

Minimalism, psychological reality, meaning and use¹

I. Introduction

A growing number of philosophers and linguists have argued that many, if not most, terms in our language should be understood as semantically context sensitive. In opposition to this trend, Herman Cappelen and Ernie Lepore (hereafter "C&L") defend a view they call "Semantic Minimalism", which holds that there are virtually *no* semantically context sensitive expressions in English once you get past the standard list of indexicals and demonstratives such as "I", "you", "this", and "that".² While minimalism strikes many as obviously false, it will be argued here that the view is more plausible than commonly assumed if one accepts the 'normative' conception of the relation between meaning and use characteristic of the literature on semantic externalism. That said, it is not clear that C&L always conceive of their minimal semantic contents in this more normative fashion, and once this framework is place, a version of 'moderate contextualism' that they do not consider comes into view.

II. Minimalism and Contextualism

There are familiar aspects of our linguistic practice that point towards contextualism. For instance, in some contexts (discussing professional basketball players) we deny that someone is "tall" if they are under, say, 6'7", while in others (discussing 2nd grade students) we claim that someone is "tall" if they are over 5'1". What counts as "tall" in some contexts does not in others, and many have argued that the semantics for "tall" should reflect this (call these "context shifting arguments"). Further, some argue that, even with the meaning of "John" fixed, in absence of a surrounding context, sentences like "John is tall" or "John is ready" can seem to have no determinate meaning at all, that is to say, we have no idea just *what* is being said by such sentences in the absence of any completing context (call these "incompleteness arguments").³

¹ I'd like to thank Herman Cappelen, Ernie Lepore and members of the 2005 Lisbon seminar on *Meaning and Communication* and two readers from Oxford University Press for comments on earlier versions of this paper.

² Other members of this list would include "he", "she", "it", "here", "now", "there", "then", "today", "yesterday", "tomorrow", "ago", "henceforth", "actual" and "present". Noticeably absent from their list are terms like "red", "tall", "ready", "every", or "know" all of which are treated as context sensitive on many currently popular semantic theories. (The view would, of course, apply to other natural languages as well.)

³ For versions of such arguments for contextualism, see, for instance, Bezuidenhout 1997, Carston 2002, Recanati 2001, 2004, Searle 1978.

By contrast, C&L present a general argument that they take to show that such commonly offered reasons in favor of the context sensitivity of terms like "tall" can't be good. This argument distinguishes 'moderate' contextualism (which takes only *some* expressions -- like "tall", "flat", "know" -- to be context-sensitive), from "radical" contextualism (which takes *every* expression to be context-sensitive). The argument runs roughly as follows:

1. Moderate contextualism entails radical contextualism.
2. Radical contextualism is incoherent.

Therefore, 3. Moderate contextualism is false.

Therefore, 4. Minimalism is true.

C&L try to establish premise number 1 by arguing that the sorts of context shifting and incompleteness arguments that support moderate contextualism can also (with a little imagination) be used to support radical contextualism, so that if they provide us with any reason to be a moderate contextualist, they provide us with reasons to be a radical one.⁴ Premise number 2 is, in turn, defended on the grounds that the truth of radical contextualism makes communication impossible, since no two people are ever in exactly the same context.⁵ Further, if radical contextualism were true, it would be impossible to even *state* the view, so the explicit endorsement of it is incoherent.⁶

Aside from the bare claim that the view is manifestly implausible (and the proverbial incredulous stare), there are three main reasons for thinking that minimalism can't be true.⁷ These are:

1. *The Processing Objection*: The sorts of invariant contents that C&L postulate violate what Recanati refers to as an "availability constraint" on what is said. Namely:

What is said must be intuitively accessible to the conversational participants.⁸

⁴ See C&L 2005, Ch. 3.

⁵ See C&L 2005, Ch. 8. The arguments here are structurally very much like those presented against holism in Fodor and Lepore 1992. For reservations about arguments of this form, see Jackman 1999, 2003.

⁶ See C&L 2005, Ch. 9.

⁷ One could also argue against Minimalism by attempting to diffuse C&L's arguments against contextualism. However, I will not be concerned with such arguments that minimalism *needn't* be true, only with the ones purporting to show that it *can't*.

⁸ Recanati 2004, p.20. (He adds the qualification "Unless something goes wrong and they do not count as 'normal interpreters'", but that will not be relevant here.)

The sorts of invariant contents the C&L postulate are not accessible in such a fashion. Indeed, what is accessible often only seems to be the more context-sensitive contents, and C&L themselves are unable to non-homophonically specify what the invariant contents of our utterances are supposed to be. Consequently, such contents do not seem like candidates for generating the context sensitive sayings and implicatures that we grasp in conversation. Minimal contents do not seem suitable to play any role in the psychological processing relevant for language understanding and thus lack "psychological reality".

2. *The Learning Objection*: C&L are able to fit semantic minimalism with our intuitions about what is said in various contexts by combining it with what they call "Speech Act Pluralism" (hereafter "SAP"), which they characterize as follows:

No one thing is said ... by any utterance: rather, indefinitely many propositions are said... What is said (asserted, claimed, etc.) depends on a wide range of facts other than the proposition semantically expressed. It depends on a potentially indefinite number of features of the context of utterance and of the context of those who report on (or think about) what was said by an utterance. (C&L 2005, p.4.)

According to SAP, the semantic content of a sentence is only one of the many things said on any occasion of its utterance, so intuitions about what was said on any given occasion are to some extent distanced from what the semantic content of the sentence uttered must be.

However, given that what is consciously accessible remains varied and context sensitive, SAP may seem to leave C&L's purportedly invariant contents unlearnable. Since the most salient thing said is usually not the minimal content, it might seem that, even if there were some invariant content that each of us associated with any given word, there would be no way to be sure that each of us attached the *same* invariant contents to the words in our languages. If I always minimally meant something like "appears red under some condition or other" by "red",⁹ while my parents always minimally meant, say, "appears red under *normal* conditions", there would be no way for me to learn that I was mistaken, since in any given context, we might still use our words in exactly the same way.

SAP allows the people who 'minimally' mean very different things to still manage to say the same thing with the same words in any given contexts, and thus make it difficult to see how communication requires that groups of speakers converge on the same minimal contents. As

⁹ C&L 2005, p. 160.

long as their differing minimal contents lead both speaker and hearer to non-minimal contents that were shared, communication would continue to go smoothly.

3. *The Supervenience Objection*: Semantic minimalism also seems to be in tension with the idea that what we mean by our terms is a function of how we use them. Our usage shows a high degree of context sensitivity, and unless one could show that there is a stable pattern behind this usage, it might seem as if semantic minimalism is incompatible with any sort of 'naturalism' about meaning. If, as the contextualist argues, there is no non-pleonastic property that all of the things that we call, say, "red" share, then it is hard to see how our use of the term could determine an invariant meaning that is constant through all of our applications of the term.

Such objections are familiar, and many find them to be, in some form or another, persuasive. However, the prospects for meeting these objections are not as bleak as most opponents of minimalism suggest, provided that one understands minimalism in the context of a certain sort of semantic externalism about mental and semantic content.

III. Semantic Externalism

Semantic externalism takes the contents of one's thoughts and utterances not to be determined by one's 'narrow' psychological states. That is to say, externalism allows that the truth conditions of what we say and think are at least partially determined by factors that do not play a role in either conscious or unconscious processing. These include, but are not limited to, the history of our terms' usage, how they are used in our society, and the actual make-up of our environment. To take the two best known examples of this, the externalist allows that (a) two people who are internally identical could mean different things by "water" if the actual substance they applied it to in their respective environments had a different underlying microstructure (while their typical phenomenal properties were the same),¹⁰ or (b) two people who are internally

¹⁰ Putnam 1975.

identical could mean different things by "arthritis" if the term was used differently in their respective environments (even if they were unaware of these differences).¹¹

The literature on semantic externalism highlights a tension within our unreflective notion of psychological states. Psychological states were taken to both be states with a *processing role* in that they were realized in the head and were the subject of empirical psychology, and states with a *cognitive role* in that they were true or false and the subject of more 'normative' disciplines like epistemology. The externalist suggested that if there were states of these two types, they could not be easily identified. Hence Putnam's original claims that "meaning ain't in the head", and that what we meant by our terms was independent of our 'psychological states' (Putnam 1975). Of course, this use of 'psychological' is somewhat tendentious, and some have insisted that 'psychological' states are 'wide' as well (Burge 1979, 1986). Consequently, I'll be here using "cognitive" for this wider sense and reserve "processing" for the narrow states.

Communication can be understood in both a 'practical' and a 'normative' way. The 'practical' way requires only that it produce successful coordination of the behaviors of the participants involved, the 'normative' way requires that the hearer actually grasp a thought with the content that the speaker was expressing. Externalism was traditionally motivated at least in part by a desire to explain our apparent ability to communicate in this more 'normative' fashion across both social and historical contexts (whether it be a botanist's ability to understand our comparatively uninformed discourse about Boston Ivy, or our ability to understand, say, Aristotle's writings about gold). For the externalist, meaning and content are more 'invariant' than they might otherwise seem, and ensuring this invariance requires accepting the idea that the truth conditions of our thoughts and utterances can often be quite different than we might (non-homophonically) specify them as being.¹² Crucially, it cannot be required for two people to grasp the same content, that their narrow psychological states have some shared non-relational property (like a shared high-level description or image). Shared narrow contents are neither necessary nor sufficient for communication in this more normative sense.

¹¹ Burge 1979.

¹² Indeed, Putnam's original presentation of semantic externalism seems to have emerged from his desire to defend the idea that the meaning of terms remained invariant through changes in our scientific theories. Also, one should note, that the commitment to invariance characteristic of externalism is a *defeasible* one, and that in some cases we do treat the meanings of our terms as having changed over time or as varying between two people in a single community. It will eventually be suggested that the invariance associated with minimalism should be understood as defeasible in this way.

Semantic Externalism has a more 'normative' conception of the relation between use and meaning, and the contents associated with this more normative framework should be thought of as distinct from whatever is needed to give a causal explanation of linguistic processing. Unfortunately, while these two stories are distinct, they are often run together in the literature in this area. On the one hand, there is a tendency to look at a plausible empirical story about language processing and treat it as if it translates directly to what we should think the more 'normative' contents of our cognitive states should be. This running together of the two types of story produces a familiar form of 'psychologism' about cognitive content. Some varieties of contextualism may, to a certain extent, be guilty of this. On the other hand, it is equally unjustified to take constraints that are placed on our accounts of cognitive content, and treat them as applying to an empirical account of language processing as well. To do so would be guilty of, for want of a better word, a type of 'philosophism' about psychological processing. For instance, our intuitive idea of cognitive content is plausibly one in which, among other things, we produce utterances that are true or false and communicate by sharing thoughts with the same content. However, it is far from clear that such requirements transfer directly, or at all, to our account of psychological processing. In particular, it is far from clear that any story of language processing needs to find a place for such invariant and shared contents.

Both psychologism and philosophism involve taking a monistic view towards content and communication, treating the psychological and philosophical story as if they should be the same, or at least isomorphic to each other. It will be argued below that minimalism is much more plausible once this sort of monism is rejected.

IV. Minimalism and Processing

The 'processing objection' has much less bite if the sorts of contents that the minimalist are talking about are understood as *externalist* contents, since such contents were never plausibly part of a simple processing story. Anyone offering the 'processing objection' against minimalism also needs to show that the relevant sense of 'psychological processing' is not one that rules out contents conceived in the more traditional externalist way as well.¹³

For instance, Recanati argues that the non-contextualist position is incompatible with truth conditional semantics given that we can't non-disquotationally specify the truth conditions for

¹³ Unless, of course, the contextualist plans to reject standard varieties of externalism as well.

sentences like "John cut the sun".¹⁴ While one might want something more than mere disquotation, an endorsement of semantic externalism requires that our ability to specify the truth conditions of our utterances may not be as extensive as may have originally been assumed. Someone who has a very thin, and largely mistaken, conception of what airfoils are, can still be treated as meaning *airfoil* by "airfoil" in virtue of deferring to experts who have mastered the term. Consequently, we need to ask (1) how 'psychologically real' such externalist contents are, and (2) what, if any, reason we could have for doubting that invariant contents aren't at least this psychologically real as well. It would be unfair to use a stronger requirement of psychological reality when evaluating the minimalist's talk about content than one does with the externalist's.¹⁵

Invariant contents may have no "psychological role" if by that we mean a certain kind of *processing* role, but if by "psychological role" we mean *cognitive* role, then they certainly do.¹⁶ Further, it is precisely such non-processing related cognitive roles that C&L stress when they discuss the importance of invariant contents.¹⁷ C&L explicitly endorse the more 'Fregean' conception of communication as requiring shared thoughts,¹⁸ and they defend the claim that invariant contents have the cognitive role of ensuring communication by allowing us to grasp thoughts with the same content as our interlocutors in spite of our different contexts.

If invariant contents are meant to play this sort of cognitive rather than processing role, then C&L's *laissez faire* "that's just a problem for the metaphysicians" attitude towards just what these invariant contents are can seem more understandable.¹⁹ It would be comparable to the claim that

¹⁴ Recanati, 2004, pp. 92-3.

¹⁵ This is not to deny that there may be conceptions of psychological reality that the externalist contents could meet while the minimal ones could not (after all, for many -- but by no means all -- externalist contents, the correct application conditions are *actually* known at least by the *experts*). However, even if such conceptions of psychological reality could be found, they are not strong enough to underwrite the processing objection.

¹⁶ Recanati makes the following remark about minimal propositions: "Let the semanticist use it if he or she wants to, provided he or she agrees that ... the minimal proposition has no psychological reality. It does not correspond to any stage in the process of understanding the utterance, and need not be entertained or represented at any point in that process" (Recanati 2001, p.89). C&L object to this remark since they take the *cognitive* role of minimal contents to show that they have "psychological reality", while Recanati takes their lack of *processing* role to show that they don't. With the exception of his use of "psychological reality", I think that the minimalist can accept Recanati's claim here, which really could be viewed as the suggestion that the minimalist can keep invariant propositions for the cognitive role, provided that they don't take them to have a processing role. Indeed, given that the roles are different, it would seem possible to be a minimalist about cognitive content and adopt the most extreme form of contextualism (what Recanati refers to as "Meaning eliminativism" (Recanati 2004, pp.146-51)) for one's processing story.

¹⁷ C&L 2005, pp. 184-5.

¹⁸ C&L 2005, p. 153.

¹⁹ See C&L 2005, Ch. 11.

it is up to the scientists to discover just what we mean by "gold" or to the judges to determine just what we mean by "culpable negligence". If the invariant contents played a causal role in processing, then it might seem as if the semanticist *should* be able to say what they were, but if the invariant contents play no processing role, then there is no reason to think that finding these contents shouldn't be more properly a problem for the metaphysicians, judges, scientists, and other investigators of the world that our words are ultimately about. Like any externalist, the minimalist will insist that the contents of our words can only be specified by investigating what they are about, not through any sort of *a priori* conceptual analysis.

Unfortunately, the fact that invariant contents need not play any role in psychological processing is often obscured by C&L's presentation of their own position. Indeed, while I've been at some pains to distinguish the cognitive from the processing role that minimal contents could be candidates to play, C&L freely admit that "If there's a difference between having a cognitive function and corresponding to a stage in processing or having psychological reality, we don't know what that difference consists in."²⁰ Having a cognitive function and corresponding to a stage in processing are, as suggested above, two different ways of having psychological reality, and just because invariant contents are psychologically real in the first sense (and most of C&L's reasons for thinking that invariant contents are psychologically real relate to this first sense), it does not follow that they will be psychologically real in the second.

Their choice of terminology does not help matters either. The term "minimalism" has a familiar use from Recanati (a minimalist view is one that minimizes the distance between sentence meaning and what is said),²¹ but while C&L are minimalists in Recanati's sense, so are many of the "moderate contextualists" (e.g.: Stanley 2000) that C&L take their minimalist view to be *opposed* to. The most obvious sense in which C&L's views could be called "minimalist" seems to be that, in comparison to most semantic theories, it *minimizes the number* of context-sensitive expressions in our semantics.²² However, there is nothing in the view that suggests that the *type* of content that is invariant should be in any way 'minimal'. Unfortunately, C&L's references to the invariant semantic contents as "minimal semantic contents" can lead one to

²⁰ C&L 2005, p. 186.

²¹ Recanati 2004, p. 7.

²² "The most salient feature of Semantic Minimalism is that it recognizes few context sensitive expressions" (C&L 2005, p.2). Their division of the field as a choice between radical contextualism, moderate contextualism and minimalism also encourages this interpretation.

misread the role of invariant contents in a way that would make their position appear vulnerable to the processing objection.

After all, given the many different ways that a term like "tall" is used, what all of these uses would have in common might seem to be very thin, general or 'minimal', and thus calling invariant semantic content "minimal" would suggest that one wanted something thin enough to be processed by everyone who used the term.²³ Consider C&L's discussion of dancing:

Some people dance by stepping, some crawl around the floor (like Martha Graham), some have music, some don't have music, some jump in the air, some wave their arms, some hold on to other people, some are alone, some slide on ice, some fly in the air, etc. What do all these activities have in common in virtue of which they are all dancing? (C&L 2005, p. 161)

While this can seem like a straightforward metaphysical question, the implicit assumption that there *is* something which all these activities have in common can make it sound as if there is something in our heads that leads us to call all of these activities "dancing". However, if invariant content bears a normative relation to use, there is no reason to think that *all* of the activities mentioned above need have *anything* in common, since it could very well turn out that some things which we called "dancing" were not, in fact, instances of dancing at all. Contextualism keeps semantic values close to actual usage, and there is no reason to think that the minimalist can, or should, do this.²⁴ The more normative/externalist conception of content allows that everyone who uses a term means the same thing by it because they all committed to the same satisfaction conditions for the use of that term, even if there is no shared 'minimal' non-relational thing that is being processed in all of their heads.

Because of this, we should be careful when considering claims like: "The proposition semantically expressed is our minimal defense against confusion/misunderstanding and it is what guarantees communication across contexts of utterance."²⁵ At one level this claim is true. Invariant contents ensure communication across contexts. Nevertheless, it is less clear how well

²³ Compare recent 'minimalist' accounts of truth (Howich 1990, Wright 2001), which take the 'minimal' concept of truth to be one that we can all agree on (the disagreements relating to whether the concept extends beyond the minimal one). Minimal contents might be (mis)understood as a generalized version of this.

²⁴ Further, this reduces the problem posed by those accounts which suggests that the application of our terms are prototype driven and thus that there may not be any objective property in common among all the items to which a single term is applied (see, for instance, Lakoff 1987). The minimalist could insist that actual use is generated by a prototype while insisting that the invariant content that is abstracted from the use so generated need not reflect the prototype's structural features.

²⁵ C&L 2005, p.185. See also: "the proposition semantically expressed is that content the audience can expect the speaker to grasp (and expect the speaker to expect the audience to grasp, etc.) even if she has ... mistaken or incomplete information [about the context]", or "The minimal semantic content is a "shared fallback content" and ... this content serves to guard against confusion and misunderstanding." (C&L 2005, pp. 184-5, 185.)

they defend against confusion and misunderstanding. They ensure communication, in that they ensure that people in different contexts will be talking about the same things, but they do not in the sense of ensuring that there is some minimal core of one's concept that will be found in each person's conception. That is to say, it ensures communication in a normative rather than practical sense. However, confusion and misunderstanding often result when people have radically different *conceptions* of what they are talking about, and externalist (and minimalist) content only ensures that people are talking about the same things, not that they conceive of them in the same way. Indeed, even with shared invariant contents in place, speech act pluralism will allow for lots of miscommunication, because even if there is a minimal content that is shared, the main thing the speaker is trying to get across typically won't be the minimal content.

Further, one should not think of the invariant contents as the starting points from which speakers can be expected to derive the other propositions expressed in a context. Speech act pluralism allows that its semantic content is just one of many contents expressed by an utterance in a particular context, but it would be a mistake to see the 'minimal' content as the more 'basic' content from which the other contents can be generated. Because of this, it seems misleading of C&L to argue:

How does it help an audience to know that this minimal proposition was expressed? ...Our response is simple: it is a starting point. Suppose, for the sake of argument, that the proposition *A is red* is trivially true [i.e. it is true if A is red on some surface or other under some condition or other] ... The audience knows that the speaker is talking about A and its redness, not, for example, about oysters, France or Relevance Theory. There's a lot of stuff to talk about in the universe. The proposition semantically expressed pares it down considerably. Knowledge that this proposition was semantically expressed provides the audience with the *best possible* access to the speaker's mind, given the restricted knowledge that they have of the speaker. It is trivial that A is red on some surface or other under some condition or other. The audience can assume that the speaker knew that this was trivial and was not interested in conveying such trivialities with his utterance and can, therefore, infer that there is work to be done in order to figure out exactly what the speaker was trying to communicate. (C&L 2005, pp. 185-6.)

If invariant content bears a normative relation to use, we shouldn't expect it to play this role. There is no reason to think that the invariant content need be psychologically available in a fashion that would allow the speaker or hearer to draw inferences about the point of an utterance using it.

V. Minimalism and Language Learning

It should also be clear that if the minimalist is a semantic externalist, the "learning" objection lacks bite. For the externalist, speakers are not "conceptually autonomous" in the sense of

having, on their own, complete command of the conditions of application for the terms in their language, since the meaning of a their terms extend beyond both their discriminatory capacities and the descriptions available to them.²⁶ This will have considerable consequences for what externalists should take language learning to require.

The idea that speakers are conceptually autonomous can lead one to think that two speakers share a language only if the two languages they independently speak are of exactly the same *type*. Learning a language thus involves acquiring autonomous mastery of a language that is type identical to the target language. However, the externalist should understand neither sharing a language nor learning a language the way that the model of the autonomous speaker suggests. In particular, Putnam's "division of linguistic labor" metaphor suggests that I can be understood as sharing a language with other members of my community because we look to each other for help when deciding how our terms are correctly applied. Two autonomous speakers, even if they spoke in *exactly* the same way, would not, in this sense, *share* a language at all. Non-autonomous speakers are understood as speaking the same language not in virtue of speaking two type identical languages, but rather in virtue of both taking part in a single (*token* identical) linguistic practice.²⁷

Further, this conception of what it is for two people to speak the same language brings with it a correspondingly different sense of what it is to *learn* a particular language. Learning a language involves being initiated into a particular linguistic practice, and it need not involve achieving anything like complete mastery of the practice in question. Just as I can mean *arthritis* by "arthritis" even if I haven't mastered every aspect of the doctor's usage, my learning to mean by "red" what my parents meant by it does not require that we use the term in precisely the same way.

²⁶ Indeed, speakers' discriminatory capacities often appear to underdetermine or even misidentify what they are talking about. A typical example of underdetermination is Putnam's use of "beech" and "elm" (Putnam 1975, p. 226), while the best known discussion of misidentification is Burge's discussion of Bert's belief that he has arthritis in his thigh (Burge 1979).

²⁷ This line of thought is developed in considerably more detail in Jackman 1998a.

Since the contents involved have a type of cognitive rather than processing reality, skepticism about whether my parents and I mean the same thing by "red" seems unmotivated. Such worries would make sense if we thought of the minimal content as some unconscious representations from which our actual use was generated. Such representations would be ontologically independent, and so the question of whether they were tokens of the same type would be a legitimate one. However, if the invariant contents of our terms are not independent of each other in this way, then endorsing minimalism should not lead to skepticism about whether I can mean by "red" what my parents do by it

VI. Minimalism and Supervenience

The supervenience objection is perhaps the most serious problem for the minimalist position, and ultimately may be what lies behind most people's reservations about the view.

Even if one accepts C&L's arguments that moderate contextualism leads to radical contextualism and that radical contextualism is incoherent, there remains a question of what such arguments manage to establish. In particular, such arguments would seem to leave the minimalist in a position familiar from the literature on vagueness. Epistemicists like Timothy Williamson and Roy Sorensen present what can seem like fairly compelling arguments to the effect that all of our terms *must* have precise extensions.²⁸ In spite of this, many philosophers refuse to take epistemicism seriously because they are convinced that the extensions of our terms just *can't* be as precise as the epistemicist requires. Our terms can only be as precise as our usage and environment make them, and nothing in what determines meaning seems like it could, for most terms in the language, produce the purportedly required level of precision.²⁹

In much the same way, the minimalist argues that we are committed to the meanings of our words having certain properties (in this case context-invariance rather than precision) that intuitively our use of language does not seem up to producing. If C&L are right, we may be committed to the existence of context invariant contents, but it doesn't follow from this commitment that there actually *are* such contents. Just as the epistemicist should back up the

²⁸ Since admitting that any of our terms have semantically borderline cases can seem to quickly commit us to explicit contradictions. (See Williamson 1994, Sorensen 2001).

²⁹ For a dissenting view of this last claim, see Williamson 1994, Jackman 2004.

claim that we are committed to meanings being precise with a story about how there could be such things, the minimalist owes us a story about how our seemingly context sensitive practice of word usage could produce context invariant meanings.³⁰

Simply being an externalist doesn't remove these problems, since some explanation of how our utterances get their cognitive content is still owed. Nevertheless, since externalism presupposes a more 'normative' relation between meaning and use, the problem can seem less insuperable than it might have otherwise. For the externalist, meanings are the types of things for which our commitments about them contribute to what they are,³¹ and if we are committed to the meanings of our terms being invariant, then that can help make it the case that they are.

While the idea that meaning is determined by use might point towards contextualism given how context sensitive the use of our terms often is, C&L rightly point out that important parts of our usage actually suggest that our words have invariant contents. These aspects of our usage include the following:

Intercontextual indirect belief reports: Paradigmatically context sensitive expressions do not allow for intercontextual indirect belief reports. For instance, if John said "I'm hot today", Peter cannot truly say the next day that "John said that I'm hot today", rather he would need to say "John said that *he* was hot *yesterday*". Purportedly context sensitive expressions like, say, "tall", do not seem to display this behavior. If John claims that he is tall in one context, we can truly claim "John said that he was tall" in another. (See C&L 2005, pp. 88-99.)

Reports under ignorance: Sometimes a person reporting on an utterance might be ignorant of the relevant contextual features of the original context of utterance. That ignorance, however, needn't influence (what we take to be) the truth value of the indirect report. (See C&L 2005, p. 93.)

Reports based on mistaken assumptions: Sometimes the reporter has false beliefs about the original context of utterance. Such false beliefs need not influence (what we take to be) the truth value of the disquotational indirect report. (See C&L 2005, p. 93.)

Collective Ascriptions: If John said "Peter is tall", and Mary said "Peter is tall", then we can say that "John and Mary both said that Peter was tall" even if John and Mary are in different contexts not only from us, but also from each other. We cannot, in a similar way, say that "John and Mary both said that it was hot today" if they each said "it is hot today" on different days. (See C&L 2005, pp. 99-104.)

These aspects of our usage are important, and often overlooked, but they are *defeasible* and one can't tell *a priori* that any adequate semantic theory for our language must endorse these aspects of our linguistic practice over others. The fact that our practice embodies a commitment to

³⁰ Of the two, the minimalist may be better off in that it might seem more plausible to think that shared context invariant meanings could supervene upon our usage than perfectly precise ones. On the other hand, the epistemicist may be in a stronger position given that our commitment to classical logic will seem to many to be stronger than our commitment to, say, the truth of collective ascriptions.

³¹ For a more extended discussion of this see Jackman 1996, 1998a, 2003.

invariance doesn't entail that it is a necessary criterion of adequacy for any semantic theory. Our commitment to invariance is just one (heavy) commitment among others, and it may have to be given up if it is incompatible with too many of the rest. The question becomes, then, which is more deeply entrenched, those aspects of our usage that point towards contextualism or those that point towards invariantism?

We see a similar tension in the literature on semantic individualism. Given that each individual uses a their words differently, the idea that meaning is a function of use might seem to point towards an individualistic semantics according to which, say, someone who used "arthritis" in the way we did but also applied it to pains in his thigh would mean something different by "arthritis" than we do. However, other aspects of our usage (interpersonal indirect belief reports, collective ascriptions, deference) point towards such a speaker still meaning what we do by "arthritis". In the "arthritis" cases most argue that those aspects of usage that support socially shared meanings trump those that support more individualistic meanings.³² However, the question of whether or not meaning should be equated with social or individual usage cannot be settled for an entire language at once. In some cases the standard meanings are *not* ascribed to the idiosyncratic speaker (whose usage may differ too radically from the standard, or who may simply be unwilling to defer to the standard usage of a particular term).³³ In much the same way, the question of whether meanings are invariant or contextually sensitive may have to be settled on a word-by-word basis, and for each particular word the question will ultimately be an empirical/metaphysical one.³⁴

Minimalism involves the very strong claim that for *every* word outside of C&L's basic set, those aspects of our usage favoring invariance will trump those favoring contextualism. The radical contextualist is committed to the equally strong claim that such invariance-favoring considerations will *never* trump the contextualist ones. The moderate contextualist, by contrast, can take the view that for some words the considerations favoring invariantism will be more central, while for others, the contextualist considerations may win out.

³² Burge 1979 being, of course, the classic expression of this view. There are, of course, some dissenters (such as Bilgrami 1992, and Davidson 1987, 1994).

³³ For a discussion of this, see Jackman 1998b.

³⁴ Consequently, it may be surprising which expression turn out ultimately to be semantically context sensitive (*pace*, C&L 2005, p. 112).

C&L claim that moderate contextualism collapses into radical contextualism because the standard arguments for moderate contextualism (context shifting, incompleteness) can be extended with a little imagination to all terms in the language. This may be true if moderate contextualism is motivated solely by such arguments, but the sort of argument for moderate contextualism considered here is different. A term might be taken to have a contextually sensitive semantics because there is no acceptable way of systematizing our usage of it that preserves the intuition that meaning is invariant. The fact that this might be true of a term like, say, "ready" would give us no reason to think that it need be true for a term like "flat" or "every". All three terms might be subject to context shifting and incompleteness arguments, but it is arguably only the first that does not seem to have a single leading candidate for what its invariant meaning should be.

How plausible an invariant semantics is for any given term may ultimately depend upon the strength of the candidates for its proposed invariant meaning. Given that I've suggested both that the questions of whether our words have invariant or context sensitive semantic contents should be settled on a piecemeal rather than wholesale basis, and that this question need not be resolvable in a purely *a priori* fashion, I won't commit myself here to how I think most of our vocabulary will ultimately play out. Nevertheless, I will note that the case for invariantism certainly seems stronger for some words than for others.

For instance, certain words (quantifiers like "every" or connectives like "and") have clear leading candidates for what their invariant meaning would be (unrestricted quantification, logical conjunction). Of course, such candidates may be out of line with much of our everyday usage, and people might be unwilling to give up the literal truth of everyday claims like "there is no beer left" in order to preserve the idea that they could truthfully report claims made using such words from other contexts. Nevertheless, this is not a problem if, like C&L, we endorse Speech Act Pluralism. SAP gives one a considerable degree of flexibility in determining the function between use and meaning, by allowing one to say, for instance, that when someone utters "there is no beer left", they have still *said* (not just implicated) something true. The minimal semantic content of the utterance (there is no beer left in the universe) may still be something false, but that is only one of the many things said in the context, and the most salient of these may be that there is no beer left in the fridge. Consequently, the intuition that we have said something true doesn't count against proposed semantic analyses that make the sentence false. SAP allows one's

metaphysics/semantics to be comparatively revisionary because it allows that everyday utterances that don't fit the proposed semantic analysis can still be used to make true claims.

'Absolute' terms like "flat" might also lend themselves to this sort of analysis. With such terms there are clear leading candidates for what the term should mean, and such candidates can serve as the norm against which the other uses are understood. SAP makes it less worrisome that most of our usage fails to live up to these norms, and the invariantist account is also more plausible when we remember that the other speech acts, (the true claim you may communicate by "the table is flat") are not generated by some sort of conscious or unconscious inference from the invariant one.

However, things are less clear with terms like "ready" or "tall" where there seems to be no non-arbitrary way to settle on a candidate for what the invariant meaning should be. It is these cases, where there seem to be *many* equally good candidates for what the invariant meaning should be, that may lie behind many people's doubts about minimalism.³⁵ A commitment to invariance might give us reason to settle on a single meaning for such terms, but there is no reason to think that anything in our usage or the world has already settled on one.³⁶ The degree of freedom one gets from a normative conception of the relation between use and meaning allows one to equate meaning with a leading candidate even if that candidate is out of line with much of our actual usage. However, if there are *multiple* equally good candidates, this freedom doesn't give one the ability to *arbitrarily* select one of the candidates and say that that is what everyone has always meant.

If there is no leading candidate for what a word's invariant meaning could be, there would still seem to be two alternatives to becoming a contextualist about the term in question: The first

³⁵ This seems to be the basic worry behind MacFarlane (2007): "Semantic Minimalism is problematic not because it does not *provide* an answer to questions about the intensions of its minimal properties and propositions, but because it requires that there *be* answers to such questions." We assume that claims like "I don't want to know whether John is tall for a basketball player or tall for a six year old, I just want to know whether he is just plain *tall!*" don't seem to make sense because it seems as if nothing in our use of "tall" would have determined a particular extension for (just plain) "tall" rather than the more specific claims.

³⁶ The similarity with vagueness is, once again, evident here. A commitment to bivalence may give us reason to precisify a term that has multiple equally good candidate interpretations, without giving us any reason to think that our usage has already done so. As with epistemicism, minimalism would be more plausible in these cases if our metaphysics of properties was comparatively 'sparse' (see Heck 2003, Lewis 1983), but such strong metaphysical commitments don't seem otherwise well motivated. As should be clear from Jackman 2004, I'm inclined to deal with this by suggesting that the epistemicist should allow subsequent precisification to be retroactively read back into current contents, and one could argue that a committed minimalist should do this as well, particularly when it comes to dealing with the 'open texture' of language.

of these would involve adopting some sort of supervaluational account for expressions that don't have clear leading candidates for their meanings.³⁷ The other would be to adopt some sort of error theory about the semantic values of the relevant terms.³⁸

Neither of these options will be explored here, but I hope to have shown that many of the *prima facie* problems for semantic minimalism are, if not eliminated, at least reduced considerably when the view is understood from within an externalist framework. It is not clear how eager Cappelen and Lepore would be to embrace such a framework for their view, but I hope to have given some reasons for thinking that they should.

³⁷ Such as that found in Field 1973, or Wilson 1982.

³⁸ See Cappelen 2005, pp. 6-7. This second response has some appeal for moral terms like "fair", where the pull of invariance is extremely strong. (We might be more willing to say that there was no such thing as fairness than we would be willing to say that what *is* fair just varies from context to context (*Pace Unger* 1995).) It seems, however, that for most terms our commitment to invariance is a defeasible one, and if no invariant meaning could be found for "tall", we would be more willing to give up on some collective ascriptions than we would be willing to give up on the idea that anything is 'really' tall.

Works Cited:

- Bezuidenhout 1997, "Pragmatics, Semantic Underdetermination and the Referential/Attributive Distinction", *Mind*, vol.106, no. 423 pp.375-409.
- Bilgrami, A. 1992. *Belief and Meaning*, Cambridge: Blackwell.
- Burge, T. 1979. "Individualism and the Mental", in French, Uehling & Wettstein. *Midwest Studies in Philosophy IV: Studies in Metaphysics*, Minneapolis: University of Minnesota Press. 1979, pp.73-121.
- Burge, T. 1986. "Individualism and Psychology", *The Philosophical Review* Vol. 95.
- Cappelen, H. and Lepore, E. 2005. *Insensitive Semantics*, Cambridge: Blackwell.
- Cappelen, H. 2005. "Reply to McFarlane", paper presented at the 2005 Pacific Division Meeting of the American Philosophical Association.
- Carston, R. 2002, *Thoughts and Utterances: The Pragmatics of Explicit Communication*, Oxford: Blackwell.
- Davidson, D. 1987. "Knowing One's Own Mind", Reprinted in Ludlow and Martin (eds.), *Externalism and Self-Knowledge*, Stanford: CLSI Publications. 1998.
- Davidson, D. 1994. "The Social Aspect of Language", in McGuinness and Oliveri (eds.). *The Philosophy of Michael Dummett*, Dordrecht: Kluwer, 1994.
- Field, H. 1973. "Theory Change and the Indeterminacy of Reference", *The Journal of Philosophy* Vol. LXX, August, 1973, 462-481.
- Fodor, J. and Lepore, E. 1992. *Holism: a shopper's guide*, Cambridge: Blackwell.
- Heck, R. 2003. "Semantic Conceptions of Vagueness", in J. Beall, ed., *Liars and Heaps* (Oxford: Oxford University Press, 2003), pp. 106-27
- Horwich, P. 1990. *Truth*, Oxford: Blackwell,
- Jackman, H. 1996. *Semantic Norms and Temporal Externalism*. Ph.D Thesis, University of Pittsburgh.
- Jackman, H. 1998a. "Convention and Language", *Synthese*, 117 (3) 1998: pp. 295-312.
- Jackman, H. 1998b. "Individualism and Interpretation" *Southwest Philosophy Review*, V.14, n. 1, Jan 1998, pp. 31-38.
- Jackman, H. 1999. "Moderate Holism and the Instability Thesis", *American Philosophical Quarterly*, Vol. 36, No. 4, October 1999, pp. 361-369.
- Jackman, H. 2003. "Charity, Self-Interpretation, and Belief" *Journal of Philosophical Research*, Volume 28 (2003), pp. 145-170.
- Jackman, H. 2004. "Temporal Externalism and Epistemic Theories of Vagueness" *Philosophical Studies*, Vol. 117, No.1-2, (January 2004), pp.79-94.
- Lakoff, G. 1987. *Women, Fire and Dangerous Things*, Chicago: The University of Chicago Press.
- Lewis, D. 1983. "New Work for a Theory of Universals", *Australasian Journal of Philosophy*, v.61, n.4.

- Macfarlane, J. (2007), "Semantic Minimalism and Nonindexical Contextualism", This volume.
- Putnam, H. 1975. "The Meaning of 'Meaning.'" Reprinted in his *Mind Language and Reality*, New York: Cambridge University Press. 1975. pp. 215-271.
- Quine, W.V.O. 1960. *Word and Object*, Cambridge: MIT.
- Quine, W.V.O. 1987. "Indeterminacy of Translation Again", *The Journal of Philosophy*, vol. LXXXIV, no. 1.
- Recanati, F. 2001. "What is Said" *Synthese*, 128: 75-91.
- Recanati, F. 2004, *Literal Meaning*, Cambridge: Cambridge University Press.
- Searle, J. 1978. "Literal Meaning" *Erkenntnis*, vol. 13, pp.207-24.
- Searle, J. 1987. "Indeterminacy, Empiricism, and the First Person", *The Journal of Philosophy* Vol. LXXXIV, no.3, 123-146.
- Sorensen, R. 2001. *Vagueness and Contradiction*, New York: Oxford University Press.
- Stanley, J. 2000. "Context and Logical Form", *Linguistics and Philosophy* 23, (2000), 391-434.
- Unger, P. 1995, "Contextual Analysis in Ethics", *Philosophy and Phenomenological Research*, 55: 1-26.
- Williamson, T. 1994. *Vagueness*. London: Routledge.
- Wilson, M. 1982. "Predicate Meets Property", *The Philosophical Review*, Vol. 91, No. 4, pp.549-89
- Wright, C. 2001. "Minimalism, Deflationism, Pragmatism, Pluralism" in Lynch, M. (ed.) *The Nature of Truth*, Cambridge: MIT Press.