truth

On the localist picture I have been defending it makes good sense to ask of an assertion whether it is true or false. An assertion is used by the speaker to propose a modification of the context set. The context set consists of disjoint possibilities, so at most one of them can be actualized. If the proposal is such as to keep the actualized possibility in the context set, the assertion is true. If the proposal is such as to eliminate the actualized possibility from the context set, the assertion is false.

In some cases conversational participants will have made a false presupposition somewhere along the way, and the actualized possibility has been excluded from the start. This makes it more difficult to assign truth-values to assertions in a principled way, but the following procedure is useful for some purposes. Imagine that you—the theorist—are eavesdropping on the conversation. You make no false presuppositions, so your version of the context set includes the actual possibility. But you otherwise do your best to have your version of the context set reflect the spirit of the conversation. When the assertion takes place, you try to interpret it by carrying out the Hermeneutic Procedure. In the special case in which you are able to identify a salient proposal for modifying your version of the context set on the basis of the assertion, the assertion can be counted as true if the proposal in question is such as to keep the actualized possibility in your version of the context set, and otherwise counted as false. (For a more sophisticated set of tools for addressing this issue, see (Yablo, 2006, forthcoming; Schoubye, 2009).)

Our characterization of truth and falsity for *assertions* can be used to give a characterization of truth and falsity for *sentences*. A sentence is true relative to an assertoric context just in case an assertion of that sentence in that context would count as true. (And similarly for falsity.)

2.6 Reference

On the localist picture I have been defending, lexical items don't usually have a reference independently of an assertoric context. The only semantic information concerning a particular lexical item that is available to speakers independently of an assertoric context is a grab bag. And unless the grab bag happens to consist of a semantic rule—as it will in certain special cases—it will fall well short of determining a reference for the term. A grab bag for 'elephant', for example, might consist of a mental image of an elephant, a few encyclopedic entries ('elephants are animals', 'elephants are big and gray, and have trunks'), and a few memories ('I once rode an elephant when I was a child'). Such a collection of mental items does not, by itself, determine a referent for 'elephant'. In fact, it is not even clear what it would mean to say that an image, some encyclopedic entries and a few memories *determine* a referent. All the grab bag can do is help render a potential referent salient to the right sort of subject, in the right sort of context.

For an assertion to be *successful* by the lights of the localist—for it to comply with the Principle of Clarity—it is not necessary that different conversational participants associate the same range of application to lexical items occurring in the sentence asserted. All that is required is that rival ranges of application generate the same division of the possibilities under consideration. Consider case 1 from section 1 as an example. Speaker and hearer can associate different ranges of hue, saturation and luminosity with 'blue' for the purposes of the assertion, and still agree that the possibility whereby the party is at the blueish gray house is to be kept in the context set, and that the possibility whereby the party is at the gray house is to be excluded. In fact, an individual subject doesn't even have to choose a specific range of application for 'blue'. As long as she decides which of the contextually relevant houses are to count as instances of 'blue', she will be in a position to identify the proposed division of the possibilities at hand.

From our current perspective, reference is nothing more than a tool. It is a device for

getting from an assertion to a division of the possibilities at hand. There is therefore no sense to be made of the question of whether a subject's assignment of reference is correct over and above the question of whether it gets the job done: whether it supplies a successful method for generating the requisite division of possibilities.

2.7 Linguistic Meaning

Philosophers tend to think that the expressions of a language like English are endowed with *linguistic meaning*, a contextually invariant semantic property with two central features. The first is compositionality: the linguistic meaning of a complex expression should be determined by its grammatical structure, together with the linguistic meanings of its constituent parts. The second is a link with truth-conditions: the linguistic meaning of a sentence should determine, relative to a given context, *what is said* by that sentence in that context.

It is important to distinguish between an expression's linguistic meaning and its *competence-meaning* (relative to a subject): the information possession of which enables the subject to count as a competent user of that particular expression. The Grab Bag Model has a lot to say about competence-meaning—and specifically about the competence-meaning of the members of the language's basic lexicon—but it has nothing to say about linguistic meaning. Notice, in particular, that talk of linguistic meaning—and talk of 'what is said' by a sentence in context—is entirely absent from the Hermeneutic Procedure. According to the localist implementation of the Grab Bag Model that has been developed here, *the notion of linguistic meaning does not have a central role to play in explaining and predicting the evolution of the context set*. (This is a substantial view, with which many philosophers disagree; see, for instance, (Stanley, 2000).)

A friend of the Specialized-Knowledge Model will see things very differently. She can be expected to think that there is a close connection between the linguistic meaning of an expression and its competence-meaning. The simplest view would be that they coincide:

an expression's linguistic meaning is a semantic rule, and it is precisely the semantic rule that a subject would need to associate with the expression in order to count as a competent speaker of the language. (On a more sophisticated version of the view, linguistic meaning would stand to competence-meaning as secondary intensions stand to primary intensions.) In addition, a friend of the Specialized-Knowledge Model can be expected to think that linguistic meaning plays a central role in explaining and predicting the evolution of the context set. To a first approximation, she will wish to claim that an assertion modifies the context set by eliminating any worlds that are incompatible with the truth-conditions that are determined by the linguistic meaning of the sentence asserted relative to the context of the assertion.

Even though it sidelines the notion of linguistic meaning, the implementation of Grab Bag Model that has been defended here has a lot to say about meaning in context. We have seen that, relative the right sort of context, one gets interesting notions of truth and reference. But it is worth noting that one also gets an interesting notion of *propositional* content. For in the right kind of context, a sentence can be used to propose a modification of the context set. So a proponent of the Grab Bag Model can identify the proposition expressed by a sentence in context with the set of possibilities that would be left in the context set after the proposed modification.

On this way of seeing things, the proposition expressed by a sentence in context should be understood as proposing a *local* division of the possibilities. In other words: a possibility *in the relevant context set* can be counted as a verifier or falsifier of the sentence in context depending on whether it is a member of the proposition expressed, but the sentence in context should be understood as taking no stand with respect to possibilities outside the context set. It is also worth noting that the proposed characterization of propositional content does not rely on a distinction between what is said by a sentence in context and what is communicated, or between literal and non-literal discourse. This is advantageous in some

ways, since the distinctions in question can be hard to pin down. But it does mean that the proposition expressed is not guaranteed to correspond to what one would ordinarily think of as the literal meaning of the asserted sentence. It will always correspond to the proposed modification of the context set.

Grab Bag Localism—as I shall hereafter refer to view that I have been defending—has a number of advantages. We will discuss some examples in the next few sections.

3 Vagueness

According to the Grab Bag Localist, a vague term like ‘blue’ can apply to different objects as used in different contexts. As used in case 1 from section 1, for example, ‘blue’ applies to the blueish-gray house; but not as used in case 2. In this minimal sense, the Grab Bag Localist might be described as a *contextualist* about vagueness. (For classic statements of contextualism see (Soames, 1999; Fara, 2000); for further discussion see (Åkerman and Greenough, 2009).)

Contextualists face an important challenge:

THE PROBLEM OF CUTOFF-POINTS

Say you have a Sorites-series of tiles: the first tile is a clear case of blue, the last tile is a clear case of green, and a normal human would find any two adjacent tiles indistinguishable in color.

Fix a particular context of utterance. It is natural to suppose that, as used in that context, ‘blue’ applies to some of the tiles but not others. What is the nature of the transition from cases of application to cases of non-application? Is there a sharp cutoff-point? (If so, what explains why the cutoff-point of ‘blue’, as used in that particular context, is located precisely where it is?)

Should one say, instead, that ‘blue’ applies to greater and lesser degrees, as used in the relevant context? (If so, is there a sharp cutoff-point between tiles to which ‘blue’ applies in full, as used in that context, and tiles to which it does not?) Should one distinguish between ‘definite’ applications of ‘blue’ and the rest? (If so, is there a sharp cutoff-point between tiles to which ‘blue’ definitely applies, as used in the relevant context, and tiles to which it does not?)

In this section I will argue that the Grab Bag Localist’s version of contextualism has a satisfactory answer to this challenge. I will also suggest a Localist answer to the Sorites Paradox. (For related discussion, see (Rayo, 2008).)

The Problem of Cutoff-Points

Unlike more traditional forms of contextualism about vagueness, Grab Bag Localism is not committed to the view that there is a linguistic meaning corresponding to ‘blue’. The only semantic information that is stably associated with ‘blue’ is a grab bag. And all the grab bag can do is help render a potential referent salient to the right sort of subject, in the right sort of context. When the context is such that the subject is faced with a Sorites-series of tiles ranging from blue to green—and when there are no further contextual cues—the subject will be unable to use her grab bag to render any particular division of the tiles salient. So she will be unable to divide the tiles into those that are to be counted as instances of ‘blue’ and those that are not. (This is evidenced by the fact that an assertion of ‘I want you to bring me the blue tiles and leave the rest on the table’ would fail to comply with the Principle of Clarity: in the absence of further contextual cues, the listener would be unable to determine which of the various collections of tiles that might be brought to the speaker would be such as to satisfy the speaker’s desires.)

Even in the case of a *successful* assertion—an assertion that complies with the Principle of Clarity—the speaker may be unable to use her grab bag to divide contextually relevant

objects into those that are to be counted as instances of ‘blue’ and those that are not. Suppose the topic under discussion is whether a Sorites-series in the next room transitions from blue to green or blue to purple. If the speaker were to say ‘the series transitions from blue to green’, the assertion can be expected to be successful. For it should be clear to the listener that the speaker is proposing to modify the context set by excluding the possibility whereby the series transitions from blue to purple and retaining the possibility whereby the series transitions to blue to green. Yet there is no reason to suppose that the listener will be in a position to divide the *tiles* that figure in the possibilities under consideration into those that are to be counted as instances of ‘blue’ and those that are not. She doesn’t need to in order to work out how the speaker proposes to modify the context set, since any sensible assignment of reference to ‘blue’ will result in the same way of dividing the *possibilities* under consideration into verifiers and falsifiers. And all it takes for the Principle of Clarity to be satisfied is for a division of the *possibilities* under consideration to be rendered uniquely salient to participants in the conversation.

The upshot is this. Independently of a particular assertoric context, it makes no sense to talk about the range of application of ‘blue’, and therefore no sense to talk about a cutoff-point for ‘blue’. In the right kind of assertoric context, the right kind of subject might be able to use her grab bag for ‘blue’ to divide contextually relevant objects into those that are to be counted as instances of ‘blue’ for the purposes of the assertion and those that are not. But this will certainly *not* be the case when it comes to a soritical context of the kind that is the focus of the present section. Accordingly, the Grab Bag Localist should respond to the Problem of Cutoff-Points by claiming that it makes a false presupposition: in the relevant context, there is no such thing as the supposed transition from instances of ‘blue’ to non-instances.

What about non-Soritical contexts? Won’t ‘blue’ sometimes be associated with a range of application? And wouldn’t the Grab Bag Localist then have to account for the transition

from instances of ‘blue’ to non-instances? Yes. But in such cases there is no reason to think that accounting for the transition will be a problem. To see this, go back to our stock example: Case 1 from section 1. All the listener needs to do to interpret the assertion in that context is decide which of the two houses to count as instances of ‘blue’ for the purposes of the assertion. And now there is a clear explanation of why the speaker will count the blueish-gray house as an instance of ‘blue’, but not the gray house. For such a classification of contextually relevant objects into instances and non-instances of ‘blue’ can be expected to stand out as *significantly more salient* than any rival, in light of the grab bag corresponding to ‘blue’ and general information about the situation at hand. This is in stark contrast with a typical Soritical context, where no single classification of tiles into instances and non-instances of ‘blue’ will stand out as uniquely salient because any potential location for the cutoff-point will seem roughly as salient as its immediate neighbors.

Localism is playing a big role here. When only the possibilities that are relevant for the purposes of the assertion need to be divided into verifiers and falsifiers, and not the entire space of possibility, there is typically no need classify every single point in the color spectrum as an instance or non-instance of ‘blue’. In most ordinary cases, it is only necessary to classify a few, contextually relevant objects. (Compare (Fara, 2000).) And one’s grab bag corresponding to ‘blue’—together with general information about the situation at hand—will often be enough to render a classification of such objects uniquely salient. As emphasized in section 2.5, the Grab Bag Localist thinks of reference as a *tool*: it is a device for getting from the assertion to a division of the possibilities at hand. So there is no reason to think that there is more to be said about the reference of a particular lexical item in a given assertoric context than is needed to make clear how the speaker proposes to divide the possibilities at hand.

Higher-Order Vagueness?

Consider the sentence ‘I want you to bring me the blue tiles and leave the rest on the table’. There are contexts in which an assertion of this sentence would satisfy the Principle of Clarity, and context in which it would not. (Compare, for example, a context in which there are two paradigmatically blue tiles on the table and two paradigmatically green ones, and a context in which there is a Sorites-series of tiles on the table.)

One might worry that this could be used to generate a problem of higher-order vagueness for the Grab Bag Localist. Consider a sequence of contexts, all of which involve an assertion of ‘I want you to bring me the blue tiles and leave the rest on the table’. In the first context, there are only paradigmatically blue and paradigmatically green tiles on the table. But successive contexts include additional intermediate tiles, until we get a context in which there is a Sorites-series of tiles on the table, ranging from paradigmatically blue to paradigmatically green. Don’t we face a higher-order version of the Problem of Cutoff-Points? The Grab Bag Localist wants ‘successful’ to apply to the assertion in the first member of the sequence, but not to the assertion in the last member of the sequence. So one would hope that she would have something informative to say about nature of the transition from cases of application to cases of non-application. And little progress would be made if she ended up postulating some sort of arbitrary cutoff-point.

Fortunately, the Grab Bag Localist does no such thing. For an assertion to be successful is for it to conform to the Principle of Clarity—for it to make clear to its audience how the speaker proposes that the context-set be modified. But there is no hidden fact of the matter about how much clarity is needed for this requirement to be satisfied. Whether or not to count an assertion as successful is a *practical* matter. It is *decision* on the part of participants in the conversation: the assertion needs to be clear enough for the purposes at hand, by the lights of the people involved. In the course of a conversation, speaker and hearer hope to strike a balance between the cost of risking a misunderstanding and the cost of burdening the

conversation with too many requests for clarification, and they can use whichever standard of clarity they feel comfortable with as a means for achieving a suitable compromise.

On a more traditional form of contextualism, in contrast, one would be committed to the claim that the sentence under consideration has a linguistic meaning. So there is room for arguing that the success of an assertion should be determined by *semantic* considerations. One might suggest, for example, that an assertion is to be counted as successful just in case the linguistic meaning of the asserted sentence delivers well-defined truth conditions with respect to the context of the assertion—which would leave us just where we started, since one might wonder whether there is a sharp cutoff-point separating contexts with respect to which the sentence’s linguistic meaning delivers well-defined truth-conditions, and context with respect to which it does not.

The Sorites Paradox

Say we have a Sorites-series of tiles. The first is a clear case of blue, the last is a clear case of green, and a normal human would find any two adjacent tiles indistinguishable in color. Then the following sentences all appear to be true:

- (1) The first tile is blue.
- (2) The last tile is not blue.
- (3) If a given tile is blue, so is its successor.

But (1)–(3) are jointly inconsistent.

The Grab Localist thinks that (1) and (2) are both true, as used in canonical contexts. Suppose, for example, that you are asked to describe the Sorites series in question to someone who has not seen it. You begin your description by asserting ‘the first tile is blue’ and ‘the last tile is not blue’. Even though your assertions will not succeed in rendering salient a

particular range of application for ‘blue’, they can certainly be expected to comply with the Principle of Clarity. And—in the absence of false presuppositions—the proposed partition of the context set would be such as to count the actualized possibility as a verifier.

The situation with respect to (3) is more complex because (3) can be used to convey very different pieces of information, even when attention is restricted to canonical soritical contexts. Suppose that it is common knowledge that you and your interlocutor are both competent logicians, and that you both interpret bare conditionals as material conditionals. Then it should also be common knowledge that (3) will be regarded as equivalent to ‘either every tile in the series is blue, or no tile is’. An assertion of (3) in such a context will certainly comply with the Principle of Clarity under such circumstances. But—unlike an assertion of (1) or (2)—it can be expected to count the actualized possibility as a falsifier, since it will leave your interlocutor with the impression that the sequence of tiles is not a Sorites series at all, if it includes any blue tiles.

Note, however, that a sentence can be used to say something about *language* even if it contains no linguistic vocabulary. ‘A baker’s dozen is 13’, for instance, can be used to communicate the thought that the expression ‘baker’s dozen’ is used to pick out collections of 13 objects. Similarly, (3) can be used to communicate something about the expression ‘blue’ even though it contains no linguistic vocabulary. Consider, for example, a context in which people familiar with the Sorites series in question are debating the range of application of ‘blue’. In such a context, (3) might be used to communicate the following claim:

MEANING-TOLERANCE

If ‘blue’ applies to a given tile in the series, it must also apply to the tile’s immediate neighbors.

And (3)’s negation might be used to communicate the claim that ‘blue’ has a sharp cutoff-point: that there is a tile such that ‘blue’ applies to it but not to one of its neighbors. (For

related discussion, see (Tappenden, 1993).)

According to the Grab Bag Localist, both Meaning-Tolerance and its negation should be rejected. For they both presuppose that ‘blue’ has a range of application in the context at hand, and the Grab Bag Localist thinks that there is no such thing.

There is, however, something pre-theoretically appealing about Meaning-Tolerance. And unless one has a story to tell about why this is so, one hasn’t fully accounted for the Sorites Paradox. Fortunately, the Grab Bag Localist is especially well-placed to tell such a story. Consider the following principle:

USE-TOLERANCE

In a typical Soritical context, a competent speaker will only be in a position to discern reasons for applying ‘blue’ to a given tile if she is in a position to discern comparable reasons for applying ‘blue’ to the tile’s immediate neighbors.

Use-Tolerance is clearly true. It is also a natural consequence of Grab Bag Localism. To see this, suppose that the only semantic information competent speakers of the language associate with ‘blue’ is a grab bag consisting of two or three mental ‘samples’ of paradigmatic shades of blue. In the right kind of context, the right kind of speaker might be able to use such a grab bag to render a particular range of application salient. But an assortment of mental samples is a pretty coarse tool. In the absence of special contextual cues, it cannot be used to discriminate amongst contextually relevant objects unless one of them is significantly bluer than another.

It is easy to make the mistake of going from the unimpeachable Use-Tolerance to the disastrous Meaning-Tolerance if one is under the grip of a certain conception of language: the idea that linguistic competence is a matter of gaining access to the linguistic meanings of our words and learning to apply them in the right sorts of ways. For—as emphasized in (Wright, 1976)—this picture makes it natural to suppose that one can uncover the linguistic meanings

governing our language by introspection, merely in virtue of mastering the language. The Grab Bag Localist thinks of matters very differently. Linguistic competence is not usually a matter of deploying specifically linguistic information; it consists in the ability to use one's grab bags to make sensible decisions about how to partition the context set in light of the situation at hand.

4 The Liar Paradox

The aim of this section is to propose a localist account of the English word 'true', and use it to develop a picture of what goes on when an English speaker asserts the Strengthened Liar Sentence:

(λ) λ is not true.

4.1 Two Challenges

Consider the following line of reasoning:

Suppose, for *reductio*, that λ is true. Then the world is as λ says it is. But λ says that λ is not true, contradicting our hypothesis. So, by *reductio*,

λ is not true.

But wait! I've just asserted the very sentence that has been shown to be untrue.

So I've committed myself to a claim that is, by my own lights, untrue.

According to *contextualism* about truth, this line of reasoning rests on a mistake. It presupposes that λ has a stable meaning, which is expressed both when λ is asserted in a given context and when λ is asserted in the course of explaining what goes on in that context. The

presupposition leads to trouble because it allows one to fudge the distinction between the unproblematic claim that one has just asserted a *sentence* that one knows not to say something true in some other context, and the disastrous claim that one has just *asserted a proposition* that one knows to be untrue. (For discussion of contextualism, see (Parsons, 1974; Burge, 1979; Glanzberg, 2004).)

Contextualism faces two important challenges. The first is that we need an account of the linguistic mechanism whereby λ comes to express different propositions in different contexts—and we would like the suggested mechanism to have some sort of independent motivation.

The second challenge concerns what is sometimes referred to as *revenge phenomena*. To see what's at issue, assume the contextualist has been granted the claim that λ expresses different thoughts in different contexts. One can still consider the question whether:

(ξ) No assertion of ξ is true.

is true. And it is not clear that the contextualist is in a position to give a stable answer to this question. (For a collection of recent essays on revenge phenomena, see (Beall, 2007).)

In what follows I will argue that both of these challenges can be addressed by a version of contextualism that is informed by Grab Bag Localism.

Here is the proposal, in rough outline. According to the Grab Bag Localist, one doesn't usually interpret an assertion by implementing a semantic rule. Instead, one uses the Hermeneutic Procedure to try to come up with a sensible decision about how to partition the context set in light of the situation at hand. The Hermeneutic Procedure leads to a number of constraints on the subject's interpretation of a particular assertion of λ . The first challenge will be addressed by noting that satisfaction of these constraints can result in different interpretations for different assertions of λ . The second challenge will be addressed by noting that there are contexts in which it is not possible to satisfy all the relevant constraints.

A Warm-Up Case

It will be useful to start with a warm-up case, which I borrow from the questionnaire at the end of (Cartwright, 1987). Suppose that an English speaker makes an assertion α by producing a sequence of phonemes of the kind that is naturally associated with following sentence:

The last word in assertion β is obscene.

But quotation marks are not usually pronounced in English. So α could also be interpreted as corresponding to the following sentence:

The last word in assertion β is ‘obscene’

or to the following:

The last word in assertion β is ‘ ‘obscene’ ’

If we knew more about β , considerations of charity might be able to help us decide between these different interpretations of α . In particular:

- If it turns out that the last word in β is obscene, charity demands that we take the last word in α to be ‘obscene’.
- If it turns out that the last word in β is ‘obscene’, charity demands that we take the last word in α to be ‘ ‘obscene’ ’.
- If it turns out that the last word in β is ‘ ‘obscene’ ’, charity demands that we take the last word in α to be ‘ ‘ ‘obscene’ ’ ’.
- And so forth.

This leads to the following constraint on our interpretation of α :

Hermeneutic Constraint

If last word in β sounds like ‘obscene’, considerations of charity demand that the last word of α be the result of surrounding the last word of β by quotation marks.

But now suppose that α and β turn out to be one and the same. Then there is a second constraint on our interpretation:

External Constraint

Since $\alpha = \beta$, the last word of α is identical to the last word of β .

And, since the last word of α sounds like ‘obscene’, there is no way of satisfying both of these constraints at the same time.

When α and β are assumed to be numerically distinct, however, the problem goes away. Considerations of charity will still be deployed to interpret α . So the Hermeneutic Constraint will remain in place. But it doesn’t lead to trouble, because it is perfectly stable in the absence of the External Constraint.

The lesson of our exercise is that numerical identity matters for interpretability. One can have two assertions that are perfect duplicates of each other with respect to a crucial set of properties—phonetic profile, in this case—but differ significantly in terms of interpretability because only one of them is subject to the External Constraint. The same sort of observation will be at the heart of our discussion of the Strengthened Liar Sentence. One can have two assertions of λ that are perfect duplicates of each other with respect to a crucial set of properties, but differ significantly in terms of interpretability because only one of them is subject to an analogue of the External Constraint.

A Truth-Predicate for Grab Bag Localists

As noted in section 2.7, the Grab Bag Localist steers clear of the notion of linguistic meaning. In the absence of such a notion, there is a natural picture of sentential truth that ceases to be available. One cannot say that for a sentence to be true (relative to a set C of contextual parameters) is for the truth-conditions determined by its linguistic meaning (relative to C) to be satisfied.

Fortunately, the Grab Bag Localist has a different way of thinking about sentential truth. She can characterize sentential truth by using *assertoric* truth: a sentence is true relative to an assertoric context just in case an assertion of that sentence in that context would count as true. And, of course, the Grab Bag Localist's characterization of assertoric truth does not presuppose a prior characterization of sentential truth. As noted in section 2.5, the Grab Bag Localist can say that an assertion is true if it is used by the speaker to propose a modification of the context set which is such as to keep the actualized possibility as a live option. Moreover, getting such a proposal across to the listener does not require a notion of sentential truth, or a notion of linguistic meaning, since speaker and hearer coordinate by way of the Hermeneutic Procedure.

This picture of sentential truth suggests a particular way of modeling our usage of the English word 'true'. The proposal is based on a *simulation theory*. In order to interpret 'ϕ is true', one starts by simulating an assertion of ϕ in a context that seems sensible in light of the situation at hand; one then takes 'ϕ is true' to be true just in case the simulated assertion would succeed in proposing a modification of the context set which would keep the actualized possibility as a live option.

If this story is along the right lines, a sensible grab bag for 'true' might consist of an instruction of the following kind:

Apply 'true' to a sentence ϕ just in case an assertion of ϕ in whichever context

seems most salient in light of the situation at hand would propose a division of possibilities that would keep the actualized possibility as a live option.

4.2 The First Challenge

If the preceding section is along the right lines, what we need to do in order to interpret the sentence ‘ λ is true’ is simulate an assertion of λ in the right kind of context.

Let us begin by considering the canonical example—an assertion α of λ in a context C_α such that the following two conditions are satisfied:

1. The context that is most salient in C_α is C_α itself.
2. Three possibilities are treated as live in C_α :

t: α is true: it succeeds in proposing a modification of the context set that would keep the actualized possibility as a live option.

f: α is false: it succeeds in proposing a modification of the context set that would rule out the actualized possibility as a live option.

d: α is defective: it fails to propose a suitable modification of the context set.

The Grab Bag Localist can show that λ fails to propose a suitable modification of the context set as used in C_α . For there are interpretative pressures coming from two different directions:

- *Hermeneutic Constraint*

One’s interpretation of λ relative to C_α is constrained by the Hermeneutic Procedure, and therefore by one’s grab bag for ‘true’.

- *External Constraint*

One’s interpretation of λ relative to C_α is constrained by one’s views about which of *t*, *f* and *d* is actualized. (If one assumes, for example, that *t* is actualized, one has thereby assumed that λ is interpreted so as to count *t* as a verifier.)

And these two constraints turn out to be at odds with each other: they cannot both be satisfied at the same time. (See appendix for details.)

The upshot is that λ is defective as used in C_α . And, of course, this entails that, as used in C_α ,

λ is not true.

But wait! I have just asserted λ . Moreover, the context of my assertion shares two crucial features with C_α : (1) the context that is treated as most salient is C_α , and (2) the possibilities that are treated as live are t , f and d .

In spite of all this, there is no reason to think that *my* assertion is defective. For even though my assertion and α are assertions of the same sentence, are aimed at dividing the same set of possibilities, and treat the same context as maximally salient, *they need not be numerically identical*. And that makes all the difference. For if the assertions are distinct, only α is concerned with itself (even though they are both concerned with α .)

Self-reference matters. According to the Grab Bag Localist, any assertion will be interpreted by following something along the lines of the Hermeneutic Procedure. So any assertion will be subject to a version of the Hermeneutic Constraint. But the External Constraint is unusual. It only comes into the picture when the relevant assertion is its own subject-matter. This unusual condition is satisfied by α , but need not be satisfied by my assertion. So one of the interpretative pressures that led to trouble in the case of α need not be present in the case of my assertion.

In our warm-up case, it would have been wrong-headed to ask: “Forget about the various assertions and focus on the sequence of phonemes [the-last-word-in-assertion-beta-is-obscene]. How should the sequence itself be interpreted?” For it only makes sense to interpret the sequence of phonemes in the context of a particular assertion. Similarly, in the case of the Liar Sentence it would be wrong-headed to ask: “Forget about the various assertions

of λ and focus on λ itself. How should the *sentence* be interpreted?” According to the Grab Bag Localist, there is no reason to think that λ has a linguistic meaning, and therefore no reason to think that it makes sense to talk about λ 's interpretation independently of the Hermeneutic Procedure.

Let us take stock. We have seen that, from the perspective of the Grab Bag Localist, different assertions of λ can differ in terms of interpretability even if they aim to partition the same set of possibilities and treat the same context as maximally salient. When the context that is treated as maximally salient is identical to the context of the assertion, λ is uninterpretable. But λ can be interpreted (and, indeed, interpreted as true) when the salient context and the context of the assertion are numerically distinct. The result is a form of contextualism because λ is taken to lack a stable interpretation across contexts. But, unlike more traditional forms of contextualism, the Grab Bag Localist doesn't explain λ 's contextual variation by appeal to a context-sensitive linguistic meaning. And, unlike more traditional forms of contextualism, the Grab Bag Localist is in a position to give an independently motivated account of the linguistic mechanism that underlies λ 's contextual variation: it flows naturally from the Grab Bag Localist's conception of truth, together with the Hermeneutic Procedure.

4.3 The Second Challenge

Our second challenge is to address revenge phenomena. The issues here are very delicate, so I will begin with a simplified version of the problem.

The Simplified Version

Our discussion of λ in the preceding section relied on an assumption. We presupposed that the *diagnostic context*—the context from which one carries out one's diagnosis of a particular assertion of λ —can be numerically distinct from the *subject-matter context*—the context

which one talks about in providing one's diagnosis. This assumption is not particularly problematic in the case of λ . But let us now consider the following sentence:

(μ) None of Marvin's assertions of μ is true.

Those of us who are numerically distinct from Marvin have no problem diagnosing Marvin's assertions of μ . On suitable assumptions, Marvin's assertions of μ must all be defective. So we can conclude that

none of Marvin's assertions of μ is true.

Poor Marvin, on the other hand, will never be in a position to articulate such a diagnosis. For that would involve asserting μ and, as we have seen, every one of Marvin's assertions of μ will be defective. The problem, of course, is that Marvin has no place to stand. He doesn't have access to a diagnostic context which is numerically distinct from each of the subject-matter contexts one would need to allude to in supplying the relevant diagnosis.

It seems to me, however, that there is nothing mysterious about Marvin's predicament. It is a bit like being engaged in the project of doing something Marvin will never do. This is a project with respect to which Marvin will never be successful. But there is no mystery: a game has been set up, and it's built into the rules that Marvin can't win.

Here is a closer analogy. Consider the project of using considerations of charity to interpret an assertion α based on the sequence of phonemes [the-last-word-in-your-favorite-assertion-is-obscene]. If α happens to be Marvin's favorite assertion, we have a project that Marvin will never be successful at. But, again, there is no mystery: a game has been set up, and it is built into the rules that someone with preferences like Marvin's can't win.

According to the Grab Bag Localist, Marvin is in a similar situation when he engages in the project of diagnosing one of his assertions of μ . An interpretative game has been set up, but this time interpretation is governed by more than just considerations of charity:

it is based on the Hermeneutic Procedure. As in the case of the analogy, there will be various interpretative constraints. In most cases the constraints work together to single out an interpretation for the assertion under consideration. But sometimes they come into conflict and it becomes impossible to assign the assertion a stable interpretation. In the analogy, conflicting constraints make it impossible for Marvin to come up with a stable interpretation for his favorite assertion; in the case of μ , the conflicting constraints make it impossible for Marvin to come up with a stable interpretation for an assertion of μ . But that's all there's to it. A game has been set up, and the rules make it impossible for Marvin to win. There is nothing mysterious or paradoxical going on.

A Different Project

When philosophers talk about the Liar Paradox, they sometimes think of it as a *logic puzzle*. (Leitgeb, 2007) articulates a nice description of this sort of approach. He suggests that it would be ideal to characterize a theory of truth satisfying every condition on the following list. (Please don't worry if you don't understand everything.)

1. The notion of truth should be expressed by a predicate.
2. If the theory of truth is added to mathematical or empirical theories, it should be possible to prove that the latter theories are true.
3. The truth predicate should not be subject to type-restrictions.
4. T-biconditionals—sentences of the form '*s* is true if and only if *p*', where '*p*' is replaced by a sentence and '*s*' is replaced by a name of that sentence—should be derivable unrestrictedly.
5. The theory of truth should be compositional.
6. The theory should allow for standard interpretations.

7. The ‘outer’ logic and ‘inner’ logic should coincide.
8. The ‘outer’ logic should be classical.

But, as Leitgeb shows, 1–8 are jointly inconsistent. So to ‘solve’ the Liar Paradox, on this way of seeing things, is to find a theory of truth that comes as close as possible to satisfying 1–8 (or show that we can live with inconsistency).

A lot of good work has gone into this kind of project. (My personal favorites are (Kripke, 1975; McGee, 1990; Field, 2008).) But this is emphatically *not* the spirit in which the Liar Paradox is being addressed here. I am *not* out to develop a theory of truth in Leitgeb’s sense. Our present discussion concerns the usage of the English word ‘true’. And, as far as I can tell, attention to our actual linguistic practice gives us no good reason to think that an account of the natural-language truth-predicate should be based on a Leitgeb-style list of constraints. (Notice, in particular, that the fourth of Leitgeb’s constraints appears to presuppose that ‘true’ has a context-insensitive linguistic meaning.)

A Leitgeb-style conception of truth would be a logician’s dream. But it is important not to confuse the project of developing a logician’s dream with the project of understanding how our actual language works. By failing to make a clear separation between the two one risks seeing paradox where there is none. It is easy to make that mistake if one is under the grip of a certain conception of language: the idea that linguistic competence is a matter of gaining access to the linguistic meanings of our words and learning to apply them in the right sorts of ways. For this sort of picture makes it natural to suppose that we can uncover the linguistic meaning governing the English truth-predicate by introspection, merely in virtue of mastering the language. (See (Wright, 1976).) And one might think that something along the lines of Leitgeb’s list could be justified on the basis of such introspection. The Grab Bag Localist sees things very differently. Linguistic competence is not usually a matter of applying our knowledge of linguistic meaning; it consists in the ability to use one’s grab

bags to make sensible decisions about how to partition the context set in light of the situation at hand.

Revenge

A lot has been said about Marvin's predicament. But don't I myself face a similar problem?

Consider:

(ν) None of Agustín's assertions of ν is true.

or, indeed:

(ξ) No assertion of ξ is true.

I would love to say more about a sentence like ξ . Unfortunately, I don't have access to the right sort of diagnostic context. If I did have access to the right context—a context outside the domain of extant assertions of ξ —I would be in a position to supply an accurate diagnosis of such assertions in the new context by asserting 'none of the relevant assertions of ξ are true' and making clear what the relevant assertions are. But I don't. So I won't.

Fortunately—as the case of Marvin makes clear—the fact that one is not in a position to diagnose an assertion doesn't entail that the assertion is paradoxical.

4.4 What has been achieved?

I have sketched a picture of how the word 'true' might work in natural language. The picture begins with a conception of *assertoric* truth according to which an assertion should be counted as true if it proposes a division of possibilities that would allow the actualized possibility to remain in the context set. Sentential truth is characterized on the basis of assertoric truth: a sentence is true relative to a given context if an assertion of that sentence in that context would count as true. This suggests that one might model the workings of the

word ‘true’ by way of a simulation theory. Speakers figure out whether to apply ‘true’ to a given sentence by simulating an assertion of the sentence in a suitable context and trying to determine whether the assertion would count as true.

When this picture is developed against the background of Grab Bag Localism, one gets two interesting results. First, one gets an account of contextual movement in Strengthened Liar contexts: a principled explanation of how a sentence that is uninterpretable in one context can be successfully used to diagnose such uninterpretability in another.

Second, one gets an account of *why* Liar-like sentences are uninterpretable in certain contexts. The Grab Bag Localist’s picture entails that there are a number of different principles constraining interpretation—including the Hermeneutic Procedure—and it turns out that not all of the constraints can be satisfied in the relevant contexts. But that’s all there’s to it. There is nothing mysterious or paradoxical going on. The illusion of paradox comes from assuming that the English word ‘true’ should satisfy constraints that seem introspectively plausible, or that a logician would find attractive. But there is no real reason to suppose that such constraints would be motivated by an account our actual linguistic practice.

5 Linguistics

The project of setting forth linguistic theses to solve philosophical problems carries a serious risk. One might end up doing bad linguistics in the hopes of doing good philosophy.

Grab Bag Localism is partly a linguistic thesis and therefore shouldn’t be accepted unless it is compatible with our best overall linguistic theorizing. But the risk of incompatibility is smaller than one might think. For linguistic theory—and even the branch of linguistic theory that is usually referred to as ‘semantics’—has much more to say about grammar than lexical meaning. And although Grab Bag Localism makes potentially controversial claims about lexical meaning, it makes no claims about grammar: it can simply piggyback on what-

ever our best linguistic theorizing has to say.

More importantly, Grab Bag Localism can be used to address some interesting problems in linguistics. So it could arguably earn its keep on linguistic grounds alone. In this section we will consider some examples. Most of the material will be drawn from an unpublished manuscript of Bob Stalnaker's from the late 1970s: 'The Autonomy of Semantics'.

Singular Reference

Consider the following examples from (Nunberg, 1978) and Stalnaker's manuscript: 'New York won, 10 to 6'; 'I have been reading Agatha Christie'; 'the newspaper you are reading has come out against gun control'; 'the chair you are sitting in is commonly found in 18th century interiors'; 'Yeats dislikes having himself read with an English accent'; 'the *Times* has decided to change its format'.

In canonical assertions of such sentences, the name or description is naturally taken to refer to something other than its paradigmatic denotation. ('New York', for example, is naturally taken to refer to a sports team rather than a city or a state, and 'Agatha Christie' is naturally taken to refer to a literary work rather than a person.) How is a suitable referent to be identified? The Grab Bag Localist is especially well placed to supply an answer. In fact, the answer has already been given: all one needs to do is apply the Hermeneutic Procedure.

Consider the first example on the list: 'New York won, 10 to 6'. The grab bag corresponding to a name like 'New York' can be assumed to consist of mental items that most readily bring to mind the city (or state) of New York. But in a typical assertion of 'New York won, 10 to 6' the audience can be expected to recognize that New York City and New York State are not a good choices of reference for 'New York' because neither of them delivers a sensible division of the open possibilities. As long as the speaker has abided by the Principle of Clarity, however, the audience should be in a position to use the grab bag for 'New York' to identify a different object as salient—a sports team, in this case—and recog-

nize that it constitutes an attractive choice of referent for ‘New York’ on the grounds that it results in an eminently sensible division of the open possibilities. (The salience of the alternative referential hypothesis may well be *parasitic*—the sports team might be salient because it is known to be based in a city which has already been rendered salient. But there are also non-parasitic routes to salience: one could, for instance, rely on the fact that the words ‘New York’ are prominently stitched on the team’s uniform.)

The Grab Bag Localist’s story is independently motivated, and has considerable intuitive appeal. It also has the advantage of not presupposing a distinction between literal and non-literal discourse. This is important because the distinction is difficult to pin down, but also because the resulting proposal makes room for an unusually economical account of non-literal discourse. Consider, for example, ‘that gazelle over there is my nephew’, as used to refer to a prancing child. A friend of the Specialized-Knowledge Model of linguistic competence might think non-literal discourse should be interpreted by way of a two-step process. One must first ascertain the literal meaning of the sentence in question, as used in the relevant context, and then go on to compute the assertion’s non-literal meaning by relying on its literal meaning. According to the Grab Bag Localist, in contrast, there is no need for a two step process. When the range of possibilities is suitably constrained—when it is common knowledge, for example, that the speaker is using the assertion to single out one of the children in the park—the hearer’s grab bag for ‘gazelle’ can be sufficient to render the right child salient with no need for intermediate steps. For even if none of the children would be rendered especially salient by the relevant grab bag independently of the relevant context, a prancing child can be expected to be rendered significantly more salient in context than any of the more sedentary children.

Cross-category movement and Polysemy

The word ‘stop’ can move across syntactic categories. It can be used both as a verb (you can stop running) and as a noun (you can come to a stop).

In addition, ‘stop’ is highly polysemous. You can stop writing; you can stop a burglar; you can stop a check; you can stop a nail hole with plaster; you can use your fingers to stop the holes of a flute; you can stop a poker into the fire; you can stop the tide by anchoring your boat. You can come to a stop; there can be a stop in your speech; you can include a stop in a telegram; you can put a stop on a camera; you can pull out all the stops on an organ.

How should one account for such grammatical and semantic versatility? One option is to postulate a complex array of semantic rules (or a few semantic rules exhibiting complex patterns of contextual variation). But Grab Bag Localism delivers a simpler answer. One can place a few key items in one’s grab bag for ‘stop’—for instance, representations that bring to mind interfering, preventing, obstructing and closing—and let common sense and sensitivity to context take care of the rest.

And there is no need to use different grab bags for different grammatical categories. A mental image that evokes obstruction, for example, can be used to render salient the action of closing a valve when interpreting an assertion of ‘she stopped the flow of oxygen’, and to render salient a particular knob on a pipe organ when interpreting an assertion of ‘she moved the stops to control the air-flow into her organ’.

6 Outstanding Issues

There are three outstanding issues I would like to mention before closing. The first concerns intensional contexts. Extant accounts of propositional attitude ascriptions and modal operators tend to appeal to the notion of linguistic meaning or its kin. (See (Richard, 2006;

von Fintel and Heim, 2010).) So one might worry that the Grab Bag Localist will end up committed to linguistic meanings after all.

The way forward, it seems to me, is to model propositional attitude ascriptions and modal operators by way of a *simulation theory*, in a manner reminiscent of (Stalnaker, 1987). Here is an example. Suppose ‘Susan believes that the party is at the blue house’ is asserted in context where two possibilities are regarded as open: the possibility that Susan believes that the party is at the blue house, and the possibility that Susan believes that the party is at the red house. To decide whether to count each of these possibilities as a verifier or a falsifier, one moves to a *derived* context-set, consisting of the possibility that the party is at the blue house and the possibility that the party is at the red house. One then *simulates* an assertion of ‘the party is at the blue house’ with respect to the derived context set, and deploys the Hermeneutic Procedure to determine how the simulated assertion would propose that the derived context-set be partitioned. The result is that the possibility that the party is at the red house is eliminated from the derived context set. Finally, one goes back to the original context set, and rules out any possibilities whereby Susan’s beliefs are compatible with possibilities that have been eliminated from the derived context set. The result is that the possibility that Susan believes that the party was at the red house gets ruled out from the original context set—which is precisely what one would expect.

A lot more work needs to be done before the Grab Bag Localist can claim to have an adequate treatment of belief ascriptions. But I hope I have said enough to make clear that an account of attitude ascriptions and modal contexts that sidesteps the notion of linguistic meaning cannot be written-off from the start. (For relevant discussion regarding epistemic modals, see (Yalcin, 2005).)

The second issue I would like to raise concerns the relationship between Grab Bag Localism and cognitive science. Grab Bag Localism takes a stand on the question of what sort of semantic information needs to be stored in a subject’s cognitive system in order for her

to be competent in the use of the lexical items of her language. And this is just the sort of empirical issue that cognitive science hopes to address.

There are cognitive scientists who think of linguistic competence in ways that are broadly in line with Grab Bag Localism. (See, for instance, (Kamp and Partee, 1995; Glenberg, 1997; Pulvermüller, 1999; Fauconnier, 2004; Pecher and Zwaan, 2005; Martin, 2007; Barsalou et al., forthcoming); see also Wilson and Sperber (2004).) If these sorts of approaches were to prevail in the long run, Grab Bag Localism would be left in a very strong position. For it would supply a model of linguistic competence that makes sense from the standpoint of cognitive science, in addition to being well motivated from a philosophical point of view.

But what if the empirical evidence were eventually to point in a different direction? Although that would certainly be bad news for Grab Bag Localism, it is important to be clear about what it would and would not show. It would show is that Grab Bag Localism is not the right account of how the mind works: it is not an accurate description of the psychological mechanisms whereby competent speakers use language. But Grab Bag Localism could still have *instrumental* value: it could still be a useful way of modeling the evolution of the context set throughout a conversation.

The final issue I would like to raise before closing concerns *thought*. Attention so far has been focused on public language. But it would be nice to have a unified account of linguistic and mental representation. So it would be good to know whether there is any prospect of developing mental analogue of Grab Bag Localism.

Although I myself am cautiously optimistic, it would be misleading to suggest that I have anything like a developed proposal in mind. There are significant differences between linguistic and mental representation, and certain aspects of Grab Bag Localism cannot be carried over to the mental case without further ado. The most obvious example concerns localism. In the case of public language, we have independent reasons for thinking that conversation takes place against the background of a fairly restricted set of possibilities. But

there is no obvious reason to assume that the same is true of thought.

My own view is that the assumption can, in fact, be justified on independent grounds. It seems to me that there are good reasons for thinking that a subject's cognitive system should be modeled as *fragmented*. It makes no sense to ask what a subject believes or desires *simpliciter*: one can only ask what the subject believes or desires *for the purposes of a particular project*. And the nature of the project supplies the needed restriction on the space of relevant possibilities. (See (Stalnaker, 1984, chap. 5).) In any case, there is no denying that significant further work would need to be in place before one could defend a mental analogue of Grab Bag Localism.

7 Conclusion

I have defended Grab Bag Localism. According to *localism*, all that is required for an assertion to be in good order is for it to succeed in the dividing the set of possibilities that are relevant for the purposes of the conversation into verifiers and falsifiers. According to localism of the *grab bag* variety, speakers carry out this division of possibilities on the basis of very meager semantic information. All it takes to be competent with the use of a basic lexical item is to have access to a suitable grab bag. The rest of the work is done by the speaker's knowledge of grammar, and by a combination of sensitivity to context and common sense.

I have argued that Grab Bag Localism can be used to give attractive accounts of the Sorites and Liar paradoxes. The accounts have a common structure. In both cases, Grab Bag Localism delivers a non-paradoxical description of the phenomenon in question. And in both cases Grab Bag Localism can be used to show that the illusion of paradox is encouraged by a misguided picture of the workings of language.

I have also argued that Grab Bag Localism can be used to address a family of stubborn problems in lexical linguistics, and that it could conceivably be seen as earning its keep on

linguistic grounds alone.

8 Epilogue: The Rock and the Sextant

Rocks are powerful tools. So are sextants. Rocks can be used to break windows and prop doors open; they can be used to keep things from being blown away by the wind. Sextants can be used to establish one's latitude at sea.

There is an important difference between what it takes to know how to use a rock as a paper-weight or a door-stop, and what it takes to know how to use a sextant to establish one's latitude at sea. In the case of the rock, all it takes is a general understanding of the situation at hand, and common sense. One doesn't need to undergo specialized rock-training before one can use a rock to keep one's notes from being blown away by the evening breeze. As long as one knows a few basic facts about the rock, and about the situation at hand, one can normally rely on general-purpose knowledge of the sort one uses to navigate the world on a daily basis. (Which is not to say that the task is trivial: think of how difficult it would be to build a robot that was able to use rocks to keep things from being blown away by the wind, and do so under a broad range of circumstances.)

Now consider sextants. In order to know how to use a sextant to establish one's latitude, a typical human would have to rely on more than just a general understanding of the situation at hand, and common sense. Specialized training is required. One needs to learn about a set of specific procedures whereby a sextant can be used to convert information about the positions of various celestial bodies into information about one's latitude.

It is tempting to think of language-use as being more like sextant-use than rock-use. That is the Specialized-Knowledge view: one masters a language primarily by learning a set of language-specific procedures—a set of specifically linguistic 'rules'—that allow for the conversion of information about the physical properties of an assertion into information about

what the assertion says. Sensitivity to context and common sense have a role to play, but their importance is secondary.

It is certainly true that one can't master a language unless one has *some* specialized knowledge: one needs information about the language's grammar, and about how particular lexical items are used. But I have argued that the role of specialized knowledge is much more limited than one might think. When it comes to the usage of particular lexical items, the lion's share of the work is done by sensitivity to context and common sense. Mastering a language is more like knowing how to use a rock to keep things from being blown away by the wind than like knowing how to use a sextant to establish one's latitude at sea.*

Appendix

Here is a detailed explanation of why λ is defective as used in C_α :

Hypothesis 1: the actualized possibility is t .

Suppose the actualized possibility is t . It follows from the definition of t that α succeeds in proposing a modification of the context set that would keep the actualized possibility as a live option. So λ must be true as used in C_α .

But let us now try to *interpret* λ in C_α by using the Hermeneutic Procedure. Since the context that is most salient in C_α is C_α itself, the grab bag for 'true' asks that 'true' be applied to λ just in case α proposes a division of the possibilities that would keep the actualized possibility as a live option—which is precisely

*For their many helpful comments, I am grateful to JC Beall, Philip Bricker, Delia Fara, Hartry Field, Barbara Grosz, Jeremy Hartman, Ricardo Mena, Jennie Pyers, François Recanati, Marcus Rossberg, Josh Schechter, Brad Skow, Bob Stalnaker, Pedro Stepanenko, Eduardo Villanueva, Bruno Whittle and Steve Yablo. I am also grateful to seminar participants at the Institut Jean Nicod and MIT, and to audiences at the Institute for Philosophical Research of the National University of Mexico, the First Annual Meeting of the Latin-American Association for Analytic Philosophy, the MIT Work in Progress Seminar, the Radcliffe Institute for Advanced Study, and the University of Connecticut. Portions of this paper were written during the tenure of an ACLS research grant at the Radcliffe Institute for Advanced Study, for which I am very grateful.

what we concluded in the preceding paragraph. So the grab bag for ‘true’ asks that ‘true’ be applied to λ . This means that the most salient assignment of reference to ‘true’ delivers the result that the sentence ‘ λ is true’ is true as used in C_α , and therefore the result that the sentence ‘ λ is not true’ is not true as used in C_α . But $\lambda =$ ‘ λ is not true’. So we are left with the conclusion that λ is not true as used in C_α , which contradicts the conclusion of the preceding paragraph.

Hypothesis 2: the actualized possibility is f .

Suppose the actualized possibility is f . It follows from the definition of f that α succeeds in proposing a modification of the context set that would exclude the actualized possibility as a live option. So λ must be untrue as used in C_α .

But let us now try to *interpret* λ in C_α by using the Hermeneutic Procedure. Since the context that is most salient in C_α is C_α itself, the grab bag for ‘true’ asks that ‘true’ be applied to λ just in case α proposes a division of the possibilities that would keep the actualized possibility as a live option—which contradicts the conclusion of the preceding paragraph. So the grab bag for ‘true’ asks that ‘true’ not be applied to λ . This means that the most salient assignment of reference to ‘true’ delivers the result that the sentence ‘ λ is true’ is not true as used in C_α , and therefore the result that the sentence ‘ λ is not true’ is true as used in C_α . But $\lambda =$ ‘ λ is not true’. So we are left with the conclusion that λ is true as used in C_α , which contradicts the conclusion of the preceding paragraph.

Hypothesis 3: the actualized possibility is d .

Suppose the actualized possibility is d . It follows from the definition of d that α fails to propose a suitable modification of the context set. So λ must be untrue as used in C_α .

But let us now try to *interpret* λ in C_α by using the Hermeneutic Procedure.

Since the context that is most salient in C_α is C_α itself, the grab bag for ‘true’ asks that ‘true’ be applied to λ just in case α proposes a division of the possibilities that would keep the actualized possibility as a live option—which contradicts the conclusion of the preceding paragraph. So the grab bag for ‘true’ asks that ‘true’ not be applied to λ . This means that the most salient assignment of reference to ‘true’ delivers the result that the sentence ‘ λ is true’ is not true as used in C_α , and therefore the result that the sentence ‘ λ is not true’ is true as used in C_α . But $\lambda =$ ‘ λ is not true’. So we are left with the conclusion that λ is true as used in C_α , which contradicts the conclusion of the preceding paragraph.

Since no hypothesis about which of the open possibilities is actualized allows for a stable interpretation of λ , λ cannot be used in C_α to propose a suitable division of the open possibilities into verifiers and falsifiers. So λ is defective as used in C_α .

References

- Barsalou, L., Santos, A., Simmons, W. K., and Wilson, C.D., forthcoming “Language and Simulation in Conceptual Processing.” In De Vega et al. (forthcoming).
- Beall, JC (ed.), 2007 *The Revenge of the Liar*. Oxford: Oxford University Press.
- Bräuerle, R. et al. (eds.), 1979 *Semantics from Different Points of View*. Springer-Verlag.
- Burge, Tyler, 1979 “Semantical Paradox.” *Journal of Philosophy*, 76, pp. 169–98. Reprinted in Martin (1984).
- Byrne, Alex and Thomson, Judith (eds.), 2006 *Content and Modality*. Oxford: Oxford University Press.
- Carston, Robyn (ed.), 2002 *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell Publishing.
- Cartwright, Richard, 1987 *Philosophical Essays*. Cambridge, MA: MIT Press.
- Chomsky, Noam, 2000a “New Horizons in the Study of Language.” In Chomsky (2000b). pp. 3–18.
- , 2000b *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Davis, S. (ed.), 1991 *Pragmatics: A Reader*. Oxford: Oxford University Press.
- De Vega, M., Glenberg, A., and Graesser, A. (eds.), forthcoming *Symbols, embodiment, and meaning*. oxford: Oxford University Press.
- Devitt, M. and Hanley, R. (eds.), 2006 *Blackwell Guide to the Philosophy of Language*. Oxford: Blackwell Publishing.

- Dietz, Richard and Moruzzi, Sebastiano (eds.), 2009 *Cuts and Clouds: New Essays on Vagueness*. Oxford: Oxford University Press.
- Evans, Gareth and McDowell, John (eds.), 1976 *Truth and Meaning*. Oxford: Clarendon Press.
- Fara, Delia Graff, 2000 "Shifting Sands: An Interest-Relative Theory of Vagueness." *Philosophical Topics*, 28, pp. 45–81. Originally published under the name 'Delia Graff'.
- Fauconnier, G., 2004 "Pragmatics and Cognitive Linguistics." In Horn and Ward (2004).
- Field, Hartry, 2008 *Saving Truth from Paradox*. Oxford: Oxford University press.
- von Fintel, Kai and Heim, Irene, 2010 "Intensional Semantics." Lecture notes, MIT; available at <http://mit.edu/fintel/IntensionalSemantics.pdf>.
- Glanzberg, Michael, 2004 "A Contextual-Hierarchical Approach to Truth and the Liar Paradox." *Journal of Philosophical Logic*, 33, pp. 27–88.
- Glenberg, A.M., 1997 "What memory is for." *Behavioral and Brain Sciences*, 20, pp. 1–55.
- Horn, L.R. and Ward, G. (eds.), 2004 *The Handbook of Pragmatics*. Oxford: Blackwell.
- Kamp, Hans and Partee, Barbara, 1995 "Prototype Theory and Compositionality." *Cognition*, 57, pp. 129–191.
- Keefe, R. and Smith, P. (eds.), 1996 *Vagueness: A Reader*. Cambridge, MA: MIT Press.
- Kripke, Saul, 1975 "Outline of a Theory of Truth." *Journal of Philosophy*, 72, pp. 690–716.
- Leitgeb, Hannes, 2007 "What Theories of Truth Should be Like (but Cannot be)." *Philosophy Compass*, 2(2), pp. 276–290.

- Lewis, David, 1979 "Score-Keeping in a Language Game." *The Journal of Philosophical Logic*, 8, pp. 339–59. Reprinted in Bräuerle et al. (1979); reprinted in Lewis (1983).
- , 1983 *Philosophical Papers, Volume I*. Oxford.
- Martin, A., 2007 "The representation of object concepts in the brain." *Annual Review of Psychology*, 58, pp. 25–45.
- Martin, R. L. (ed.), 1984 *Recent Essays on Truth and the Liar Paradox*. Oxford: Oxford University Press.
- McGee, Vann, 1990 *Truth, Vagueness and Paradox: An Essay on the Logic of Truth*. Indianapolis/Cambridge: Hackett.
- Nunberg, Geoffrey, 1978 *The Pragmatics of Reference*. Bloomington, IN: Indiana University Linguistics Club 1978.
- Parsons, Charles, 1974 "The Liar Paradox." *The Journal of Philosophical Logic*, 3, pp. 381–412. Reprinted in Parsons (1983).
- , 1983 *Mathematics in Philosophy*. Ithaca, NY: Cornell University Press.
- Pecher, C. and Zwaan, R. (eds.), 2005 *Grounding cognition: The role of perception and action in memory, language, and thought*. New York: Cambridge University Press.
- Preyer, Gerhard and Peter, Georg (eds.), 2004 *Contextualism in Philosophy*. Oxford: Oxford University Press.
- Pulvermüller, F., 1999 "Words in the brain's language." *Behavioral and Brain Sciences*, 22, pp. 253–336.
- Åkerman, Jonas and Greenough, Patrick, 2009 "Hold the Context Fixed, Vagueness Still Remains." In Dietz and Moruzzi (2009).

- Rayo, Agustín, 2008 “Vague Representation.” *Mind*, 117, pp. 329–373.
- Recanati, F., 1989 “The Pragmatics of What is Said.” *Mind and Language*, 4, pp. 97–120.
Reprinted in Davis (1991).
- Richard, Mark, 2006 “Propositional Attitude Ascription.” In Devitt and Hanley (2006).
- Schoubye, Anders J., 2009 “Descriptions, Truth Value Intuitions, and Questions.” *Linguistics and Philosophy*, 32, pp. 583–617.
- Soames, Scott, 1999 *Understanding Truth*. Oxford and New York: Oxford University Press.
- Stalnaker, Robert C., 1979 “Assertion.” *Syntax and Semantics*, 9, pp. 315–322. Reprinted in Stalnaker (1999), pp. 78–95.
- , 1984 *Inquiry*. Cambridge, MA: MIT Press.
- , 1987 “Semantics for Belief.” *Philosophical Topics*, 15. In Stalnaker (1999), pp. 117–129.
- , 1999 *Context and Content*. Oxford: Oxford University Press.
- , typescript “The Autonomy of Semantics.”
- Stanley, Jason, 2000 “Context and Logical Form.” *Linguistics and Philosophy*, 23, pp. 391–434.
- Tappenden, Jamie, 1993 “The Liar and Sorites Paradoxes: Toward a Unified Treatment.” *The Journal of Philosophy*, 90, pp. 551–577.
- Travis, Charles, 1985 “On What is Strictly Speaking True.” *Canadian Journal of Philosophy*, 15, pp. 187–229.
- Wilson, D. and Sperber, D., 2004 “Relevance Theory.” pp. 607–632. In Horn and Ward (2004).

Wright, Crispin, 1976 “Language-Mastery and the Sorites Paradox.” In Evans and McDowell (1976), pp. 223–47; reprinted in Keefe and Smith (1996), pp. 151–173.

Yablo, Stephen, 2006 “Non-catastrophic Presupposition Failure.” In Byrne and Thomson (2006).

———, forthcoming “Aboutness.”

Yalcin, Seth, 2005 “Epistemic Modals.” *Mind*, 116, pp. 983–1026.