

The Dynamic Lexicon in a Truth-Conditional Framework; or How to Have Your Cake and Eat It

John Collins

Abstract

A fundamental principle of all truth-conditional approaches to semantics is that the meanings of sentences of natural language can be compositionally specified in terms of truth conditions, where the meanings of the sentences' parts (words/lexical items) are specified in terms of the contribution they make to such conditions their host sentences possess. Thus, meanings of words fit the meanings of sentences at least to the extent that the stability of what a sentence might mean as specified in a theory is in accord with the stability of what a word might mean as similarly specified. In this paper, I shall be concerned with Ludlow's (2014) idea that, in fact, there need be no such sympathy between words and sentences. He proposes that we can square what he calls a *dynamic lexicon*, where word meaning is not stable at all, with a traditional truth-conditional approach of the kind indicated, where sentence meaning is delivered via 'absolute truth conditions'. I share Ludlow's aspiration to accommodate dynamic features of word meaning with a truth conditional approach, but not his belief that the marriage is an easy deal. Thus, I shall present a problem for Ludlow's position and show how resolving this problem leads to an alternative picture of how the meaning of a sentence may be truth-conditionally specified with all relevant dynamic features of the lexicon retained.

Keywords: dynamic lexicon; Ludlow, Peter semantic underdetermination; truth conditional semantics

1: Introduction

Truth-conditional semantics comes in various philosophical and formal flavours. Some of the differences will be pertinent to what follows. What is agreed upon by all, however, is that the

meanings of sentences of natural language (park for the moment whether one's primary concern is sentence types or utterances thereof) is to be compositionally specified in terms of truth conditions for the sentences, where the meanings of the sentences' parts (words/lexical items) are specified in terms of the contribution they make to such conditions their host sentences possess. As intimated, there are many ways of spelling out this basic idea, but what appears to be essential is that the meanings of words fit the meanings of sentences at least to the extent that the stability of what a sentence might mean as specified in a theory is in accord with the stability of what a word might mean as similarly specified. In this paper, I shall be concerned with Ludlow's (2014) idea that, in fact, there need be no such sympathy between words and sentences. He proposes that we can square what he calls a *dynamic lexicon*, where word meaning is not stable at all, with a traditional truth-conditional approach of the kind intimated, where sentence meaning is delivered via 'absolute truth conditions'. I share Ludlow's strategic aspiration to accommodate dynamic features of word meaning with a general truth-conditional approach, but demur when it comes to his tactics. Thus, I shall present a problem for Ludlow's approach and show how resolving this problem leads to an alternative picture of how the meaning of a sentence may be truth-conditionally specified with all relevant dynamic features of the lexicon retained. Thus, I come in peace, and want Ludlow to have his cake and eat it; I just don't want him to eat his cake exactly how he proposes.

2.1: Truth Conditions

For present purposes, let a semantic theory for language L comprise an assignment of (parameterised) truth conditions t to each sentence s of L such that t holding of s is a function of the syntactic composition (at the appropriate level of analysis—LF for the nonce¹) of the lexical items that constitutes s along with appropriate valuation of the open parameters. If such is the desideratum, then the theory has to assign to each lexical item l whatever property of l contributes to the truth conditions of every s that hosts l . Think of this property as simply the

semantic value of *l* allied with the appropriate type category, where we may ignore whether these values are extensions or intensions. So far so familiar.

Note two consequences of this picture. First, if we assume that any lexical ambiguity is syntactically resolved (e.g., *bank* [side of river] and *bank* [financial institution] are distinct items and are treated as such syntactically) and that all context-sensitivity is parameterised, then the semantic values of lexical items will apparently be uniform. They would fail to be so only if their extensions (/intensions) are such as to vary in ways not recorded by the theory, which would thus fail to deliver determinate truth conditions. In other words, the standard picture is that all factors to which changes of extension might be sensitive are already parameterised. Secondly, since the theory is abstracted from speaker-hearers of a language save as they might feature as values for parameters (speakers, hearers, addressees, etc.), the semantic properties of the language will be static, or timeless, as far as the theory is concerned; that is, the theory does not accommodate any shift of extension over time. In sum, a theory on the model on offer presents lexical items as possessing fixed determinate semantic values across speaker-hearers and times.

Some such view as this is perfectly standard in philosophical and semantic circles and is reflected in how lexical items are treated as invariant semantic atoms, whose contribution can be specified disquotationally or via some uniform algebraic relation (again, parameterisation apart).² It remains a separate question, however, whether lexical items are so accommodating; on the dynamic view espoused by Ludlow, they most certainly are not. If that view is right, and I think it is, the pressing question is: how can, if at all, the standard truth-conditional approach accommodate a dynamic lexicon? First, then, I want to commend Ludlow's view.

2.2: The dynamic lexicon

A lexical item counts as dynamic on Ludlow's view because of two features: negotiation of semantic value and underdetermination. The first feature is that the extension of lexical items is always open to negotiation or litigation by speakers of the language. When realised, such negotiation can lead to quite radical results, such as Pluto no longer being counted as a planet (see Ludlow, 2014, pp. 42-51, for a delightful discussion of the ins and outs of the decision to reclassify Pluto). Most often, though, lexical negotiation will involve what is to count as a cake or doll (to use Ludlow's examples), or a sport rather than a pastime, or one style of music as opposed to another, and so on indefinitely. On reflection, one should recognise that such decisions on extensions are ubiquitous; their frequency, however, is not the point. *Any* lexical item (at least the open class ones³) have extensions that are open to negotiation and decision by speaker-hearers. In this sense, the meaning of a lexical item, insofar as it is supposed to determine an extension, is not accommodated by the traditional static picture of the truth-conditional framework. Such a framework tells us what the extension of a lexical item is (modulo context-sensitivity) in one shot, as it were; yet if such extensions are essentially apt to be negotiated, then no theory could specify lexical extensions once and for all, and so couldn't do so for truth conditions.

A similar conclusion is reached upon consideration of underdetermination, the second feature Ludlow takes to characterise lexical items. The underdetermination thesis is that the extension of a lexical item is underdetermined by the standing meaning of the item cognised by speakers of the relevant language, i.e., different, incompatible extensions can be taken by a lexical item on different occasions of its tokens, consistent with a single standing meaning (the item is not ambiguous). So, even if no negotiation or decision is in the offing *vis-à-vis* some given lexical item, its extension is not statically determined by the language alone, for diverse extensions are available compatible with the understanding of the item of any given interlocutors.⁴ Ludlow leads with the example of *book*. In answer to the question, 'How many

books have you written?', different criteria may apply that variously include or exclude co-authored books, edited collections, Ph.D. theses, unpublished book-length manuscripts, on-line copy, etc. In the absence of some further specification as to what is to count as a book in the present context, there is no right answer as to what falls under the extension of *book*. Analogous reasoning can be readily concocted for much about any open class item one may consider (cp., Asher, 2011)

The import for the traditional static conception is much the same as detailed above for the negotiated aspect of lexical extensions. Absent some specification of what is to count as falling into the extension of the given lexical item in the present context of utterance, one cannot say that the item has any determinate extension, and so it follows that sentences lack determinate truth conditions. The problem for the traditional conception, of course, is simply that it does not accommodate any further specification of extensions beyond what can be parameterised, but the kind of underdetermination exhibited does not fall under a contextual parameter, at least not the kind anyone has ever proposed.⁵

My purpose in this paper is not to defend a dynamic conception of the lexicon, so I shall not here address the myriad responses that might be imagined on behalf of the traditional conception. My intent is to see if the dynamical conception is, after all, compatible with a truth-conditional framework, albeit, obviously, shed of its static pretence. Still, it is worthwhile to pause on a general point Ludlow rightly emphasises.

Granting that lexical items are open to negotiation and underdetermined certainly shows that aspects of meaning (whatever determines extension) are multiple and non-static. So much, though, does not entail the absence of a basic or primitive meaning, which fixes a canonical extension. The openness there is to lexical items, therefore, might pertain to the availability of secondary meanings without affecting the putative basic meaning. Ludlow

(2014, pp. 81-3) endorses *semantic egalitarianism*, which is the simple denial of this suggestion: of a range of potential extensions for an item, there is not a basic or canonical one. I endorse semantic egalitarianism, too, but Ludlow says little in defence of the thesis; so, a brief discussion is appropriate.

Some form of semantic privilege is very tempting, but I'd wager that its intuitive pull is due to a conflation of privilege with two perfectly innocent and plausible theses that are not in tension with a dynamic conception of the lexicon.

It is undeniable that the meanings of lexical items have a history, as it were, where an original meaning is retained or blended with novel construals. Take Ludlow's own example of *book*. Transparently, any notion of an e-book or an on-line manuscript is historically derivative of a hard copy (published or not). For any given pair of construals, one might be able to tell some such history of dependence. Let's assume one could. The thought, encouraged, therefore, is that there is an asymmetric dependence between construals, and that the dependent construal is non-basic, open to the kind of change and triggering the dynamic view seeks to generalise erroneously.

This thought is fundamentally mistaken, even granting that one could tell some such asymmetric dependence story across the board. The problem is simply that the relevant notion of lexical meaning must be local to the speaker-hearers whose language is our topic of inquiry. Consider again the case of *book* just presented. The notion of an e-book is historically grounded in the notion of a hard copy, but why should that matter to an all-too-modern speaker-hearer more *au fait* with computers than libraries? In essence, the hypothesis tabled is merely a sophisticated version of the etymological fallacy, whereby the history of a word is taken to be constitutive of its content. So, sure, there are lots of historical asymmetric dependencies, such is just what it is for something to be historical. What is required to spike the guns of the dynamic

view is a synchronous asymmetric dependence. Maybe a story could be told that makes such relations both pervasive and plausible. Still, it is wholly unclear how such a story would go without falling back into etymology.

The second source, I think, of the intuitive appeal of the idea of a unique basic meaning is a conflation of invariance with privilege. Diverse work in lexical semantics and syntactic theory has given great plausibility to the idea that lexical items possess a core invariant meaning. Ludlow (2014, p. 99) is happily swayed by this plausibility.⁶ What such work entails, however, is that *not every* aspect of meaning is dynamic in the relevant sense; that is, not every aspect of lexical meaning is open to negotiation/litigation or is underdetermined. It is crucial, though, to be clear on the structural nature of the kind of invariances plausibly uncovered by linguistic theory. Mostly, they pertain to argument structure—how many arguments, what kind of arguments, and what grammatical relations arguments may take to the syntactic projection (e.g., alternation options). These invariances do not fix a unique extension for any lexical item, but rather constrain possible extensions. Indeed, Grimshaw (2005) has claimed that, from this perspective, all lexical items belonging to a class of the same argument properties amount to linguistic synonyms. As Ludlow notes, this conception of invariance is perfectly compatible with a dynamic view of the lexicon, which is not intended to be a species of Humpty-Dumptyism. The dynamic view, at least as propounded by Ludlow, does not insist on such invariances, but if they do obtain, then they serve as constraints on the diversity of extensions a lexical item might express. There might be richer invariances of the kind that underwrite various analyticities, sometimes captured by meaning postulates. Again, however, it is unclear why any such invariances would support a privileged extension. Suppose, for example, that it is analytic that books contain words, can be read, have authors, and so on. Such relations hold across the diversity of extensions *book* may take.

Hereafter, then, let's assume that lexical items are as Ludlow depicts them to be. Given the apparent dissonance of such a conception with the traditional static truth-conditional framework, a natural move would be to seek to dispense with the received view in favour of some other conception. Ludlow pursues a more nuanced path, which will be in focus for the remainder of the piece.

3: Squaring a dynamic lexicon with a truth-conditional framework

My talk about a truth-conditional framework has failed to disambiguate what for some is a crucial bifurcation, *viz.*, whether the framework appeals to *truth in a model* or *absolute truth*. The difference can be sourced to Tarski's different approaches to the semantics of formal languages (cp., Milne, 1999; Etchemendy, 1990). The model-theoretic approach effectively treats language as an uninterpreted symbol system. A model consists of (i) a set of entities of the relevant kinds and (ii) an interpretation function that maps such entities onto the elements of the language. Further functions are called upon to interpret the effect of various permissible compositions within the language. Truth conditions for the sentences of the language are thus definable only relative to a model that assigns various extensions to the linguistic expressions. A change of model, for example, may preserve the distribution of truth and falsehood over the language's sentences, but issue in radically distinct truth conditions (thus, logical notions are definable in terms of conditions holding across all models).

A framework that trades in absolute truth is one that takes the language to have a fixed interpretation, which the theory seeks to capture (just how such facts are fixed is a matter that can, in principle, be settled in various ways consistent with the theory). Thus, the sentences of the language have fixed truth conditions (assuming some process of disambiguation), and so the concept of truth is understood to be invariant or absolute, for there is just the one

interpretation at issue. As with the question of what fixes the facts of interpretation, what makes a sentence true can be settled in various ways consistent with the general framework.

There is a small literature arguing that an absolute conception of truth is the only one suitable for semantic theory for a natural as opposed to formal language (Lepore, 1982; Higginbotham, 1988, 1993). Glanzberg (2014), however, is right to be sceptical of the supposed Homeric struggle, as it were, insofar as, if natural as opposed to formal language is at issue, then the so-called model-theoretic approach simply collapses into the absolute conception, for a specific intended interpretation is always presupposed. Only if one is seeking to characterise some notion of logical truth or consequence might different models be relevant. Such a putative collapse goes in the direction of the absolute conception. Still, a feature of the model-theoretic approach is commonly assumed with or without explicit commitment to the conception. As mentioned, a model consists of a universe of objects, some means of composing such objects, and a function that assigns objects to linguistic material. Most familiarly, we take the entities of the model to be ur-elements of a set-theoretic universe, and sets to be further entities, or, equivalently, to be functional specifications of such entities (characteristic functions). The important point for our purposes is that such entities have precise individuation conditions—they are static, fully determined. Such a view naturally goes over to the absolute conception, for if truth conditions on such a view are ways the world might be such that a sentence might be true or false, then they too will presumably involve fully determinate entities insofar as the world as it is anyway is not subject to our negotiation or construal, only what we say by an utterance might be. At first pass, then, the distinction between a model-theoretic and an absolute conception of the truth-conditional framework seems irrelevant to our purposes: on either conception, a dynamic conception of the lexicon is incompatible with the basic truth-conditional approach. Not so, suggests Ludlow.

The matter turns on the basic theory form in which an absolute conception is couched. A model-theoretic form assigns linguistic expressions entities from the model. Very roughly, for example:

$$(1)a \text{ } [[\mathbf{Kansas}]]^j = \mathbf{Kansas}'^7$$

$$b \text{ } [[\mathbf{flat}]]^j = \lambda x. \mathbf{flat}(\mathbf{x})'$$

Statements such as these say that the interpretation function j assigns to the lexical items some function designated in the metalanguage via an expression with the same extension as the relevant object-language expressions. For the sentence, after some manipulation, we have:

$$(2) \text{ } [[\mathbf{Kansas is flat}]]^j = \mathbf{T} \leftrightarrow \{x: x = \mathbf{Kansas}'\} \subseteq \{x: \mathbf{flat}(\mathbf{x})'\}$$

The meanings of the lexical items are given via model-theoretic objects, such as functions and sets, and truth conditions are similarly given via operations on sets. The language of the theory (the metalanguage) thus specifies what sentences of the target language mean via a function that effectively maps abstract entities onto the linguistic items of the language. Insofar, then, as we take sets, functions and other such denizens of models to be precisely individuated, fully determinate entities, a marriage of a model-theoretic semantics and a dynamic lexicon is out. Why should we think a different conclusion is on the cards once the dynamic view is paired up with an absolute conception? The answer, for Ludlow, rests on the theory form appropriate for the absolute conception. Ludlow presents a standard form along the following lines:

$$(3)a \text{ } \forall [x, \mathbf{Kansas}] \leftrightarrow x = \mathbf{Kansas}$$

$$b \text{ } \forall [x, \mathbf{flat}] \leftrightarrow x \text{ is flat}$$

By a relevant compositional principle and rules of substitution, we have the following example a theorem (see Ludlow, 2014, pp. 107-8):

(4) $\forall[T, \textit{Kansas is flat}] \leftrightarrow \textit{Kansas is flat}$

The crucial point here is that the theory does not assign entities to lexical items and sentences, but states conditions under which the items/sentences have certain semantic properties by way of employing the very language at issue in the metalanguage of the theory (it wouldn't matter, of course, if German was used). So, the VP *is flat* isn't assigned the set of flat things as its extension, but we are told that to be the semantic value of the VP is equivalent to being a flat thing. In this sense, the metalanguage simply inherits the conceptual content of the object-language, and specifies semantic properties in such terms without introducing its own generic ontology of sets or functions. As Ludlow (2014, p. 108) puts it:

But now notice that we did all this without introducing the usual machinery (utilizing properties, sets, functions) that is supposed to be precise (and determinate). What is interesting about our axiom for 'flat' is that it takes advantage of the underdetermined meaning of 'flat' in the metalanguage... I am not making a claim against contemporary formal semantics. There is no question but that formal semantics has been extremely fruitful and that it has provided many profound discoveries and insights into the nature of language... None of these discoveries, however, rests upon the assumption of a precise or determinate metalanguage... The point is that for the most part all the actual results of semantics are preserved if we reject the idea of precision in the metalanguage.

The difference between the two approaches, therefore, might be expressed in terms of the different weight the metalanguage of the respective theories carries. On the absolute conception, the metalanguage simply inherits its concepts from the object-language, as well as providing various logical and semantic machinery. The relevant expressions of the metalanguage are specified as referring to whatever the corresponding object-language expressions refer to. The model-theoretic approach, by contrast, introduces novel ontological categories, such as sets and functions, to serve as extensions for the object-language

expressions. The relevant expressions of the metalanguage, therefore, can be said to refer to what the corresponding object-language expressions refer to only given a weighty theoretical claim to the effect that the metalanguage actually provides the right ontology for the object-language rather than simply specifying what the object-language expressions mean. To this extent, the absolute conception *appears* to be fundamentally less committing than the model-theoretic view, for the former takes no position on how to unpack the ontological significance of an expression's contribution to the truth conditions of its host sentences. I hedge this claim for the matter will turn on how the two conceptions are to be best construed. At the level of ideology, at least, the absolute conception is neutral in a way the model-theoretic conception is not, and it is precisely this difference Ludlow seeks to exploit.

Ludlow's basic idea is that since the metalanguage inherits its concepts from the object-language—at any rate, the concepts employed to specify the extensions of the object-language expressions—the metalanguage will inherit any properties of the object-language. In particular, therefore, the dynamic properties of the object-language lexicon will be inherited by the metalanguage lexicon. If a semantic theory along the lines decreed by the absolute conception meets whatever adequacy criteria are appropriate, then it will continue to do so on the assumption that the lexicon of the language in question is dynamic, for the character of the lexicon is irrelevant to the theory, insofar as it simply passes up whatever properties the lexical items possesses from the object- to the metalanguage. This, of course, is just the policy Davidson (1984, p. 31) commends: 'the mystery is transferred from the word [...] in the object language to its translation in the metalanguage'. Neither Davidson nor Ludlow think that there really is any mystery hereabouts, only a phenomenon that an adequate theory does not explain.

At first blush, this line of reasoning might seem too quick by half. True enough, it might be thought, the dynamic properties can be inherited from the object- to the metalanguage, but that postpones the problem rather than resolves it. The problem is how to square the essential

dynamic features of the lexicon with the compositional stability of truth conditions. No such concern arises, beyond the usual cases, for a marriage of a static lexicon with static truth conditions, but the problem appears insuperable, if one seeks to marry the static and the dynamic. Of course, one might think that a dynamic lexicon ineluctably leads to dynamic truth conditions, but if so, it is unclear how the standard truth-conditional format Ludlow favours is supposed to specify such conditions. A world that contains books on one construal of *book* (an e-book) need not be a world that contains books on another construal of *book* (concrete particulars). On the first construal, in some given world, *The book is black* would be false or nonsense, whereas on the second construal it might well be true, but the bare specification of truth-conditions Ludlow commends does not differentiate between the cases, for '*The book is black* is true iff the book is black' will be the result on either construal. Ludlow has an answer to this kind of quandary, albeit one motivated by a slightly different concern pertaining to bivalence, to which I shall return in §4 (see Ludlow, 2014, pp. 111-2).

Ludlow does not think of a truth-conditional theory as targeting a static language; instead, a theory is effectively about the kind of conversational exchanges between speaker-hearers in which dynamic features of the lexicon are realised. To this extent, the theory does trade in dynamic truth conditions, but not because it explains or otherwise encodes the dynamic features of the lexical items, but because the theory takes sentences as used in dynamic conversations as its object, not sentences as abstract types. Thus:

I take the semantic theory to be a theory which computes the semantic values of utterances (or, if you prefer, tokens)—not sentences construed as abstract objects (this is a distinctively anti-Kaplanian assumption which I won't defend here) (Ludlow, 2014, p. 112).

The crucial move now is to restrict the object-languages of utterances amenable to a truth-conditional approach to be those for which the question of truth can be settled. Ludlow (*op cit.*) offers the following principle:

Microlanguage S-admissibility: No utterance *u* of a sentence *S* is admissible in a microlanguage *L*, unless discourse participants (tacitly) agree that the terms of *S* are modulated so that an utterance of the sentence will be determinably either true or false.

So, sentence types will be indeterminate and underdetermined in various respects, but as any such type is tokened, interlocutors must at least recognise an assertion being made, and so treat the token as possessing a truth value under conditions they may tacitly converge upon. The position has obvious sympathies with Davidson's (1986) conception of a 'passing theory', which Ludlow (2014, pp. 113-4) readily acknowledges. Davidson, however, notoriously concludes that 'there is no such thing as a language' as an invariant structure or a stable body of information. Such a conception goes well beyond what Ludlow at least recognises as the proper remit of a semantic theory, for there might well remain all manner of semantically significant invariances even after one has factored out what is peculiar to filling out the conditions involved in Ludlow's principle of admissibility. Prudentially, Ludlow (*ibid.*, p. 114) observes:

[T]here must be stable elements to the construction of a passing theory; otherwise it is difficult to see how it could be possible. It is a safe bet that most of the nonterminal semantic rules are stable across these shifts (they may well be stable across all human languages). Thus the real dynamic portion would be in constructing the terminal (lexical) rules on the fly.

In essence, the suggestion here is that dynamic factors affect the lexicon alone, leaving invariant syntax and any other purely structural conditions on meaning a semantic theory properly targets.

4: A friendly alternative

There is so much that is accurate and insightful in both the detail and general arrangement of pieces in Ludlow's position that my objections skirt churlishness. My aim, however, is not to nit-pick, but to show how all that is crucial to Ludlow's general position can be preserved while dropping the notion that the object of a semantic theory is a microlanguage; indeed, I shall suggest that what motivates the introduction of the microlanguage notion is a construal of semantic theory that is too wedded to the very static conception Ludlow renounces, i.e., an undue concern for bivalence as a property of a theory of sentence types.

First, let's consider the question of what a semantic theory is about. The obvious answer is 'language', of course, but that leaves everything more or less open, given that language can be broadly or narrowly construed. Let's take our lead from generative linguistics, however, and understand language narrowly as a lexicon plus syntactic operations defined over it. On this conception, a semantic theory being about language would amount to it being about the set of conditions syntax and lexical items impose upon meaning, which pretty much amounts to Ludlow's observation quoted at the end of the last section. But what is meaning? Let's reformulate slightly and say that a semantic theory is about the conditions language alone (lexicon plus syntax) imposes upon what one can literally say. The basic linguistic phenomena, therefore, are of the following form:

(BLP) A can use S to say P in C

A semantic theory is a theory of the values of ‘*S*’ as conditions upon the values of ‘*P*’. Such a characterisation fails to decide if the conditions *S* imposes upon *P* are often or even ever determining of *P*. Consider the following case:

(5) Bill borrowed a red pen

The properties of *S* alone—the type *Bill borrowed a red pen*—will place fixed, non-negotiable conditions upon what one can say with (5), such as the following: (i) Bill is the recipient of a borrowing; (ii) there is an unspecified lender; (iii) the object transferred is a pen; and (iv) the pen counts as red for the speaker. This is quite a bit, but other factors are left underdetermined, such as what it is for the pen to count as red: a pen with a wholly red surface or partially red surface (does the interior of the pen count?); a pen that writes red, but could have a surface of any colour; a pen with a red plastic cap; a pen classified as red (belonging in the red tray); a pen that merely looks red in certain lighting; and so on. Examples such as this can be multiplied indefinitely (cp., Lahav, 1989; Travis, 1997). The general moral is that narrow linguistic properties will fix conditions on possible propositions expressible via sentence tokens, but other crucial ingredients will be left unfixed by the language alone, to be determined on an occasion of use by extra-linguistic conditions.⁸

Ludlow’s realm of the dynamic falls squarely within the extra-linguistic as I am here conceiving of it. A semantic theory targets the type properties of linguistic material as invariant conditions on what can be literally said with such material. Dynamic features are, *eo ipso*, not invariances of content, but matters of construal selected by speakers on the occasion of use, or at least features highly sensitive to variant conditions of use and open to negotiation. Since such dynamic aspects of the lexicon are not part of the language narrowly conceived, they are not features a semantic theory will specify or otherwise explain. Content, on this view, as what can literally be said with a sentence, will be left undetermined by fixing the semantic properties

of the language. The theory still properly counts as truth-conditional, however, because it specifies constraints upon what can be said, and so upon truth conditions. It is thus perfectly proper and inevitable that the theory will specify semantic properties in terms of truth conditions, the conditions will simply be less than determining.

It is in this sense that Ludlow's idea, after Davidson, of a semantic theory inheriting the lexical features of the object-language is correct. It is just another way of saying that the theory fails to specify the range of construals the lexical item of the object-language makes available, but does constrain whatever values the item does take to fall within whatever the item may contribute to what is said. So, what's the beef?

In spelling out the broad outlines of the object and purpose of a semantic theory, dynamic features were accommodated, but no appeal to microlanguages or concomitant principles were required. Microlanguages appear to be redundant, a needless posit. A semantic theory remains perfectly general, applying to linguistic types and specifying the conditions they invariantly impose upon what can be said with their tokens. What is peculiar to the level of the microlanguage, as it were, is something on which a semantic theory is silent, precisely because it is essentially extra-linguistic and so goes beyond what the relevant linguistic types can determine. Of course, the microlanguage phenomena are not inconsistent with the right semantic theory; the phenomena peculiar to particular discourse situations simply fill out the properties of what is said, and so what can be true, which the semantic properties of the type only partially determine.

If the above is correct, then why should Ludlow or anyone else find appeal in microlanguages? As far as I can see, the only motivation is due to a concern over bivalence Ludlow (2014, pp. 111-2).⁹ Let's suppose that we presume every sentence that meets the relevant conditions (declarative, non-vague, etc.) to be either true or false (parameterised for

context). Dynamic factors appear to give the lie to bivalence, for (6) is indeterminate in truth value, regardless of how the world actually is, until some speaker has settled on how the adjective is to be construed in the relevant context:

(6) Kansas is flat

Yet, a semantic theory is also supposed to tell us that

(7) 'Kansas is flat' is true iff Kansas is flat

It would appear, therefore, that since the right-hand side of (7) is indeterminate, the left-hand side will be, too. Yet that is just to give up on bivalence, for now it is indeterminate whether a sentence is true. So, all instances of the T-schema will now be given up for any target sentence that is indeterminate. One response to this quandary, the one Ludlow entertains, is to uphold bivalence, but take a theory's target to be microlanguages, where a level of 'passing' determinacy is settled on to allow for utterances to be counted as true or false. This solution, I think, throws the baby out with the bathwater, and is not required either.

A semantic theory is not supposed to be a theory of interpretation and so not a theory of what can be settled in a discourse setting. It is supposed to be a theory of the semantic properties of the language alone. The distinction is not mere whimsy, but reflects the facts of general structural conditions invariantly holding over discourses where variously different propositions are produced and consumed. In this light, an adequate theory precisely must remain general, and so above the contingencies of microlanguages, if it is explicitly to record those semantic properties fixed at the level of the type. A theory targeting microlanguages could *implicitly* record whatever is constant across microlanguages, but then the situation would simply call for an explicit account of such features that each speaker-hearer brings to the discourse, and such a theory is precisely the one we have already that targets sentence types. So, the microlanguage

posit appears to lose exactly what we want to capture, which can, to boot, be fully caught by eschewing microlanguages in the first place and keeping with sentence types.

The problem of bivalence remains, or so it appears. Let's suppose we wish to retain bivalence. An amendment we need to make given the linguistic facts is that sentence types are neither true nor false, but that is no loss, for they are simply not the kind of things that can possess truth value. Ludlow is right: what has truth value are the utterances with their peculiar construals that may take place in a discourse setting. What a theory of truth conditions records is not the precise conditions under which any utterance would be true, but schematic or partial information of the kind that constrains rather than determines what is sayable with a sentence of the type. So, in specifying truth conditions for a sentence type, general conditions are given for a sentence type free of content-determining factors. Both sides, as it were, are thus reciprocally underspecified such that, say, no speaker ever intends to speak of a red pen as such, for there can be no such thing, but, for all that, linguistic properties of *red pen* constraint what he can say. Likewise, the world cannot contain red pens as such, even though the various ways for a pen to count as red all count as instances of red pens. Thus, to swap examples, (7) can be perfectly true, even if (6) as a type is indeterminate. The point is that both sides share a lack of specificity but count as equivalences, if construed jointly. Thus, one might render (7) as

(8) 'Kansas is flat' is true, on some construal, iff, construed the same way, Kansas is flat

Of course, such a paraphrase is not part of the theoretical machinery, and that is no loss. The paraphrase is simply a way of marking that the theory specifies general structural information that constrains what can be said with sentences without determining any specific proposition, and so the truth conditions are real, albeit schematic. The situation, of course, is perfectly familiar from the study of formal uninterpreted languages. We can say that this or that formula

is true, but mean to talk of instances of the formula. Just so here, save that what is to count as an instance is not a matter that can be nailed down model-theoretically, but remains a matter to discover given the various construals the lexicon makes available or which can be negotiated. Still, bivalence is retained.

5: Conclusion

I commend my alternative proposal to Ludlow and anyone else who admires the truth-conditional approach to semantics while being equally struck by the irresolvable underdetermination of linguistic content. Ludlow has identified the right problem and offered a nice solution. I think a better solution is available that costs less and does a better job. Still, in philosophy, asking the right questions is the important thing. All the answers have a tendency to be kinda wrong. I endeavour to be less wrong than Ludlow.¹⁰

Notes

¹ Think of LF as a syntactic level that realises certain structural or logical features of meaning, such as argument structure, scope assignment, variable assignment, and so on. See May (1977, 1985) and Chomsky (1981) for seminal discussion. It doesn't matter to the following whether LF is viewed as an independently constituted syntactic level or as simply a cipher for whatever syntactic properties fix the relevant features of semantics.

² The basic thought here is that a semantic theory does not indulge in general conceptual analysis (bracket the issue of meaning postulates), but simply maps metalanguage expressions onto object-language ones with an understood preservation of extension. Of course, a semantic theory should and does do more insofar as it captures algebraic aspects of argument structure, functional properties (tense, modality, etc.), and relations of scope and quantification. Still, every proposal in the truth-conditional tradition has what we might think of as a primitive

component that simply maps unitary concepts onto unitary lexical items. See Glanzberg (2014) for an insightful discussion of this issue.

³ Ludlow dwells on kind nominals (artefact and natural), but examples can be readily developed for verbs and adjectives, and all kinds of other nominals. Closed class items appear insusceptible, presumably because they express functional/structural relations, which fall outside of negotiations over what is to count as an instance. See note 2.

⁴ For broad discussion of underdetermination, see Recanati (2004) and Carston (2002). The basic insight, however, is long-standing; see, for example, Austin (1962) and Ziff (1972).

⁵ Again, see Asher (2011).

⁶ Perhaps too swayed. There is a growing literature, initiated by Borer (2005a, 2005b), that rejects lexical complexity in favour of lexical roots as unstructured atoms. My present point, however, simply presupposes that all parties accept complex lexical items.

⁷ Here, I treat a name as if it simply picks out an object. An standard Montagovian treatment takes a name n to pick out a set of properties possessed by some entity ($\lambda P.P(n)$), such that predicating a property R of n is true iff R is one of the properties P possessed by n . There are other ways of proceeding.

⁸ Lahav, for one, pitches underdetermination as a problem for compositionality. Such a putative quandary has attracted a number of solutions via the semantic accommodation of context (cp., Reimer, 2002, and Szabó, 2001). Similarly, Travis has been taken to raise issues for compositionality and been dealt with in a similar manner (cp., Predelli, 2004; Kennedy and McNally, 2010; Lasersohn, 2012). Regardless of how successful or otherwise justified such responses might be, my current appeal to the style of reasoning Lahav and Travis popularised pertains just to the identification of underdetermination phenomena. For what it is worth, I think it is clear that some notion of compositionality remains unmolested by underdetermination, which only putatively shows that type-level meaning doesn't fix a context-

invariant proposition. Independent work must show whether or not the relevant contextual factors are genuinely extra-linguistic or somehow under a linguistic license. For present purposes, I am simply assuming the latter.

⁹ Ludlow thanks David Braun for raising the issue. Lepore and Ludwig (2005, pp. 138-42) address essentially the same concern in review of Davidson's 'mystery'-transferring stratagem whereby lexical properties of the object-language are simply inherited by the metalanguage.

¹⁰ My thanks go to Herman Cappelen for the kind invitation to contribute to this symposium and to Peter Ludlow, for various discussions of these issues over the years.

References

- Asher, Nicholas. 2011. *Lexical Meaning in Context: A Web of Words*. Cambridge: Cambridge University Press.
- Austin, John L. 1962. *How to Do Things with Words*. Oxford: Oxford University Press.
- Borer, Hagit. 2005a. *Structuring Sense, Volume I: In Name Only*. Oxford: Oxford University Press.
- Borer, Hagit. 2005b. *Structuring Sense, Volume II: The Normal Course of Events*. Oxford: Oxford University Press.
- Carston, Robyn. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Chomsky, Noam. 1981. *Lectures on Government and Binding: The Pisa Lectures*. Dordrecht: Foris.
- Davidson, Donald. 1984. *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.
- Davidson, Donald. 1986. "A Nice Derangement of Epitaphs." In *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, edited by E. Lepore, 433-446. Oxford: Blackwell.

- Etchemendy, John. 1990. *The Concept of Logical Consequence*. Cambridge, MA: Harvard University Press.
- Glanzberg, Michael. 2014. "Explanation and Partiality in Semantic Theory." In *Metasemantics: New Essays on the Foundations of Meaning*, edited by A. Burgess and B. Sherman, 259-292. Oxford: Oxford University Press.
- Higginbotham, James. 1988. "Contexts, Models, and Meanings: a Note on the Data of Semantics." In *Mental representations: The Interface of Language and Reality*, edited by R. Kempson, 29-48. Cambridge: Cambridge University Press.
- Higginbotham James. 1993. "Grammatical Form and Logical Form." In *Philosophical Perspectives, 7: Language and Logic*, edited by J. E. Tomberlin, 173-196. Atascadero, CA: Ridgeview Publishing.
- Kennedy, Chris. and McNally, Louise. 2010. "Color, Context, and Compositionality." *Synthese* 174: 79-98.
- Lahav, Ran. 1989. "Against Compositionality: the Case of Adjectives." *Philosophical Studies* 57: 261-279.
- Lasnik, Peter. 2012. "Contextualism and Compositionality." *Linguistics and Philosophy* 35: 171-189.
- Lepore, Ernie. 1982. "What Model Theoretic Semantics Cannot Do." *Synthese* 54: 167-187.
- Lepore, Ernie, and Ludwig, Kirk. 2005. *Donald Davidson: Meaning, Truth, Language, and Reality*. Oxford: Oxford University Press.
- Ludlow, Peter. 2014. *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. Oxford: Oxford University Press.
- May, Robert. 1977. *The Grammar of Quantification*. Ph.D. dissertation, MIT.

- May, Robert. 1985. *Logical Form: Its Structure and Derivation*. Cambridge, MA: MIT Press.
- Milne, Peter. 1999. "Tarski, Truth and Model Theory." *Proceedings of the Aristotelian Society* 99: 141-167.
- Predelli, Stefano. 2004. "Painted Leaves, Context, and Semantic Analysis." *Linguistics and Philosophy* 28: 351 - 374
- Recanati, Francois. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.
- Reimer, Marga. 2002. "Do Adjectives Conform to Compositionality?" *Philosophical Perspectives* 16: 183-198.
- Szabó, Zoltan Gendler. 2001. "Adjectives in Context." In *Perspectives on Semantics, Pragmatics and Discourse: A Festschrift for Ferenc Kiefer*, edited by I. Kenesei and R. M. Harnish, 119-146. Amsterdam: John Benjamins.
- Travis, Charles. 1997. "Pragmatics." In *A Companion to the Philosophy of Language*, edited by B. Hale & C. Wright, 87–106. Oxford: Blackwell.
- Ziff, Paul. 1972. *Understanding Understanding*. Ithaca: Cornell University Press.