

CHEAP CONTEXTUALISM¹

Peter Ludlow
Northwestern University

0. Introduction

In Ludlow (2005) I made a case for a version of contextualism that involved an implicit “L-marked” position for standards of knowledge. Along the way I offered some positive considerations in favor of the contextualist view, and some arguments to rebut recent anti-contextualist arguments from Hawthorne and Stanley.

Since that paper, Stanley (2005; ch. 3) has offered some new arguments which have given me pause about aspects of my original proposal. Specifically, Stanley has insisted that the “by standards of x” locution does not appear to behave like a genuine argument or L-marked position in either the syntax or semantics of natural language. While I think that there are plenty of implicit L-marked positions in knowledge reports for degrees of justification and subjective certainty, I think Stanley is right to insist that the “by standards of x” locution isn’t one of them. I also think that he is probably right (ch. 4) that implicit positions for grades of justification or certainty aren’t sufficient to secure the case for contextualism (see his discussion of Cohen).

On the other hand, I think the facts that Stanley highlights don’t make the case against contextualism so much as they make the case for a different version of contextualism—a version that he calls “contextualism on the cheap”. I prefer the term “inexpensive,” but whatever we call it I think that Stanley was too quick to dismiss the view. My goal in this paper is to defend a version of it.

My plan of attack is as follows. In part 1, I’ll lay out some preliminaries, including a working definition of contextualism and a survey of different strategies for executing contextualism—ranging from implicit L-marking to cheap contextualism.

In part 2, I'll review some of the arguments that I made against Stanley and Hawthorne in Ludlow (2005). The considerations raised there will be relevant to my defense of cheap contextualism.

In part 3, I'll turn to the arguments in Stanley (2005) and lay out why I think they don't undermine contextualism so much as they make the case for cheap contextualism. In particular, I'll argue that when someone utters a sentence like (1)

- (1) By the standards of Unger's seminar, I don't know I have hands

the talk of standards really has to do with picking out a usage of 'knows' that has been coined for application in Unger's seminar. I'll be arguing that the 'by standards of x' locution does not specify an L-marked phrasal position of the core predicate, so much as it engages in a form of "code switching"—swapping in a predicate from, in effect, another language. For example, consider one of Stanley's examples.

- (2) By the standards of Chemistry, the stuff in the Hudson River isn't water

Stanley argues correctly that the mechanisms in (1) and (2) are basically the same, but I think Stanley is misdiagnosing the nature of the mechanism. I will argue that in each case we have a kind of metalinguistic operation that involves reference to lexical items that are not, strictly speaking, part of our current negotiated language—the referred-to lexical items will include sharpenings of the terms 'know' and 'water' from other conversational contexts.

In part 4, I'll draw on other work I have done, including Larson and Ludlow (1993) and Ludlow (2000, 2006) to show how this fits neatly into a proposal I have made about the nature of the lexicon—in particular that the lexicon is dynamic and that we routinely negotiate little micro-languages on the fly as we move between groups and contextual settings. In doing so, we routinely coin new terms and flesh out or sharpen terms that are common coin (but are massively underdetermined).

In part 5, I'll consider some of the objections raised against cheap contextualism by Stanley (ch. 4) and show that his core objection banks on a widely shared assumption in epistemology—that "absolute" sharpenings of words (e.g. the sharpening in which 'flat' means absolutely flat, and 'known' means known with absolute certainty) are the core meanings of terms and that usages which diverge from these are forms of "loose talk" that, like saying it's three o'clock when it is 3:03, are sometimes permissible, but not strictly speaking true (thus assimilating cheap contextualism to views like pretense theory or a pragmatic story about assertability conditions). I challenge this assumption, observing that on the dynamic lexicon proposal

the absolute sharpening is not privileged—it is just one possible sharpening among many.

In part 6, I'll return to my initial (2005) discussion of Hawthorne and Stanley and review some of the considerations from the perspective of the dynamic lexicon.

In part 7, I'll offer some general discussion of how this view sits in comparison with Stanley's more general plea for the importance of practical interests in knowledge attribution (I'll argue that the views are compatible) and the semantic minimalism of Cappelan and Lepore (also compatible). I'll suggest that some of the neat labels we have for positions in this domain may be misleading—that the move space is more subtle and malleable than we have usually supposed.

Finally, in part 8, I'll suggest that while there are multiple ways of sharpening and fleshing out word meanings, the process is not “up for grabs” but is normatively guided. I'll suggest that talk of concepts (e.g. the concept of knowledge) is best understood as talk about word meanings. Discussions about the concept of knowledge, for example, are nothing more than normatively constrained arguments about how we should define ‘knowledge’.

1. Preliminaries

Hawthorne (2004) offers a working definition of contextualism in which the following two components hold.

- C1. [context-sensitive semantic values]: A given sentence, say ‘Chesner knows that he has feet’ has different semantic values relative to different contexts of utterance, (and this at least in part due to contextual parameters connected to the verb ‘know’ itself ...). In brief, the contextualist claims that the epistemic standards required for someone to count as meriting a positive knowledge ascription varies from ascriber to ascriber, with the result that one ascriber may truly say ‘He knows that he will be in Syracuse,’ referring to a given person at a given time, and a different ascriber may say ‘He doesn't know that he will be in Syracuse’, speaking of the same person at the same time.
- C2. [ascriber calls the shots]: According to standard contextualist semantics, the ascriber calls the shots, so to speak: the standards of application for the verb ‘know’ are determined by the ascriber and not by the subject (unless the subject happens to be identical to the ascriber).

As I observed in Ludlow (2005), (C2) leaves matters open in a way that is not benign. For example, can the ascriber stipulate the context of utterance, or

is it out of the ascriber's hands altogether? Is it the context alone that does the work? In this case (C2) would be precisified as in (C2*)—this appears to be the precisification that Lewis (1996) was working with.

C2*. According to contextualist semantics, the *ascriber's context of utterance* calls the shots, so to speak: the standards of application for the verb 'know' are determined *by the context in which the ascription is made* and not by *the context in which subject appears* (unless the subject happens to be identical to the ascriber).

I say that this appears to be Lewis's view because if you are in a high-standards context (say an epistemology class with Unger) there is nothing you can say or do or whistle to invoke the low stakes context.

Alternatively, you might think that even if I am in Unger's epistemology class, I can invoke a lower stakes context by stipulating it—"well by the standards of Ludlow's class I know that P". If this is our view, then we sharpen up (C2) as (C2**).

C2**. According to contextualist semantics, the ascriber's context of utterance calls the shots, so to speak: the standards of application for the verb 'know' *are either explicitly stated* or are determined by the context in which the ascription is made and not by the context in which subject appears (unless the subject happens to be identical to the ascriber).

Now however we chose to sharpen up (C2) there are a lot of choice points remaining. For example, whichever sharpening of (C2) we endorse we still leave open the question of whether there are going to be argument positions for standards of knowledge or whether there is a kind of lexical sensitivity instead.

One wishes that this distinction could come off cleanly, but there are a number of factors that require us to move carefully here. For example, I think it is a bit sloppy to say that my (2005) paper made the case for implicit *argument* positions. In that paper I made the case for what Chomsky has called "L-marked phrases." In the remainder of this section I'll explain a bit about the notion of L-marking, and then once this idea is clear I'll lay out some of the possible ways in which one can be a contextualist.

One way to illustrate the traditional argument/adjunct distinction is by thinking about the argument structure of a sentence like (3).

(3) John hit Bill

In our introductory logic courses, we tell our students that translated into the predicate calculus, (3) will have an argument structure like that in (3').

(3') hit(John, Bill)

If we introduce Davidson's event analysis of action sentences, then we will propose an argument structure like that in (3''), incorporating a third argument position for events.

(3'') $(\exists e)$ hitting(e, John, Bill)

When we introduce adverbs we have the option of introducing them as additional argument positions, or we can think of them as being *adjuncts*, as in (4')—we say they are adjuncts because they are not part of the core event structure of the verb, but are adjoined to that structure by a conjunction.

(4) John hit Bill reluctantly

(4') $(\exists e)$ hitting(e, John, Bill) and reluctant(e)

Finally, we might suppose that sentences can also have *implicit* arguments. To illustrate, consider (5).

(5) John ate

Does 'ate' in this sentence merely have only one argument? Or does it have an *implicit* argument as well (for the stuff that John ate)? How about an argument for the *instrument* he used for eating (fork, fingers, etc).

All this may seem obvious enough, but more recent work in event logic has raised the question as to whether the argument/adjunct distinction hasn't collapsed. For example, in Castañeda (1967), Parsons (1990), it is proposed that so-called arguments are actually linked to the core event predicate via relational predicates called *thematic roles*, so that strictly speaking, everything except the event argument is actually a kind of adjunct and not part of the core argument structure of the predicate. The correct analysis of (4) is thus not as in (4'), but rather (4'') in which both the agent and patient are operating more as adjuncts than as arguments.

(4'') $(\exists e)$ (hitting(e) & agent(John, e) & patient(Bill, e) & reluctant(e))

Maybe it's better to say that verbs *select for* certain phrases (whether arguments or adjuncts). Following Chomsky 1986, we can call these selected phrases *L-marked phrases*. Very informally, the basic idea is that verbs favor having certain phrases with thematic roles associated with them—usually an agent, sometimes a patient, sometimes an instrument, etc.

When we turn to the topic of contextualism then, the question will not be whether a verb like 'knows' has an argument position for standards of knowledge, but whether it L-marks for a prepositional phrase that indicates

standards of knowledge; does the verb naturally select for (L-mark) a standards of knowledge phrase?

Of course when we say that the verb selects for phrases of a certain type, we are not saying that the phrase will be explicitly voiced. It may well be an *implicit* L-marked phrase—something that is linguistically represented but not articulated.

What kinds of tests are available to determine when a verb L-marks for a certain kind of phrase? Larson (1988), discussing implicit *arguments*, suggested that implicit elements include optional but non-iterable phrases such as phrases of source, path, goal and phrases of instrumentality, as in (6), where the material in parentheses is not articulated, but is nevertheless implicit. (Other tests are discussed in Ludlow (1995, 1996).)

- (6) a. John ran (Source from the house) (Goal to the store) (Path along the river).
 b. John cut the salami (Instrument with a knife).

Shortly, we will address the question of how these implicit phrases are to be represented, and as we will see there is a broad range of views on the matter. First, however, we should connect these considerations with the issue of contextualism. Given this understanding of (possibly implicit) L-marked phrases, we can address the issue of contextualism by revisiting a sentence like (1).

- (1) By the standards of Unger's seminar, I don't know I have hands

The question is, does the verb 'know' in a sentence like (1) L-mark for a standards of knowledge phrase?—in this case, the phrase 'by the standards of Unger's seminar'.

Linguists can be cagey about what is required for an argument or L-marked position to be realized. Some of them might hold that the position has to be syntactically realized, others will hold that an argument slot in the thematic structure of the verb is all we are really committed to. There is a range of options here, and each option suggests a different way in which contextualism might be realized in the grammar. Here is a first stab at it.

- (i) 'knows' is a context sensitive predicate without an L-marked position for standard or degree of knowledge and knowledge reports have no operators representing standards of knowledge, but knows is nevertheless a context-sensitive predicate.

(Like tense morphemes for A-theorists—they are context sensitive but there is no explicit argument place for times. See Ludlow (1999). Also see Kamp (1975) on degree modification).

- (ii) ‘knows’ does not L-mark a standards-of-knowledge phrase, but knowledge reports do have an explicitly represented operator position for standard or degree of knowledge. See Cinque (1999) on adverbs and functional heads.
- (iii) ‘knows’ has an implicit argument, but it is not syntactically realized apart from in the thematic structure of the verb.

Implicit arguments are not the mysterious shadowy presences they are sometimes made out to be. They are really nothing more than the argument slots in the argument structure, A ‘weak’ θ -criterion is all that is needed to give implicit arguments, since these are nothing more than unlinked argument roles. (Williams 1985: 314)

An implicit argument is a conceptual argument that is neither expressed syntactically nor bound to an argument that is expressed syntactically. (Jackendoff 1987: 409)

- (iv) ‘knows’ L-marks positions for standard or degree of knowledge, and these (sometimes implicit) positions are syntactically represented, although the evidence for them does **not** include binding facts.
- (v) ‘knows’ L-marks for standards of knowledge and the evidence for implicit positions *does* include binding facts.

That taxonomy may sound excessively complete, but it doesn’t even cover the position I’m defending in this paper. (i) is close, but it’s not really that the predicate is context sensitive so much as new predicates are negotiated on the fly, or in some cases we invoke predicates that have been negotiated in other settings and other contexts. The resulting position might be characterized as in (vi):

- (vi) The content of the lexical item ‘knows’ (like all lexical items) is dynamic and is negotiated on the fly from context to context. Sometimes meanings negotiated in other contexts (e.g. Unger’s epistemology class) are invoked and deployed in new contexts.

This is pretty sketchy, I know, and for this reason I will return to the topic of the dynamic lexicon in section 4. First, it will be useful to review some important issues regarding the behavior of embedded knowledge reports as discussed in my exchanges with Stanley and Hawthorne.

2. Stanley and Hawthorne

2.1 *The binding test* (Stanley 2000, 2002a, 2002b)

In the context of taking aim at versions of contextualism like (v) above, Stanley has drawn attention to the contrast between (7) and (8). (7) suggests a

kind of bound reading in which each person goes to a bar local to him/her, but (8) doesn't seem to have a similar bound reading—that is no reading corresponding to (8').

- (7) Everyone went to a local bar
- (8) Everyone knows that Chesner has feet.
- (8') Every x knows by x's standards that Chesner has feet

In Ludlow (2005) I argued that this isn't the sort of reading that a contextualist should expect—given that if it existed it would violate clause (C2) of our characterization of contextualism—that the context of the *ascriber* sets the standards. Why on earth would a contextualist expect a bound variable reading if the ascriber is the one who sets the standards?

Real tests would presumably require quantifying over knowledge *ascribers*:

- (9) Everyone_i asserted that Chesner knew_{f(i)} he had feet. (standards can vary)

For example, we might imagine that an utterance of (9) is followed by a more detailed report, as in (9).

- (10) Everyone asserted that Chesner knew that that he had feet. A said that Chesner knew with certainty that he has feet, B said that Chesner knew with some assurance that has feet, C said that Chesner in effect knows that he has feet, etc.

The definition of contextualism I borrowed from Hawthorne above is actually mute on whether we should expect this sort of phenomenon, but you might take this sort of fact as making a case for an extreme form of contextualism in which we can even locate binding facts that support the existence of comparison classes—in effect, lending support to position (v). But as we shall see, it also supports position (vi).

2.2 Hawthorne on Embedded Knowledge Reports²

Hawthorne (2004) claims that there is a problem for any contextualist account, and it turns on the following argument. Consider first, the following premises.

The True Belief schema

If x believes that P, then x's belief is true if and only if P

The False Belief schema

If x believes that P, then x's belief is false if and only if it is not the case that P.

Disquotational Schema for 'knows'

If an English speaker x sincerely utters a sentence of the form 'I know that P' and the sentence in the that-clause means that P, then x believes that he knows that P.

Hawthorne reasons as follows. Suppose a speaker S utters (11):

(11) I know I have feet

Hawthorne claims that "since the semantic value of that utterance is true, the belief you manifest by sincere acceptance is a true belief. So if [H] is a (standard) contextualist [H] is committed to saying that [S's] belief is true. But the Disquotational Schema enjoins [H] to say..."

(12) You (S) believe that you know you have feet

but then, by the True Belief Schema H deduces:

(13) You (S) know you have feet

Hawthorne concludes that "standard contextualism, in combination with the True Belief Schema and the Disquotational Schema would have [H] conclude that [S knows S has] feet. But this conclusion is altogether forbidden by the standard contextualist. For were H to sincerely accept 'You (S) know you have feet', then [H] would have a false belief since, in the scenario envisaged, the semantic value of the latter sentence is false." The problem, of course, is that I can only ever read (13) as, roughly

(13') You know-by-my-standards you have feet

The problem with Hawthorne's argument is with the disquotational schema, which doesn't for if the matrix predicate (the predicate that appears before the that-clause, as 'knows' does) is an indexical. A couple of examples illustrate why.

Consider Mr. There, who is somewhere distant from us and who utters 'Here it is the case that it is hot'. Consider the plausibility of the disquotational schema as uttered in a different location from the quoted utterance:

If an English speaker x sincerely utters a sentence of the form 'Here it is the case that it is hot' and the sentence in the that-clause means that it is hot, then x believes that here it is the case that it is hot.

Or consider Mr. Then, who utters ‘It is now the case that it is hot’ at some time before our tokening of a corresponding instance of the diquotational schema. Consider the plausibility of the disquotational schema in this case:

If an English speaker *x* sincerely utters a sentence of the form ‘It is now the case that it is hot’ and the sentence in the that-clause means that it is hot, then *x* believes that it is now the case that it is hot.

In the former case, the schema fails for the locative indexical, in the latter case it fails for the temporal indexical. Hawthorne himself fudges on the first person indexical in his own examples, swapping a variable *x* for the indexical ‘I’. Similarly, any contextualist will argue that an instance of the schema that utilizes the predicate ‘knows’ in the matrix will fail in precisely the same way.

Hawthorne suggests (p.c.) that there is something different about the ‘knows’ case—maybe in general we are more apt to disquote than we are in the case of indexicals like I, ‘here’, and ‘there’. This is surely true, however, I think this just suggests that we disquote only when the indexical parameter can be held fixed. For example, if I am the one who utters ‘I believe that I have feet’ I am quite happy to disquote ‘I’. If Mr. There is here, I am quite happy to disquote his utterance of ‘Here it is the case that it is hot’, and if Mr. Then is talking right now, we are happy to disquote his use of ‘now’. Clearly some of these happen more than others, because some of these parameters are more apt to have flipped on us than others. Will the indexical parameter in knowledge reports be less apt to flip than the person parameter in an utterance of ‘I’? Most likely. Will it be less likely to flip than the place parameter in an utterance of ‘here’? Maybe. The problem is that frequency is beside the point; the issue is that we will resist the disquotational schema in precisely those places that count—those in which we have reason to think that the standards of knowledge have shifted.

3. Stanley’s (2005) Objection

Let’s return to examples (1) and (2) from the introduction.

- (1) By the standards of Unger’s seminar, I don’t know I have hands
- (2) By the standards of Chemistry, the stuff in the Hudson River isn’t water

Stanley wants to suggest that these are not so very different in kind—that there is really nothing special about the standards talk with epistemic predicates. We get that sort of standards talk with every predicate (counting common nouns as predicates here).

The interesting thing about these constructions is that the standards talk appears to be a form of “non-restrictive modification”. It patterns with (14) and not (15).

- (14) The woman, who is a lawyer, entered the room
- (15) The woman who is a lawyer entered the room

The comma intonation suggests that ‘is a lawyer’ is not modifying ‘the woman’ but is being thrown in as a parenthetical aside.

It is instructive to note that examples like (1) and (2) are less smooth if the comma intonation is left out.

- (16) I don’t know I have hands by the standards of Unger’s seminar
- (17) The stuff in the Hudson River isn’t water by the standards of chemistry

This provides some evidence that there is no L-marking involved in these cases, because it suggests that the standards phrase is not “built in” to the thematic structure of the verb.

Another observation is that we are tempted to add emphasis on the verb ‘knows’ and ‘water’ in these case

- (1a) By the standards of Unger’s seminar, I don’t **KNOW** I have hands
- (2a) By the standards of Chemistry, the stuff in the Hudson River isn’t **WATER**

We could accomplish the same thing by using scare quotes or using the “pair of fingers” gesture to indicate scare quotes when we are speaking:

- (1b) By the standards of Unger’s seminar, I don’t “know” I have hands
- (2b) By the standards of Chemistry, the stuff in the Hudson River isn’t “water”

What is going on here? I want to suggest that there is a kind of meta-linguistic operation going on here—that we are swapping in a non-standard usage, or a homophone from a different language if you will.

In the next section I’ll say a bit about micro-languages, and the idea that when we engage with conversational partners we flesh out or sharpen up the meanings of some of the terms we are using—terms that are ordinarily massively underdetermined in meaning. Each time we converge on a sharpening or have made things sharp enough for conversation to proceed we can say that we have created a micro-language.

The idea I am developing in this paper is that not only can we create micro-languages on the fly, but we can borrow from other micro-languages

we have participated in or are aware of, so long as we signal what we are up to. One way of signaling that this is what is going on is by using expressions like ‘by the standards of’ or ‘in so-and-so’s terminology’, or ‘if you ask a Chemist...’ or ‘in the *parlance* of the Ungerian’.

The phenomenon involves what linguists call “code-switching”—incorporating terms from another language. So, for example, if I utter (18), I am code switching in at least two spots. First, I am using a bit of French when I deploy ‘*parlance*’ (either because of the shade of meaning it affords or because I’m a snob and a show-off). Second, I am code switching with the word ‘know’.

(18) In the *parlance* of the Ungerian, I don’t know I have hands

Code switching isn’t quite the same as borrowing something from one language into another. The dead giveaway is that pronunciation is usually normalized in the case of borrowing, but not code switching. If I want to borrow the word for the Italian snack called ‘bruschetta’ I will say [broo-shetta], if I am code switching I will say [broo-sket’-ta]. The claim here is that we are code switching on homophones from different micro-languages (although if people in an Unger seminar pronounced ‘knowledge’ as [newledge] we might be inclined to adopt their pronunciation, as in (18’)).

(18’) In the *parlance* of the Ungerian, I don’t [new] I have hands

That’s the basic idea, but of course this requires some further discussion of the idea of micro-languages.

4. Microlanguages³

In Larson and Ludlow (1993), Ludlow (2000, 2006), and subsequent work I’ve made the case for what I have called micro-languages that are negotiated on the fly by discourse participants. The headline idea is that the common coin view of language is badly mistaken and that discourse participants routinely mint new linguistic items and also that what common coins there are are typically “thin”—in the sense that their meanings are underdetermined and fleshed out on a case-by-case basis. Likewise in some instances we come into conversations with fleshed out word meanings that need to be thinned out for purposes of the conversational context.

Most philosophers of language suppose that the words we use are “common coins” in a broad social context—that when you learn a language, among other things you learn a set stock of shared word meanings. In the papers cited above, I’ve suggested that we reject this picture and opt instead

for the idea that many of the terms that we use are invented on the fly or at least fleshed out during individual conversations.

In effect this is a generalization of conclusions that have been reached by psycholinguists (e.g. Garrod and Anderson 1987, Brennan 1986, Brennan and Clark 1996, and Clark 1992) and their study of lexical “entrainment”—a process whereby the use of certain words—sometimes novel words—are introduced on the fly by discourse participants.⁴

Studies on entrainment undermine the myth of a common-coin lexicon by showing that individuals who overhear or witness a conversation are much less in a position to understand what is being said than are the participants. Schober and Clark (1987), for example, show that discourse participants are in a much better position than are eavesdroppers for understanding what is being said because participants are involved in the introduction of the lexical items that will be employed in the evocation of certain concepts in the conversation.

Consider, for example, how much of a lecture you can comprehend by dropping in on a course in the middle of the term. If you are not familiar with the subject matter you may well be quite lost, and not just because you lack familiarity with the objects under discussion (if it is a philosophy class you might have dropped in on an unintelligible discussion of whether tables and chairs exist). One obstacle you may face is that you are unfamiliar with the terminology in play (of course, grasp of the terminology and knowledge of the subject matter are not so easily separated). You were not involved in the entrainment process whereby certain terms were introduced into the course. In such situations you may dismiss the terms being used as “jargon,” but this is just a way of saying that you don’t understand the terms being deployed.

One important result of the entrainment experiments is that those common coins that do exist are not always in circulation, and indeed, are strategically retired and placed back into circulation depending upon the demands of the micro-language under construction. The situation is analogous to the position of the traveler who finds that various combinations of US Dollars, Euros, Yen, and Argentinean Pesos are accepted in different settings. Some are more widely accepted than others, and some can be introduced in the odd transaction with a bit of cajoling, but at the end of the day there are still establishments where only a Peso will do. Linguistic common coins are like this too, but their deployment is more strategic.

The experiments on entrainment are particularly illuminating here because they show that additional common coins are introduced into the micro-language in response to conversational demands on the discrimination of the concepts being deployed. If similar concepts are being deployed (and there is greater need to discriminate concepts and kinds of objects), there is increased pressure to reissue certain coins.

4.1 *Common coins are thin*

Linguistic common coins, whether in circulation frequently or rarely, are “thin.” By that I mean that the shared part of the lexicon consists of just hints and clues (like one finds in dictionary entries) that may help us to deploy cognitive resources to flesh out the word meanings, and the way we flesh them out will vary according to contexts and social settings. A classic illustration would be dummy terms like ‘whatchamacallit’ and ‘thingamajigger’, which are reissued often but typically with different denotations each time they are reissued.

Another example of this is the meaning of the term ‘good’. This is a widely shared common linguistic coin, but there is much to its meaning that is underdetermined. For example, it is a typical phenomenon of sports talk radio to debate which of two sports stars is better. Was Mickey Mantle better than Barry Bonds at baseball? Well, one of them hit more home runs, but the other was on more championship teams. One of them may have cheated by using steroids. Should that be a factor? What is really up for grabs here is the question of what counts as a “good” baseball player—it is about the meaning of ‘good’.

Jamie Tappenden (1999) offers a formal example of this phenomenon, introducing a language in which some meanings are open-ended and to be precisified at a later time. The language leaves “certain objects as ‘unsettled’ cases of a given predicate, in that it is open to the speakers of the language to make a further stipulation that the object is, or is not, to be counted as having the property in question.”

As Tappenden (1999) notes, these cases happen frequently both unintentionally and intentionally outside of formal languages, with an example of intentional cases coming from the realm of law:

This happens with some frequency in law: it may be convenient to stipulate a condition for only a restricted range, leaving further stipulation for the future. There have been many different reasons for such reticence: courts have wanted to see how partial decisions fly before resolving further cases, higher courts may want to allow lower courts flexibility in addressing unexpected situations, legislatures may be unable to come to the needed political compromises without leaving ‘blanks’ for courts to fill in.

Tappenden is thinking of cases in which matters are intentionally left open, but we can imagine lots of reasons why aspects of word meaning might remain open as a kind of natural default state—it may simply be too complicated to specify everything (even for an expert) or it may be that crucial aspects of word meaning depend upon facts about the world that remain open. Or it may just be that the language faculty is only accidentally suitable for communication and that for no reason in particular it just happens not to fix robust lexical meanings.

It would be a mistake, I think, to try to assimilate these cases of open meanings to that of vague predicates like ‘bald’. Many of the disputes that arise have little to do with vagueness. To see this, consider the following case from a dispute I heard on WFAN (a sports talk radio station in New York) when *Sports Illustrated* announced its “50 greatest athletes of the 20th Century.” Some listeners called in complaining that a horse—Secretariat—had made the list, while host Mike Francesa defended the choice. Clearly this is a dispute about what should be in the extension of ‘athlete’, and the callers wanted to argue that a horse had no place here. It is not as though the dispute would be resolved if Secretariat were a little bit faster or could throw a baseball, so it seems hard to imagine that these are vagueness cases.⁵

This is also a good example of a case where fleshing out the meaning of the term is up to us and our communicative partners. So, even when we are deploying a common coin term (like ‘athlete’, for example) the extension of the term (i.e. the set of all athletes) within a given context may be up for grabs and may require some form of coordination strategy - in the sports talk radio case the coordination took the form of a debate where discourse participants argued their respective cases.

At least in this narrow instance there is an obvious similarity to the legal realm, where competing parties may come together to resolve a dispute - in this case the way in which the term is to be understood with respect to the new cases in question (think of question of whether an existing patent “reads on” (applies to) some new technology). The key difference is that rather than taking place in a formal courtroom setting, these debates play out in less formal realms, ranging from sports talk radio to arguments with colleagues, friends, and partners.⁶

4.2 *Assigning meanings to common coins by jurisdiction*

Tappenden’s metaphor of court decisions can be extended in fruitful ways. Disputes over the best baseball player or whether a horse counts as an athlete are often just wheel spinning, but sometimes a consensus is achieved. This might be due to a series of rational arguments or it might be a simple case of someone asserting a claim and other participants deferring. In a bit we will look at how this kind of deference works, but first it is worth noting that when these disputes are resolved there are often jurisdictional limits.

When courts come to a decision on a particular dispute they set a precedent which may carry over into other jurisdictions. On the other hand it may not. Similarly, we may resolve a dispute or coordinate on the meaning of a term, and expect that to carry over into other micro-languages that we form. We may be disappointed to find we have to re-argue our point of view, or re-establish our credentials.

Sometimes we may not be involved in litigating the meaning of a term, but we may rather defer to someone else's usage (perhaps in the conversation, or perhaps in the greater community). To use the famous example from Putnam, we may defer to an expert on the proper individuating conditions of the expressions 'beach tree' and 'elm tree'. On the other hand, we may not. For example, we are within our rights to reject the usage of the materials scientist who tells us that the glass in our windows falls in the extension of 'liquid' and the botanist who tells us that the tomatoes in our salad are fruits and not vegetables. In the terminology of Ludlow and Friend (2003) we are resisting the "semantic reach" of these domain experts.

That's just a sketch of the project, but maybe it is detailed enough for us to see how it intersects with the issue of knowledge attributions. The idea would be that the term 'knowledge' is fleshed out in different ways in different conversational contexts—in different microlanguages. In Unger's philosophy seminar, for example it may be understood to entail extremely robust justificatory requirements. In Carl's seminar, it might have explicitly weaker requirements, and in some contexts the justificatory requirements may remain undetermined.

The kind of construction that Jason has highlighted in examples (1) and (2) suggest a resource that we have to import terms negotiated in other contexts and deploy them in our own. They thus serve as resources that enable a form of "code switching" for us. I can swap in Unger's use of 'knows', but it doesn't thereby undermine my own usage. If you will, that term is a homophone from another micro-language.

This, by the way, helps us to avoid the sort of conclusion that Lewis was driven to. On this story, knowledge isn't elusive at all. It is true that there are ways of fleshing out the term 'knowledge'—for example, Unger's—according to which none of our beliefs fall into the extensions of the verb, but if we have negotiated another meaning for 'knowledge' we can certainly talk about Unger's fleshing-out of the basic term without taking on the corresponding standards of justification ourselves.

But which one of us is really talking about *knowledge*? We both are of course! But when we say that, we are winding back to an unsharpened usage of the term 'knowledge'. We can certainly argue about admissible sharpenings and someone can dig in their heels and say that only Unger's sharpening is admissible, but this is just semantic chauvinism, not so different in kind from someone who insists that the only admissible sharpenings of 'baby' or 'person' must include fetuses.

It should also be noted that the move to the dynamic lexicon requires that we sharpen up the definition of contextualism (from Hawthorne) that we began with.⁷ If we adjust the definition to reflect our understanding of the dynamic nature of the lexicon the result is the following (where changes are italicized).

- C1'. [context-sensitive semantic values]: A given sentence *form*, say 'Chesner knows that he has feet' has different *sharpenings and hence different* semantic values relative to different contexts of utterance. In brief, the contextualist claims that the epistemic standards required for someone to count as meriting a positive knowledge ascription varies from ascriber to ascriber, with the result that one ascriber may truly *utter the form* 'He knows that he will be in Syracuse,' referring to a given person at a given time, and a different ascriber may truly *utter the form* 'He doesn't know that he will be in Syracuse', speaking of the same person at the same time.
- C2**'. According to *cheap* contextualist semantics, the ascriber's context of utterance calls the shots, so to speak: *sharpenings and hence* the standards of application for the verb *form* 'know' *are either explicitly stated* or are determined by the context in which the ascription is made and not by the context in which subject appears (unless the subject happens to be identical to the ascriber).

The changes are necessary because it is open to dispute as to whether we are talking about the same lexical items (under different sharpenings) or distinct lexical items. The talk of forms is a way of staying neutral.

5. The Error of Privileging "Absolute" Sharpenings

The theory I have sketched above would count as "contextualism on the cheap" by the lights of Stanley (2005), but it is nuanced for all that, and I think it escapes the objections that Stanley (ch. 4) raises against cheap contextualism.

First of all, let's see why it is cheap contextualism. It is "cheap" by Stanley's lights because the shift in standards does not trade on a context sensitive argument position (implicit L-marked phrase) like I proposed in Ludlow (2005), but rather because the meaning of the lexical items shift from context to context (or rather, from discourse situation to discourse situation). The view is cheap in the sense that there is nothing special about the context sensitivity of knowledge claims—they are context sensitive in the way that all of our lexical items are. The dynamic lexicon applies to all words, not just our epistemic vocabulary.

Stanley offers three objections to cheap contextualism. First, he argues that if cheap contextualism takes the form of a simple lexical ambiguity, we want to know why the ambiguity is not reflected by distinct words in some other languages. We can call this argument "Kripke's Cross-linguistic manifestation argument." The idea (from Kripke 1977) is that canonically ambiguous words like 'bank' typically have their distinct meanings translated as different words in other language. We should expect the same from any ambiguity we are prepared to posit.

I don't think this is an argument so much as guiding methodological principle, but for all that it is a pretty good methodological principle and I'm happy to endorse it here. I don't think the proposal I'm making here flouts it, however, since I'm not suggesting a traditional ambiguity thesis—one in which the two meanings associated with the expression are radically different (and probably come from the accidental phonological convergence of two distinct borrow words). It would be absurd to think that a particular fleshing out of a linguistic common coin should have a distinct morphological realization in some other standard language of the world.

Stanley's next two objections take the form of a dilemma. Once you set aside the ambiguity thesis, then you appear to be endorsing a "loose talk" version of contextualism. But loose talk how? Either you are thinking that there is a literal meaning and the loose talk diverges from that meaning but is licensed somehow, or you are opting for some sort of semantical notion of loose talk in which the meaning itself is somehow loose.

Stanley (p. 82) rightly points out that the first horn of the dilemma—saying that knowledge ascriptions are literally false but we are still entitled to assert them because they are pragmatically useful—effectively gives up the game. It is precisely the position that many skeptics like Unger take: Knowledge claims are literally false but assertable.

If we avoid the first horn of the dilemma and take the second, then Stanley is ready to drop the hammer on us, arguing that semantic "loose talk" becomes nothing more than deviation from the literal meaning (like saying 'It's 3 o'clock' when it is really 3:03) and if that's the case, we still collapse into something like a pretense theory or a warranted assertability theory.

Stanley's argument rests on a nearly universally held assumption. The assumption is that any notion of semantic loose talk is talk that is somehow parasitic on or generated from a more precise "correct" meaning (perhaps with a loosey-goosey operator which means something like "this rounds off to the truth"). So for example, '3:00 o'clock' really means 3:00 o'clock on the nose, but we understand that for certain purposes 3:03 is close enough; we apply the loosey-goosey operator and what we say is true. Alternatively, we might think that the meaning of the term itself could be loosened up. For example, we could say that 'flat' means absolutely flat, but for certain purposes the meaning of the term can be relaxed to allow irregular surfaces within the extension. If that is how the loose talk story goes Stanley has a point, even if more work is necessary to nail down the conclusion. It is difficult (though maybe not impossible) to see a difference between the role played by a loosey-goosey operator and a pretense operator. It is also difficult to see the difference between an operation that loosens up word meanings and one which keeps word meanings absolute, but allows them warranted assertability conditions in cases where they don't literally apply. These certainly look like distinctions that could collapse under

serious philosophical scrutiny. The problem for Stanley is that the theory I am advancing does not buy into the basic assumption at work in this argument.

The theory I am offering diverges from such a ‘loose talk’ story in important ways. First, I reject the idea that the core meaning is the absolute meaning (e.g. absolutely flat, or perfectly hexagonal, or known with Cartesian certainty)—it is more consistent with my theory to suppose that the core meaning is an underdetermined meaning—for example that the predicate ‘flat’ pertains to some underdetermined class of surfaces, ‘hexagonal’ to an underdetermined class of shapes, and ‘knows’ to an underdetermined class of psychological states, and that the precisifications *include* technical uses like “absolutely flat,” “perfectly hexagonal,” and “Cartesian knowledge.”

For the record I think this point even extends to cases like saying it is 3:00 o’clock. The expression ‘3:00 o’clock’ doesn’t mean precisely 3:00 o’clock on the nose unless we sharpen it to mean that for certain purposes. The meaning that it is 3:00 on the nose is just one sharpening of the expression and not the core meaning. This point even extends to expressions like ‘now’. We needn’t take ‘now’ to mean a time slice at some exact time of utterance, it can just as easily include vast stretches of time (as in the ‘Universe is expanding now’). The “looser” meanings of ‘now’ are not derived from the precise meanings, but they are co-equal sharpenings of the meaning of the expression. In other words, there is nothing specious about “the specious present.”

Of course it is natural enough to see why we expect core meanings to be precise. If we think of a semantic theory as a kind of scientific theory that gives the meanings of expressions, then it is natural to expect a kind of scientific precision in the meaning-stating vocabulary of our theories. If we think of the theory of meaning as looking like a T-theory, then we expect the right hand sides of the axioms and theorems to deliver precise meanings. And we not only expect the theory to deliver meanings that are precise, but meanings that are simple or absolute in some sense. But I believe this expectation is unjustified. On the dynamic lexicon approach, if we provide a semantic theory for a community of language users (or even one language user), then the meanings delivered by the right hand sides of our axioms and theorems will have to be underdetermined. It is no virtue to introduce precision if it serves to mischaracterize the meanings of the expressions we are investigating. A good semantic theory must lift the indeterminacy of the object language into the meta-language of the theory (just as tensors lift tense into the metalanguage). Adding precision to axioms and theorems is simply to mis-state the general meanings of the expressions, which are not sharp, although they will be sharpened up (or loosened) to the degree necessary in various discourse contexts.

6. Embedded Attitudes Redux

Now let's return to the issue of embedded attitude reports, beginning with the arguments raised by Stanley in examples (7) and (8), repeated here.

- (7) Everyone went to a local bar
 (8) Everyone knows that Chesner has feet.

On the version of cheap contextualism I've been advocating, the meaning of 'knows' is negotiated between the ascriber and his/her audience. It doesn't shift between values of the quantifier 'everyone'. Accordingly, you aren't predicated to get a reading like (8').

- (8') Every x knows by x 's standards that Chesner has feet

My proposal also dovetails with clause (C2) of our characterization of contextualism—that the context of the *ascriber* sets the standards. This isn't exactly right—it isn't the context so much as the discourse participants in the context that establish the meaning of the term 'knows' and ipso facto the relevant standards.

What of cases like (9) where we explicitly quantify over subscribers?

- (9) Everyone _{i} asserted that Chesner knows _{$f(i)$} he has feet. (standards can vary)

As I indicated I do think that this has a reading as in (10), but maybe it would be more accurate to gloss the idea as in (9') where we make it explicit that each reporter is deploying a different sense of 'knows'.

- (9') Everyone asserted that Chesner knows that that he has feet. Unger said that Chesner knows^u that he has feet, Lewis said that Chesner knows^L that has feet, Carl said that Chesner knows^G that he has feet, etc.

Here the superscripted verbs don't indicate terms from the respective idiolects of these agents, but rather terms from micro-languages that they may have established with different audiences at the times of their respective reports.

In this case we can think of the mechanics of (9) working as follows: 'everyone' in conjunction with the verb 'asserted' introduces a kind of substitutional quantification taking homophones from different microlanguages as the substituends, so that for each person there is a version of 'knows' such that that person asserted that Chesner "knows" in the relevant sense that he has feet. The picture is as in (9*)

- (9*) For everyone x , there is a verb Y in the microlanguage of a knowledge report by them that is a fleshing out of the common coin verb ‘knows’, s.t. x asserted that Chesner Y ’s that he has feet

Example (9*) is in effect saying that if we borrowed terms from these microlanguages we could generate a true report of the form ‘Chesner “knows” that he has feet’ for each person x .

We can similarly rethink our response to Hawthorne’s argument by generalizing it in the following way. It isn’t just that the disquotational schema fails for indexicals. The disquotational schema only holds if our vocabulary is relevantly stable across micro-languages. Inferences based on utterances by ascribers with different sharpenings of the meaning of ‘knows’ are precisely the kinds of cases where the disquotational schema will fail.

7. Rethinking the Move Space

The theory of the dynamic lexicon predicts that a predicate like ‘knows that P ’ is going to have a broad range of precisifications (likewise any other predicate) and that some of these will incorporate more stringent standards of knowledge and others will incorporate lower standards of knowledge (it will doubtless vary on other dimensions as well).

If this is right, then it raises questions about the usual way we as philosophers carve up the move space. Is the view somehow incompatible with Stanley’s Interest Relative Invariantism? Not necessarily; there is nothing to say that the way we sharpen the predicate ‘knows’ isn’t sensitive to practical interests, some or all of the time. The invariantism part is less clear, but notice that if we think of each micro-language specific precisification as being a genuine lexical item, then we can think of each of these fleshed out lexical items as being semantically invariant.

Similar considerations apply to the “minimal semantic” approach advocated by Cappelan and Lepore (2004)—we needn’t introduce hidden indexicals or even allow that a lexical item, once fleshed out, is in any sense shifty. Once fleshed out, it means what it means once and for all.

This having been said, I don’t doubt for a minute that there are plenty of implicit L-marked positions in natural language, not least because I still believe in implicit comparison classes in adjectival constructions as I proposed in Ludlow (1989). I agree it is more controversial, but I am not entirely convinced that there aren’t also implicit arguments in knowledge reports—possibly including parameters for standards of evidence, justification, and perhaps even degree of knowledge. These are posited for linguistic reasons, however, and not to account for context sensitivity phenomena. Furthermore, I can easily imagine that our knowledge reports are sensitive *both* to the ways we have fleshed out the meaning of ‘knows’ for purposes of conversation *and* to various parameters as I argued in Ludlow (2005).

All of this means that the battle lines as currently drawn in discussions of language and epistemology are not as sharp as we might have supposed. That might be frustrating, but perhaps it also leads us to more productive work in this area.

8. Negotiated Word Meanings and the Nature of Concepts

Ram Neta (this volume) offers several objections to cheap contextualism, one of which I will touch on here because it provides a segue into my concluding discussion on negotiated word meanings and nature of concepts. Neta observes correctly that sauce for the goose is sauce for the gander and that if the dynamic lexicon is applicable to terms like ‘knowledge’ and ‘water’ then surely it is just as applicable to terms like ‘meaning’, ‘precisification’, ‘micro-language’ etc.—in effect to all of the theoretical vocabulary introduced in this paper.

Does my position cannibalize itself? Here is one way to make Neta’s objection vivid: Someone—say Unger—might refuse to bite on the dynamic lexicon. He might argue in the following way: “when *I* talk about the meaning of a term I am using the term ‘meaning’ in a way that just doesn’t allow this sort of nonsense. The meaning of a word is what it is, I’m telling you about the meaning of ‘knows’, and it has an invariant meaning. This isn’t something that is up for grabs.”

One might anticipate a related version of this objection as well (this was first brought to my attention by Jennifer Lackey). It goes like this: Ludlow’s dynamic lexicon is fairly promiscuous—it looks like any sort of meaning goes. So while Ludlow might deploy a wimpy contextualist version of ‘knows’ there is no reason for Unger to do so. How on Earth are we going to move Unger (or any skeptic) off of his position? For that matter, why should we follow Ludlow on the meaning of ‘meaning’ or ‘precisification’ or ‘microlanguage’?

Both versions of the objection miss the mark. The mere fact that there is variation in meaning of these expressions does not mean that anything goes and certainly does not mean that anyone is entitled to stubbornly dig in on the meaning of a term. To the contrary, the process by which we become entrained often involves argumentation, and argumentation is a normative activity. That is, norms govern the way we negotiate (litigate) word meanings.

While we routinely deploy thin coins that have to be fleshed out, we also deploy coins that need to be clipped and trimmed from time to time. Typically, the latter happens when two discourse participants come to a conversation and discover that they are deploying incompatible definitions of a term (like ‘torture’ or ‘person’ or ‘knowledge’ or ‘meaning’). It’s not that they need to amplify their definitions as in the Tappenden cases; they are genuinely in conflict and someone has to give if the conversation is going to make progress.

There are a lot of ways in which disputes about meaning could be resolved. Power relations could play a role (one participant might defer to a professor or religious leader or political leader). Alternatively both sides could dig in and we might have an argument (just as the sports talk radio listeners did about ‘best baseball player’).

The fact that the issue is fundamentally a matter of which definition to deploy does not mean there is nothing normative about the process by which we argue and (hopefully) come to a resolution. Hopefully, if the definition is important (Does a fetus fall under the extension of ‘person’? Does water boarding fall under the extension of ‘torture?’), it won’t just be power relations that resolve our lexical disputes.

A key element to the theory of the dynamic lexicon is the idea that when we are in conflict on an important term we argue out the best definition. There has been much work on the theory of argumentation by logicians in Europe recently and it would be impossible to offer even a glimpse at what is now a robust body of literature, but here is one theme that emerges from the literature. Part of what we do when we argue (for example, about the meaning of a term) is try to wind back to shared assumptions and show how the shared assumptions dictate that we opt for one conclusion over another.

For example, if it is a shared assumption that Ted Williams and Ty Cobb (famous baseball players who were on no championship teams) are among the greatest baseball players of all time, then we might argue that a definition of ‘great ballplayer’ that requires the player to have won championships would not be a good definition—it would exclude these individuals that we take to be canonically great ball players. Similarly, we might appeal to shared beliefs or philosophical assumptions in arguing for one definition over another.

Another objection due to Jennifer Lackey (p.c.) now becomes salient. Given that we are in a position of arguing about the right definition of ‘knowledge’ how have we advanced our position over the traditional debates in epistemology where contextualists and invariantists argued about the *concept* of knowledge? What’s the difference? There is a two part answer to this question.

First, we are often told things to the effect that contextualism is not really a theory about knowledge—the concept of knowledge—but is merely about how the word ‘knowledge’ is *used* in certain contexts. But one tenant of the dynamic lexicon proposal (not argued for in this paper) is that such reasoning is mistaken since there really is nothing more to the concept of knowledge than there is to deciding how we are going to define ‘knowledge’ relative to a background of shared interests. There is no mysterious concept of knowledge in Plato’s heaven to which we as philosophers have special access. In effect, we take away one primary argument against contextualism—the argument that it’s only about language and not about knowledge itself.

The second part of the answer is sociological. Once we puncture the myth of the concept of knowledge and understand that we are in a conflict over the best way to refine the meaning of the term ‘knowledge’ relative to some shared tasks or interests, we are more apt to be critical of arguments that rest heavily on the weight of authority. Philosophers since Descartes may have used the term ‘knowledge’ in an invariantist way, but why should we? What special claim do these experts have on us, so that we must feel compelled to reject the usual contextually sensitive uses of knowledge when we are in philosophical discussion? There may be a Cartesian tradition about the term ‘knowledge’ which takes it to have an invariantist meaning, but that tradition counts for little—or at least we can rightly ask why it counts for anything in our current discussions in epistemology.

This doesn’t mean that the dynamic lexicon by itself leads directly to contextualism about knowledge, but it does three things for us. First, it provides us a plausible version of “cheap contextualism” that escapes the difficulties that have afflicted other versions of contextualism. Second, it undermines the claim that contextualism is defective because it only speaks to linguistic usage. Finally, it puts us in a position to challenge the semantic reach of philosophical tradition. The net result is that while contextualism doesn’t fall out for free from the dynamic lexicon, it is much less expensive than we might have otherwise supposed.

Notes

1. Earlier versions of this paper were presented at the Aberdeen Scotland conference on Language and Epistemology (2007), at the University of California, Davis (2008), at Northwestern University (2008), and at the 2008 Pacific Division Meeting of the APA. In addition to the audiences at those venues, I am indebted to David Braun, John Hawthorne, Jennifer Lackey, Mark Phelan, Ram Neta, Jason Stanley, Tim Sundell, and Chris Tilman for helpful discussion.
2. Thanks to Cody Gilmore for discussion that led to a sharper formulation of these examples.
3. This section borrows from material in Ludlow (2006).
4. The term ‘entrainment’ appears to be a metaphorical extension of the use of the term to describe the process by which coupled oscillators become synchronized. See Watts (1998; 225–226) for discussion. We will come back to this idea a bit later.
5. John Hawthorne suggested to me that maybe the relevant Sorites scale is along another dimension: “maybe if Secretariat were a bit more human.” Suppose Secretariat could talk like Mr. Ed, do basic math, run on two legs, etc. Would this change minds? Um, I think not. And I have no idea if Hawthorne was serious or just messing with my head at 2 AM.
6. There are technical issues that I am avoiding here, not least of which is the logic of underspecification. How do inferences work when meanings are incomplete or underspecified? For a sample of work on this topic see van Deemter and Peters (1996).
7. The need for this revision was brought to my attention by Carl Ehrett.

References

- Brennan, S., 1996. "Lexical Entrainment in Spontaneous Dialog." Proceedings, 1996 International Symposium on Spoken Dialogue, 41–44. Philadelphia, PA: ISSD-96.
- Brennan, S. and Clark, H., 1996. "Conceptual Pacts and Lexical Choice in Conversation." *Journal of Experimental Psychology: Learning, Memory and Cognition* 22, 1482–1493.
- Bresnan, J., 1982. *The Mental Representation of Grammatical Relations*, Cambridge: MIT Press.
- Cappelen, H., and E. Lepore, 2004. *Insensitive Semantics. A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.
- Castaneda, H-N., 1967. "Comments." in Nicholas Rescher, ed., *The Logic of Decision and Action* Pittsburgh: University of Pittsburgh Press.
- Chomsky, N., 1986. *Barriers*. Cambridge: MIT Press.
- Clark, H., 1992. *Arenas of Language Use*. Chicago: CSLI Publications and The University of Chicago Press.
- Davidson, D., 1967. "The Logical Form of Action Sentences." In N. Rescher (ed.), *The Logic of Decision and Action*. Pittsburgh: University of Pittsburgh Press.
- Davidson, D., 1985. "Adverbs of Action." In Vermazen, B., and M. Hintikka (eds.), *Essays on Davidson: Actions and Events*. Oxford: Oxford University Press.
- DeRose, K., 1995. "Solving the Skeptical Problem." *Philosophical Review* 104, 1–52.
- Garrod, S., and Anderson, A., 1987. "Saying What Your Mean in Dialogue: A Study in Conceptual and Semantic Co-ordination." *Cognition* 27, 181–218.
- Jackendoff, R., 1987. "The Status of Thematic Relations in Linguistic Theory," *Linguistic Inquiry* 18, 369–411.
- Kamp, H., 1975. "Two Theories of Adjectives." In E. Keenan (ed.): *Formal Semantics of Natural Language*. Cambridge UP, Cambridge.
- Kripke, S., 1977. "Speaker's reference and semantic reference." In *Contemporary Perspectives in the Philosophy of Language*, edited by P. French, T. Uehling, and H. Wettstein, 255–276. University of Minnesota Press.
- Larson, R., 1986. "Implicit Arguments in Situation Semantics." *Linguistics and Philosophy* 11, 169–201.
- Larson, R., and P. Ludlow, 1993. "Interpreted Logical Forms." *Synthese* 96, 305–355.
- Lewis, D., 1979. "Score Keeping in a Language Game." *Journal of Philosophical Logic* 8, 339–359.
- Lewis, D., 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74, 549–567.
- Ludlow, P., 1989. "Implicit Comparison Classes." *Linguistics and Philosophy* 12, 521–533.
- Ludlow, P., 1995. "Logical Form and the Hidden Indexical Theory: a Reply to Schiffer." *Journal of Philosophy* 92, 102–107.
- Ludlow, P., 1996. "The Adicity of 'Believes' and the Hidden Indexical Theory." *Analysis* 56, 97–102.
- Ludlow, P., 1999. *Semantics, Tense and Time: An Essay in the Metaphysics of Natural Language*. Cambridge: MIT Press.
- Ludlow, P., 2000. "Interpreted Logical Forms, Belief Attribution, and the Dynamic Lexicon." In K.M. Jaszczolt (ed.) *Pragmatics of Propositional Attitude Reports*. Elsevier Science, Ltd.
- Ludlow, P., 2005. "Contextualism and the New Linguistic Turn in Epistemology." In G. Preyer and G. Peters (eds.) *Contextualism in Philosophy*. Oxford: Oxford University Press.
- Ludlow, P., 2006. "The Myth of Language." *Croatian Journal of Philosophy* 18, 385–400.
- Ludlow, P., and S. Friend, 2003. "Disagreement and Deference: is Diversity of Opinion a Precondition for Thought?" *Philosophical Perspectives 17: Language and Philosophical Linguistics*, 115–139.
- Neta, R., this volume. "How Cheap Can You Get?" *Philosophical Issues*, 130–142.

- Parsons, T., 1990. *Events in the Semantics of English*. Cambridge: MIT Press.
- Parsons, T., 1990. *Events in the Semantics of English*. Cambridge: MIT Press.
- Schober, M., and Clark, H., 1989. "Understanding by addressees and overhearers." *Cognitive Psychology* 21, 211–232. Reprinted as chapter 6 of Clark (1992).
- Stanley, J., 2000. "Context and Logical Form", *Linguistics and Philosophy* 23, 391–434.
- Stanley, J., 2002a, "Making it Articulated", *Mind and Language* 17, 149–68.
- Stanley, J., 2002b. "Nominal Restriction", in *Logical Form and Language*, G. Peters and G. Preyer, ed., Oxford: Oxford University Press, 365–388.
- Stanley, J. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stanley, J., and Z. Szabo, 2000. "On Quantifier Domain Restriction", *Mind and Language* 15, 219–261.
- Tappenden, J., 1999. "Negation, Denial and Language Change in Philosophical Logic" In D. Gabbay and H. Wansing, (eds.) *What is Negation?* Dordrecht: Kluwer.
- Unger, P., 1984. *Philosophical Relativity*. Minneapolis: University of Minnesota Press.
- van Deemter, K., and S. Peters, 1996. *Ambiguity and Underspecification*. Stanford: CSLI Publications.
- Vendler, Z., 1967. *Linguistics in Philosophy*. Ithaca: Cornell University Press.
- Vendler, Z., 1972. *Res Cogitans: an Essay in Rational Psychology*. Ithaca: Cornell University Press.
- Watts, D., 1998. *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton: Princeton University Press.
- Williams, E., 1985. "PRO and the subject of NP," *Natural Language and Linguistic Theory* 3, 297–315.
- Wittgenstein, L., 1969. *On Certainty*. New York: Harper and Rowe.