

How Cheap Can You Get?¹

Ram Neta
UNC-Chapel Hill

According to a contextualist account of knowledge ascriptions, it's possible for both Skeptic's assertion of "Moore doesn't know (at a particular time t_0) that he has hands" and Normal's simultaneous assertion of "Moore does know (at t_0) that he has hands" to be true, so long as these assertions are issued in different contexts. That's because the truth-conditions of such knowledge ascriptions (or denials) are fixed partly by features of the context in which those ascriptions (or denials) are issued.

Now, there are many different ways of implementing this general contextualist strategy. What I will be concerned to do here is to assess Peter Ludlow's way of implementing the strategy, or what Ludlow, following Jason Stanley, calls "cheap contextualism". What makes Ludlow's contextualism "cheap" is that it is not at all *ad hoc*: it predicts the context-sensitivity in the truth-conditions of knowledge ascriptions not by appeal to any special feature of *knowledge ascriptions*, but rather by appeal to a very general phenomenon that Ludlow claims to find in language. This phenomenon is described in Ludlow's theory of "the dynamic lexicon". The cheapness of Ludlow's contextualism is a result of the generality of the phenomenon from which he derives the semantic context-sensitivity of knowledge ascriptions. The idea is that the more general the phenomenon from which a theorist claims to derive contextualism, the cheaper the resulting contextualism. It's "cheaper" in the sense that the theorist need not be committed to anything very specific about *knowledge ascriptions*. The "cheap" contextualist can thus say very little about knowledge ascriptions in particular, and still earn the right to be a contextualist.

So, what is this general phenomenon that is described in Ludlow's theory of the dynamic lexicon, and from which Ludlow claims to derive contextualism? According to the theory of the dynamic lexicon:

"the common coin view of language is badly mistaken ... discourse participants routinely mint new linguistic items ... what common coins there are are 'thin' – in the sense that their meanings are underdetermined and fleshed out on a case-by-case basis." (Ludlow forthcoming, section 4.0)

What Ludlow means by "the common coin view of language" is the view, held by most philosophers of language, which says that the words of our public language have, by and large, determinate meanings independently of the specific context in which they are used, and, in general, when we use these words for communicative purposes, the contribution that these words make to the meaning of what we say is fixed in part by the determinate contextually invariant meanings of those words.

Ludlow offers a number of considerations intended to show that this common coin view of language is false. For my purposes here, I will sort the considerations that

he offers into two groups. The first group consists of data from entrainment studies. What such data reveal, according to Ludlow, is 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.” (Ludlow forthcoming, section 4.0)

Now, I do not doubt that discourse participants are, in general, in a much better position than are eavesdroppers for understanding what is being said in a discourse. But there are many possible explanations for this, and most of these possible explanations are completely independent of whether the theory of the dynamic lexicon is true. For example, discourse participants may have better epistemic access than eavesdroppers do to the facial and gestural expressions of the speaker, or to the history of the speaker, or to the subject matter of the conversation. If discourse participants do indeed have better epistemic access to these things than do eavesdroppers, then, whether or not our lexicon is dynamic in the way that Ludlow hypothesizes, this difference in epistemic access will put discourse participants in a better position to figure out anaphoric antecedents, to figure out the referents of deictic expressions (including discourse deictic expressions), and to figure out a wide range of implicatures being generated by the speaker. Also, the speaker probably has much better epistemic access to the facial and gestural expressions of other discourse participants than she does to the facial or gestural expressions of bystanders – and other discourse participants know this, and can therefore take this fact into account (along with the facts about what their own facial and gestural expressions are) when figuring out why the speaker says what she says. So, independently of whether or not the lexicon is dynamic, there are many different ways in which discourse participants are in a better epistemic position with respect to a conversation than are eavesdroppers. It’s not clear why the task of explaining entrainment data puts any rational pressure upon us to accept Ludlow’s theory of the dynamic lexicon.

But the considerations from entrainment data are just one group of considerations that Ludlow uses to argue for the dynamic lexicon. There is another group of considerations to which Ludlow appeals. This second group of considerations involves appeal to examples of discourses that may seem very directly to confirm the theory. In these examples, a discourse begins with some term or other (e.g., an ordinary term such as “good” or “athlete”, or a term from a formal language or a legal code) having what seems to be a less than fully determinate meaning, thereby giving latitude for a variety of fuller specifications of its meaning later on. Thus, Ludlow is arguing for the dynamic lexicon by appeal to seemingly clear *examples* of what we may call “lexical dynamics” (the process of fleshing out the meaning of a thin common coin word, so as to make it more fully determinate, at least for the purposes of a particular discourse).

Since I have already expressed my doubts about Ludlow’s use of the entrainment data against the common coin view of language, I will not discuss that argument any further. For the rest of this paper, I will treat Ludlow’s whole case against the common

coin view of language as resting on his appeal to examples of lexical dynamics: if his argument for the dynamic lexicon is successful at all, it is successful there.

Now, what I want to argue here is that Ludlow's argument for the dynamic lexicon is subject to a trilemma. The trilemma that I have in mind does not show that Ludlow's theory of the dynamic lexicon is false. Rather, it shows that, *if* Ludlow's theory is true, *then* there is an obstacle to Ludlow's attempt to argue for the theory by appeal to examples. More precisely, if Ludlow's theory is true, then his way of arguing for it, by appeal to examples, cannot give much support to that theory, and cannot give much support to his effort to earn contextualism on the cheap. Or so I will now argue.

Before embarking on this argument, I should stress that I will *not* be arguing that Ludlow's theory is false, or even that we should reject Ludlow's theory. I am somewhat inclined to think that Ludlow's theory of the dynamic lexicon is *true*. So the truth of Ludlow's theory is not at all what is at issue here. Rather, what is at issue is whether or not Ludlow can effectively argue for his theory by appeal to the kinds of examples that he cites. What I will argue here is simply that Ludlow cannot effectively argue for his theory in this way. If Ludlow's theory is to be defended, it must be defended in a different way. In order to understand this point, it might help us to draw an analogy between Ludlow's theory of the dynamic lexicon and the scientific theory of the electron. It has long been established that electrons are negatively charged subatomic particles. But this fact could not have been established simply by appeal to examples of negatively charged electrons. A physicist at the turn of the twentieth-century could not possibly have shown that electrons are negatively charged subatomic particles by reasoning as follows: "here's one electron that's a negatively charged subatomic particle, and here's another electron that's a negatively charged subatomic particle, and here's yet another one ...". What showed that electrons are negatively charged subatomic particles was not any such appeal to examples of electrons that are negatively charged subatomic particles: rather it was, *inter alia*, J.J. Thomson's experimental argument (using a glass tube containing an anode and a cathode) for the conclusion that cathode rays are negatively charged streams of particles, whose mass to charge ratio is much smaller than the mass to charge ratio of a hydrogen ion. Since electrons are not observable, their properties must be established by means of such indirect tests. Analogously, I claim, if Ludlow's theory of the dynamic lexicon is true, then it is not *observably* true, and so establishing its truth requires indirect tests. If Ludlow's theory of the dynamic lexicon is true, then it cannot, I will argue, be established simply by appeal to examples.

I'll begin by supposing that Ludlow's theory is true, and then I'll consider what follows from that supposition. So let's suppose that Ludlow is right when he says, in the passage above, that "what [lexical] common coins there are are 'thin' – in the sense that their meanings are underdetermined and fleshed out on a case-by-case basis." This is a very broad generalization, and, to judge from the range of examples that he gives, Ludlow means for this generalization to apply to a great many words of our public language. In fact, this generality is just what makes Ludlow's contextualism so "cheap". But just how sweeping is Ludlow's generalization? Or, to put it another way: just how cheap is Ludlow's contextualism? The trilemma that confronts Ludlow's position is built

around the following question: are the meanings of such expressions as “meaning”, “underdetermined”, and “fleshed out on a case-by-case basis” – are these meanings themselves underdetermined, and to be fleshed out on a case-by-case basis? The answer to this question must be either “yes” or “no” or “sort of”. Let’s consider each of these three options.

First horn: the meaning of “meaning” is underdetermined

Suppose the answer to our question is “yes, the meanings of ‘meaning’, ‘underdetermined’, and ‘fleshed out on a case-by-case basis’ are themselves underdetermined, and to be fleshed out on a case-by-case basis.” If this is the answer to our question, then let’s re-evaluate Ludlow’s argument for the dynamic lexicon. (Recall that, for the rest of this paper, I’m setting aside the argument from entrainment data, and focusing exclusively on his argument from examples.) Ludlow argues for the dynamic lexicon by appeal to several examples, some from natural language, some from the semantics of formal languages, and some from legislative practice. Consider, for instance, Ludlow’s discussion 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’.” (Ludlow forthcoming, section 4.0)

I take it that this passage is intended to serve as a consideration in favor of Ludlow’s theory of the dynamic lexicon by describing an example of which that theory seems clearly to hold true. The example is this: prior to a particular debate about the relative merits of Mickey Mantle and Barry Bonds, there is nothing about the meaning of “good” (or, correlatively, of “better”) that fixes it that one of the two men counts as a “better” baseball player than the other. But the meaning of “good” (and thus of “better”) can be so fleshed out in the course of a particular conversation that, at least within the context of that conversation, the question of which of them is a better baseball player acquires a determinate answer. Of course, to say that the meaning of “good” is underdetermined is not to imply that there is nothing about the meaning of “good”, prior to this particular conversation, that fixes it that Barry Bonds is a better baseball player than, say, *I* am. No matter precisely how we flesh out the meaning of “good”, so long as it is a fleshing out of that meaning and not an altogether new coinage, Barry Bonds will very clearly count as a much, much better baseball player than I.

Now why should we accept Ludlow’s claim that, prior to the aforementioned debate about Mantle and Bonds, the meaning of “good” is not determinate enough to fix it that one of the two players is “better”? And why should we accept his claim that, in the course of the debate, the meaning of “good” can be fleshed out so as to fix it? Instead of

accepting these claims, why not say, for instance, that the meaning of “good” *is* fully determinate prior to the conversation about Mantle and Bonds, and that what happens in the course of the conversation is simply that the conversational participants are debating, or perhaps trying to figure out (with whatever level of success), which of the two players really is better? Ludlow doesn’t explicitly answer these questions. I assume that’s because he thinks that his own claims about the case are evidently true. In other words, (I assume) Ludlow thinks that it is obvious – at least to the unprejudiced eye – that the meaning of “good” is not determinate enough to fix it that one of the two players is “better”, and that we can flesh out the meaning in the course of a conversation so as to fix it that one of them counts as “better”, at least on the fleshed out meaning.

Let’s refer to Ludlow’s description of the case in question as the “dynamic” description, and let’s refer to the competing description – the description according to which the meaning of “good” is fully determinate before the Mantle/Bonds conversation, and the conversation is simply a (possibly inconclusive) debate about which of the two men is the better baseball player – as the “static” description of the case. I assume that Ludlow thinks that the dynamic description of the case is evidently correct, and that the static description of the case is evidently incorrect, and that’s why he is able to adduce the case as a data point in favor of his theory of the dynamic lexicon. Let’s contrast the dynamic and the static descriptions of the Mantle/Bonds case with a third description, which I’ll call the “magical” description. The magical description says that the truth-value of the sentence “Bonds is a better baseball player than Mantle” is determined not by the meanings of our words, nor by anything that happens in our conversations, but rather by the wishes of invisible elves.

Now, the magical description of the Mantle/Bonds case is obviously ludicrous. In fact, it’s as obviously ludicrous as the claim that Bonds is no better a baseball player than I am. But neither the static description nor the dynamic description of the Mantle/Bonds case is obviously ludicrous, or even obviously false. However, if Ludlow’s theory of the dynamic lexicon is true, and if the meanings of “meaning” and of “underdetermined” and of “flesh out” are all themselves underdetermined (as they are, according to the present horn of the trilemma), then why should we regard the choice between the dynamic description of the case and the static description of the case as *any more settled in advance by the meaning of “meaning”* than the choice of Bonds or Mantle as the better baseball player is settled in advance by the meaning of “good”? Recall that, on Ludlow’s view, prior to the Mantle/Bonds conversation, even if the meaning of “better” is determinate enough to fix it that Bonds is a better baseball player than I am, still, the meaning of “better” is not determinate enough to fix whether Bonds is a better baseball player than Mantle is. The meaning of “better” need not be fleshed out in order for there to be a truth of the matter whether Bonds is better than I am, but the meaning of “better” must be fleshed out in order for there to be a truth of the matter whether Bonds is better than Mantle. But then, if the meanings of “meaning” and of “underdetermined” and of “fleshed out” are all themselves underdetermined, then, even if their meanings don’t need to be fleshed out in order for there to be a truth of the matter as to whether the magical description of the Mantle/Bonds case is correct, nonetheless, don’t their meanings need to

be fleshed out in the course of a conversation in order for there to be a truth of the matter as to whether the dynamic or the static description of the Bonds/Mantle case is correct?

Perhaps Ludlow will say that, although the meanings of “meaning”, “underdetermination”, etc. do need to be fleshed out in the course of a conversation in order for there to be a truth of the matter as to whether the dynamic or static description of the Bonds/Mantle case is correct, in fact the field of lexical semantics *has* successfully fleshed out the meanings of these terms. Lexical semanticists have thereby made it the case that there is a truth of the matter as to whether the dynamic or static description of the Bonds/Mantle case is correct. Now, this claim may be true. But first, it would have to be shown. And second, if the claim were shown to be true, then Ludlow might have grounds on the basis of which he could argue that the static descriptions of the relevant cases are not true descriptions. But then what is bearing the weight of Ludlow’s argument for the theory of the dynamic lexicon is not his dynamically described examples, but rather his independent argument that static descriptions are not true of those examples. And so, even if Ludlow goes this route and claims that lexical semanticists have successfully negotiated a determinate meaning for “meaning”, “underdetermine”, and so on, it still turns out that his appeal to dynamically described examples does not provide the argument for his theory of the dynamic lexicon. I will return to this point in my discussion of the second and third horns of the trilemma below.

In sum, if the meanings of “meaning” and of “underdetermined” and of “flesh out” are themselves underdetermined, then it’s not clear why the choice between Ludlow’s dynamic description of the Bonds/Mantle case and the competing static description of the case is any more settled in advance than is the choice between Bonds and Mantle as the better baseball player.

Precisely the same consideration applies to the other examples that Ludlow gives in support of his theory of the dynamic lexicon: in order for those examples to support his theory, they must satisfy a certain description, and if his theory of the dynamic lexicon is true, then it’s not clear why it would be settled in advance that those cases do satisfy that description. This applies just as much to the examples from the semantics of formal languages, and from legislative practice, as it does to the examples (like the Mantle/Bonds case, or the case of the dispute about whether the race horse Secretariat is an “athlete”) from ordinary language. Even if we don’t stipulate a fully determinate meaning for a given expression in a formal language, or for a given expression in the legal code, it doesn’t follow that the expression lacks a fully determinate meaning. Meanings of expressions in natural language do not generally come from stipulations alone, so why can’t expressions in a formal language, or in a legal code, acquire fully determinate meanings from things other than our stipulations? If we allow that such expressions can acquire fully determinate meanings from things other than our stipulations, we can then offer static descriptions of each of the examples that Ludlow uses to support his theory. And even if we do not allow this, then Ludlow will at most be able to cite a few unusual cases – cases that involve explicit semantic stipulation – as examples that demand a dynamic description, and cannot plausibly receive a static

description. Of course, it doesn't follow from this that cases which don't involve explicit semantic stipulation nonetheless demand a dynamic description.

Second horn: the meaning of "meaning" is not underdetermined

So perhaps Ludlow would be better off saying that the meanings of "meaning" and of "underdetermined" and of "flesh out" are *not* underdetermined. Ludlow might say that their meanings are determinate, and that they determinately fix it that the dynamic description of the Mantle/Bonds case is correct. In that case, Ludlow's theory of the dynamic lexicon is not quite as general as we might have thought, and his contextualism not quite so cheap.

But now how did these particular expressions (i.e., "meaning", "underdetermine", "flesh out") manage to acquire determinate meanings, when, according to Ludlow's theory, so many other expressions do not have determinate meanings? What is it about these particular expressions that endowed them with such determinate meanings, in advance of Ludlow's use of them in issuing his dynamic descriptions of various cases? Ludlow doesn't tell us. Of course, *if* Ludlow's theory is true, then we might be able to arrange, stipulatively or otherwise, for these expressions to acquire determinate meanings. But, by Ludlow's own lights, such determinate meanings would be an acquisition, not a birthright. And, for all that Ludlow says, and for all that I'm aware, nothing has ever happened to endow these semantic expressions (i.e., "meaning", "underdetermined", "flesh out") with more determinate meanings than the other expressions in our public lexicon.

Now, in response to what I've just said, I anticipate an objection that was briefly mentioned in my discussion of the first horn of the trilemma: a great deal has happened in the last century or so to endow the semantic expressions in question with more determinate meanings than at least some of the other expressions in our public lexicon. In particular, what's happened is the advent of a science of lexical semantics, a science that employs such terms as "meaning", but sharpens up their meaning for scientific purposes. So how should we understand the term "meaning" in the context of scientific semantics? Perhaps the most popular answer is this: when we speak of the "meaning" of a sentence, we should be understood to mean its *truth-conditions*, and when we speak of the "meaning" of a sub-sentential expression, we should be understood to mean its *systematic contribution to the truth-conditions* of the sentences in which that expression occurs. So, we might assume, what Ludlow means when he talks about an expression's "meaning" is just its contribution to the truth-conditions of the sentences in which it occurs. Doesn't this provide "meaning" with a more determinate meaning than the other expressions of our public lexicon have?

Well, the answer to this question would plausibly be "yes" if the expression "truth-conditions" had a more determinate meaning than the other expressions of our public lexicon have. And the expression "truth-conditions" might plausibly (though not necessarily) have such a determinate meaning if the expressions "truth" or "true" had such a determinate meaning. But do they? Again, I see no evidence that the expressions

“truth” or “true” have any more determinate meaning than the other expressions in our public lexicon. Although a lot of work in semantics and philosophy of language makes *use* of the expressions “true” and “truth”, I do not see that the use that has been made of these expressions suffices to give them a determinate meaning: as far as I can see, it’s just as likely that the various uses to which these expressions have been put place *competing* pressures on their meanings, and that no single meaning for these expressions can enable them to do all the work that philosophers of language have asked them to do.

Just as, in ordinary discourse, we modify the verb “to know” in various ways (someone can know something for sure, or know it with certainty, or know it for all intents and purposes, or know it by the standards of Unger’s seminar), so too we can modify the adjective “true” in various ways (a statement can be true for all intents and purposes, pretty much true, or completely true, or true by the standards of Unger’s seminar). If the applicability of such modifiers to the verb “to know” suggests that the verb is semantically context-sensitive, then does their applicability to the adjective “true” indicate that the adjective is also semantically context-sensitive?

Of course, I am not suggesting that there is anything wrong with two-valued logic, or with the use of “true” to denote a designated value in two-valued logic. But what I am talking about here is the use of the term “true” to denote an empirically detectable semantic property of sentences in a natural language: is there a single such property that is denoted by means of this term, or are there many such properties? What reason is there to insist on the former, as opposed to the latter, possibility?

Recall that our guiding question was this: are the meanings of such expressions as “meaning”, “underdetermined”, and “fleshed out on a case-by-case basis” – are these meanings themselves underdetermined, and to be fleshed out on a case-by-case basis? So far, I have argued that, if Ludlow answers this question in the affirmative, then he is under pressure to admit that it is not clear why it isn’t indeterminate whether the examples that he gives in support of the theory of the dynamic lexicon really do satisfy the description that they need to satisfy in order to support that theory. And I have also argued that, if Ludlow answers this question in the negative, then he is making an exception to the very generalization that he is propounding, and I do not see what grounds he might have for making this exception. (In particular, I don’t know of any developments in the history of lexical semantics that would justify making that exception.) But there is a third way for Ludlow to go.

Third horn: the meaning of “meaning” is sort of underdetermined

Could Ludlow avoid both problems by answering the question “sort of”? More specifically, could Ludlow say that, while it’s not quite true to say that the meanings of these expressions are determinate (so these expressions don’t *quite* count as exceptions to his generalization), it’s also not quite true to say that the meanings of these expressions are underdetermined (so they may be determinate *enough* to fix it that his examples – e.g., the Bonds/Mantle example – satisfy, at least *more or less*, his dynamic descriptions

of them)? Perhaps this is the most promising line for Ludlow to adopt, given his overall argumentative strategy.

Now, if Ludlow adopts this strategy, then he is committed to claiming that his examples (e.g., the Bonds/Mantle example, the example of whether Secretariat is an athlete, the examples involving incompletely stipulated meanings) *more or less* satisfy his dynamic description of them. But are there other examples, of which a dynamic description will be exactly true? If he grabs the third horn of the present trilemma, then Ludlow is not giving us examples of which a dynamic description will be exactly true. Since such examples would be the strongest support for his theory, then his not offering such examples would indicate that he doesn't believe that there are any. That is to say, if Ludlow accepts the third horn of the present trilemma, then the most charitable way to understand why he doesn't offer examples of which a dynamic description is exactly true is that he doesn't think that there are such examples – rather, he thinks that dynamic descriptions are, at best, only more or less true of a case. So, I conclude that if Ludlow accepts the third horn of the present trilemma, then he thinks that dynamic descriptions are, at best, only more or less true of any particular case. Now, if Ludlow's theory of the dynamic lexicon *is* exactly true, then, I take it, there must be at least some examples of which a dynamic description will be exactly true: if the Bonds/Mantle example happens not to be such an example, then there will be others. (A generalization cannot be exactly true, I assume, unless there are some instances of which it is exactly true.) But I've just argued that, if Ludlow accepts the third horn of the present trilemma, then he thinks that dynamic descriptions are, at best, only more or less true of any particular case. It follows that, if Ludlow accepts the third horn of the present trilemma, then he is committed to claiming that his theory of the dynamic lexicon is not exactly true, but only *more or less* true.

So, assuming that Ludlow's theory of the dynamic lexicon is only more or less true, I want to know: is Ludlow's theory of the dynamic lexicon something that can be known to be true? Of course, it is possible for someone to know that p only if p is true. But *how true* must p be in order for someone to know it? Is it good enough for knowing that p, if p is only *sort of* true? In particular, can we have knowledge directed towards the propositional content of Ludlow's theory, when that content is not *quite* true, but only *sort of* true?

If the answer to this question is “no”, then obviously Ludlow's argument for his theory of the dynamic lexicon can only give limited support to the theory itself. His argument cannot give us knowledge of the truth of the theory, because *nothing* can give us knowledge of the truth of the theory: the theory isn't true enough to be the content of anyone's knowledge.

Ludlow might avoid this conclusion by claiming that knowing that p doesn't require that p be more than merely sort of true. Knowing is not *factive*, Ludlow might say, but is only *sort of* factive. Perhaps Ludlow would wish to go so far as to say that “knows” is context-sensitive not simply in the degree of certainty that is required to satisfy the predicate, and in the degree of justification that is required to satisfy the

predicate, but also in the degree of truth that is required to satisfy the predicate. This strikes me as a very interesting and promising view, one which might perhaps be at least somewhat true. The question I want to raise now is: can examples help to establish this general view?

Consider the generalization “blueberries are round”. If there are any generalizations that are only *sort of* true, then this is one of them. Note also that this generalization can be supported by producing examples of blueberries that are only *sort of* round. The same is true for other generalizations that may be only sort of true, for instance “flying airplanes make lots of noise”, “soccer players suffer fewer injuries than hockey players”, “engineers earn more than philosophy professors”. If any generalizations are only sort of true, then these are. And notice that all of these generalizations seem to be established by appeal to examples. Can we also use examples to establish the theory of the dynamic lexicon?

The following consideration suggests that we cannot. When we appeal to examples of *roundish* blueberries to support the generalization “blueberries are round”, we can easily observe that particular blueberries deviate from being precisely round, but they do not all deviate from it to a large extent, and in more or less the same way. (The same point applies to “flying airplanes make lots of noise”, “soccer players suffer fewer injuries than hockey players” and “engineers earn more than philosophy professors”.) Contrast “blueberries are round” with “strawberries are round”. The latter generalization is *so* false that it cannot, in any context that I can imagine, be truthfully said to be known to be true. Strawberries deviate from roundness to a large extent, and in more or less the same way. So the question for Ludlow is: is the theory of the dynamic lexicon true to as great an extent as “blueberries are round” is true? Or is the theory as false as “strawberries are round”? If the former, then perhaps, at least relative to some contexts, we can truthfully be said to know the theory of the dynamic lexicon to be true. But if the latter, then we cannot in any context be truthfully said to know the theory of the dynamic lexicon to be true. So which is it: the former, or the latter? Is the theory of the dynamic lexicon more like “blueberries are round” or more like “strawberries are round”?

I don’t know. And furthermore, I don’t see how examples – Ludlow’s examples or any other examples – are going to tell us the answer to this question. Unlike the roundness of berries, the semantical properties of particular speech acts are not (or at least not easily) observable. What is necessary in order to support Ludlow’s theory of the dynamic lexicon is not merely examples that *can* be given a dynamic description, without straining credulity. What is necessary is rather an independently well-supported theory that predicts when, and why, and to what extent, a dynamic description of particular examples is a correct description. Perhaps such a theory can itself be supported by appeal to examples – but it will not be supported simply by appeal to the fact that such examples can be given a dynamic description.

Recall our analogy above: in order to ascertain that electrons are negatively charged subatomic particles, it does not suffice to point to examples of electrons that are negatively charged subatomic particles: we need to appeal to an independently well-

supported background theory in order to ascertain that the examples to which we are pointing are indeed negatively charged subatomic particles. And while the background theory to which we appeal might itself be supported by particular data points, the data points that support it cannot be merely the examples of electrons that are negatively charged subatomic particles. Analogously, examples by themselves, unsupplemented by such an independently well-supported theory, cannot tell us that the theory of the dynamic lexicon is true enough to be “known” by any reasonable standards. They cannot tell us that the theory of the dynamic lexicon is in the same epistemic boat as “blueberries are round” (and so knowable, at least by low standards), rather than being in the same epistemic boat as “strawberries are round” (and so not knowable, by any reasonable standards).

In sum, if Ludlow grabs the third horn of the trilemma that I’m sketching here, then he will not be able to argue for his theory of the dynamic lexicon simply by appeal to examples. What he needs is not merely examples, but an independently well-supported theory that helps us to see how and why the examples support the theory of the dynamic lexicon. (He might have such a theory in his pocket, but then that is something that needs to be brought into the clear light of day.) Unless we can observe the semantical properties of speech acts in the way that we observe the roundness of berries, examples which only *sort of* satisfy Ludlow’s dynamic description cannot lend very strong support to his theory of the dynamic lexicon.

Conclusion

I conclude that, if Ludlow’s theory of the dynamic lexicon is true, then Ludlow cannot give much support to that theory simply by appealing to alleged examples of lexical dynamics. Would such weak support be good enough for knowledge?

¹ I am grateful to Peter Ludlow, Mark Phelan, and Jonathan Schaffer for helpful discussion.