Routledge
Taylor & Francis Group

Inquiry
An Interdisciplinary Journal of Philosophy

ISSN: 0020-174X (Print) 1502-3923 (Online) Journal homepage: https://www.tandfonline.com/loi/sinq20

Idealisation in semantics: truth-conditional semantics for radical contextualists

Gabe Dupre

To cite this article: Gabe Dupre (2020): Idealisation in semantics: truth-conditional semantics for radical contextualists, Inquiry, DOI: 10.1080/0020174X.2020.1712225

To link to this article: https://doi.org/10.1080/0020174X.2020.1712225

INQUIRY
https://doi.org/10.1080/0020174X.2020.1712225

Taylor & Francis
Taylor & Francis Group

Idealisation in semantics: truth-conditional semantics for radical contextualists

Gabe Dupre

Department of Philosophy, University of California, Los Angeles, CA, USA

ABSTRACT

In this paper, I shall provide a novel response to the argument from context-sensitivity against truth-conditional semantics. It is often argued that the contextual influences on truth-conditions outstrip the resources of standard truth-conditional accounts, and so truth-conditional semantics rests on a mistake. The argument assumes that truth-conditional semantics is legitimate if and only if natural language sentences have truth-conditions. I shall argue that this assumption is mistaken. Truth-conditional analyses should be viewed as idealised approximations of the complexities of natural language meaning. From this perspective, disparity between the scientific model and its real-world target is to be expected. I elaborate on what such an approach to semantics would look like.

ARTICLE HISTORY Received 10 March 2019; Accepted 10 December 2019

KEYWORDS Philosophy of language; philosophy of linguistics; natural language semantics; contextualism; minimalism; scientific models

1. Introduction

The received view of semantics centres around the claim that the meaning of an assertoric sentence specifies a mapping from an utterance context to the conditions under which such an utterance would be true. Formal semantics has largely consisted in the production of formal systems wherein the truth-conditions of formal analogues of natural language sentences can be compositionally derived. However, theorists of many different stripes have argued that this approach rests on a mistake.

Sceptics of truth-conditional semantics[1] have argued, using a wide range of examples, that sentences fail to provide such a mapping. While

CONTACT Gabe Dupre gdupre@humnet.ucla.edu

[1]Such as generative linguists (Chomsky 2000; Hornstein 1984, 1989; McGilvray 1998; Pietroski 2005, 2006, 2018), relevance theorists (Carston 2013; Sperber and Wilson 1995), contextualists of certain stripes (Bach 1994; Recanati 2010; Searle 1980; Travis 1985, 1996, 1997), ordinary language philosophers (Austin 1950; Strawson 1950) and others who are more difficult to categorise (e.g. Rayo 2013 or Wilson 2008), to mention just a few.

a sentence may identify some contextual features which must be identified in order to determine a truth-condition, the influence of context on truth-conditions outstrips this.

Standard defences of truth-conditional semantics, such as moderate contextualism[2] and semantic minimalism[3], aim to close this gap between meaning and truth-conditions. The moderate contextualist aims to match the context-sensitivity of natural language with context-sensitivity in its formal analyses. The minimalist denies that the natural language expressions are as context-sensitive as they seem, relocating this context-sensitivity into pragmatics.

Both proponents of this argument and defenders of truth-conditional semantics share an assumption: truth-conditional semantics is legitimate only if the semantic properties of natural language sentences are in general sufficient to determine a truth-condition in a context. Proponents of the argument claim that they are not, and so truth-conditional semantics is mistaken. Defenders of this approach use various means to show that, contrary to appearances, they are, and so truth-conditional semantics is saved. In this paper, I shall argue that this assumption is mistaken. By viewing the formal systems of truth-conditional semantics as *idealised* versions of their natural language targets, I show how truth-conditional semantics can be legitimate even though the semantic values it posits do not perfectly reflect the meanings of the natural language expressions it targets. While the semantic properties of natural languages are indeed subtly context-sensitive, by abstracting away from this messiness we can produce approximations in the form of truth-conditional models. These approximately true models can then enable us to better predict and explain a variety of linguistic phenomena. In particular, I shall argue that high-level generalisations can be captured by more minimalist models, but that these generalisations will have exceptions, and that more contextualist models enable us to predict and explain these exceptions.

This approach serves to align the methodology of linguistics more closely with that of the special sciences. As we shall see, context-sensitivity places opposing stresses on formal approaches to semantics. On the one hand, if the effects of context on meaning are genuinely as widespread as radical contextualists have argued, one might feel that an adequate semantic theory must thereby also incorporate massive amounts of complexity in order to be empirically adequate. On the other, given the extent

[2]E.g. Stanley (2007).
[3]E.g. Borg (2004) and Cappelen and Lepore (2008).

of such effects, one might be tempted to exclude context-sensitivity almost entirely, with the aim of providing simple and general pictures of the determinants of meaning common to a wide range of contexts, leaving the distorting properties of context out of the inquiry. Much of the debate in the philosophical literature has aimed at defending one of these two extreme positions. However, this is rarely the approach taken in other fields which aim to understand the behaviour of complex systems. Instead, the standard approach involves producing a wide range of overlapping but distinct representations of the target, each aiming to capture some aspects, while leaving out others.

2. The argument

The argument against truth-conditional semantics challenges the assumption that sentences provide the requisite information to determine truth-conditions. If the goal of semantic theory is to isolate the stable meaning properties of natural language expressions, leaving the ways in which sentences interact with the contexts in which they are uttered to pragmatics, then the undermining of this assumption would show that truth-conditions are not the proper target of semantic theory.

Consider sentence (1), from Travis (1997):

1. These leaves are green.

Travis describes a scenario in which Pia has painted the red leaves of her Japanese maple green. She is visited by an artist and a botanist. The artist is producing a collage and asks Pia if she has anything green that she can contribute. The botanist is conducting a study of photosynthesis in green plants and asks if Pia has any green leaves for her study. Pia responds to both requests with sentence (1). Travis' intuition is that Pia has said something true in the former case, but false in the latter.[4] As she has, in both cases, apparently applied the same predicate to the same entities, it is hard to see, on the traditional picture, how this could be. Such examples aim to show that the influence of context on truth-conditions goes beyond any facts about the semantic properties of the sentence. On the assumption that the meaning of the relevant expressions (i.e. 'green') are invariant between contexts, this is taken to show that truth-conditions

[4]Hansen and Chemla (2013) provide experimental evidence that such intuitions are widespread.

are not the meanings of sentences. Thus semantic theory should not aim to assign truth-conditions to sentences.

The conclusion drawn from such cases is that semantic theory needs reshaping. If natural language sentences don't determine truth-conditions, then a semantic theory that assigns them truth-conditions must be flawed. I will refer to advocates of this position, who take the influence of context on truth-conditions to outstrip that specifically called for by lexically encoded variables, 'radical contextualists'.[5]

3. Standard responses

While the argument discussed above has convinced many theorists, truth-conditional semantics has had substantial success in accounting for a wide range of linguistic phenomena. Our understanding of the behaviour of linguistic expressions such as quantifiers, modals, attitude reports, etc. has been increased massively in the past half century or so in virtue of the formal semantic approach of investigating the interpretations of such expressions by stating the ways in which they contribute to the truth-conditions of sentences in which they occur. For this reason, it might reasonably be worried that a wholesale rejection of the approach would leave semantics in a bad place. Fortunately, many theorists have come to truth-conditional semantics' aid and responded to the above argument. I will now discuss two of the main ways in which this has been done.[6]

The fact that certain sentences determine a truth-condition only in context is not news to formal semantics. Kaplan (1989) provided a method for accounting for such phenomena within the framework of truth-conditional semantics. He distinguished between the *character* of an expression, its stable meaning, and the *content*, what it contributes to the truth-conditions of an utterance in which it occurs. With this distinction, we can see how truth-conditional semantics can account for context-sensitivity. For certain expressions, such as 'I', the character can be stated as a rule for determining how a content is determined in a context. Kaplan gives the rule: "'I" refers to the speaker or writer' (505). This tells us that the content expressed by a use of 'I' is the person who used it. Generalised, semantic theory aims to determine, not truth-conditions *per se*, but functions from specific aspects of contexts to truth-conditions.

[5]Note that I am using the term more widely than is typical. The term 'linguistic pragmatism' is also often used to describe this position. See e.g. Neale (2004).

[6]There is a third main strategy that I will not, for the purposes of this paper, discuss. This is the strategy of various forms of relativism (e.g. Gross 2005; MacFarlane 2014; Predelli 2004).

The first main response to the above arguments, the *moderate contextualist* approach, takes this as its lead. It is argued that, just as truth-conditional semantics is not undermined by the fact that the sentence 'I am in Paris' can be used to say something true in some contexts but not in others, it should not be troubled by the fact that 'These leaves are green' can express something true in one context but false in others.

While the context-sensitivity of indexical expressions like 'I' is obvious, other expressions may be more *covertly* context-sensitive. But, the fact that they contribute different things to the truth-conditions of utterances in which they occur indicates that, given a proper syntactic/semantic analysis, they will be context-sensitive in just the same way. If this is so, then the fact that sentences containing 'green' can express different truth-conditions in different contexts is no more problematic than the fact that sentences containing 'I' can. The arguments above work to the extent that the contextual features required to take speakers from sentence meaning to truth-conditions are extra-semantic. However, Kaplan's character/content distinction shows how context-sensitivity can be accounted for within semantics proper. If, like 'I', the stable semantic value of 'green' specifies specific aspects of the context needed to determine its contribution to sentential truth-conditions, then the role of context can be accounted for within semantics proper. The moderate contextualist thus reduces the context-sensitivity of 'green' to the better-understood examples of context-sensitivity like 'I'.

Various theorists have applied this approach to the example above.[7] Szabó argues that the application of 'green' involves two contextually determined features: part and contrast class. Here is a Szabó-inspired lexical entry for 'green':

2. $[\![green]\!] = \lambda c \lambda x.\ \text{green}'(x, \text{part}(x, c), \text{class}(c))$

This is a function that takes a context (c) as an argument, and then returns a function from objects (x) to truth-values, which will output true iff a contextually salient part of the object is green relative to a contextually salient contrast class. This can allegedly handle the Travis case because when talking to the artist, the contextually salient part of the leaf is its green outer surface; whereas in conversation with the botanist, the salient part is the red surface under the paint.

[7]E.g. Hansen (2011), Kennedy and Mcnally (2010), Rothschild and Segal (2009) and Szabó (2001).

The second main strategy for defending truth-conditional semantics is *semantic minimalism*. The minimalist, like the moderate contextualist, takes as her starting point a universally acknowledged feature of the history of semantics. It has long been noted that native speakers' intuitions do not directly reflect the purely semantic properties of the sentences they are considering. In any actual conversational context, an utterance will convey, as well as the literal meaning of the expression uttered, various kinds of additional information. Pragmatic features of this kind often confound semantic judgments. For example, most people, on reading sentence (3), will interpret it as saying that Michelle got out of bed and *then* put on her shoes.

3. Michelle got out of bed and put on her shoes.

However, many semanticists and philosophers take the temporal interpretation of 'and' to be *pragmatically implicated*, rather than given by the semantics proper.[8] Strictly, it is claimed, this sentence says only that Michelle both got out of bed and put on her shoes. Utterances of (3) typically convey temporal ordering, but in virtue of general constraints on rational communication, not the expression's linguistic meaning. In general, it is cautioned that when doing semantics, we must be wary that intuitions about the meanings of utterances can lead us astray due to their sensitivities to extra-semantic influence, including pragmatic notions such as communicative appropriateness.

The minimalist takes this as her paradigm. While our intuitions in the above cases suggest that one and the same sentence can express different truth-conditions in different contexts, we should not infer from this that the truth-conditions of this sentence do in fact vary between contexts. Pia's leaves are either green, in which case (1) is true whether talking to an artist or a botanist, or they are not, in which case (1) is false in both cases.[9] Our intuitions about the changing commitments of these claims should not be taken as the final word, clouded as they are by notions of communicative intent and cooperation. It would be so *misleading* and *uncooperative* for Pia to make this assertion in response to the botanist that we are tempted into thinking that it is literally false. The way in which these leaves are green is not at all what the botanist is looking

[8]Although this has been challenged. See, e.g. Lepore and Stone (2014).
[9]For my purposes I will assume that, upon final metaphysical analysis, Pia's leaves are indeed green, but nothing will turn on this assumption. What matters is just that the interpretation is invariant.

for, and so communicative norms rule out the acceptability of her saying this. But, the minimalist claims, what she said was nonetheless true.

The problems with these approaches are well known. The moderate contextualist position appears *unstable*. Its main motivation is precisely the kinds of intuitive differences in truth-conditions seen in sentence (1). However, our intuitions seem sensitive to an indefinitely large range of contextual features. Radical contextualists love to contrive novel situations in which the correct application of a predicate appears to depend on factors not yet incorporated into semantic theory. Consider, for example, situations in which Pia's leaves are unpainted but covered with green moss, or painted with glow-in-the-dark paint, and so on.[10] As the moderate contextualist is, by the logic of her theory, committed to accounting for such data there seems to be no non-arbitrary limit to what her theory must incorporate. The Kaplanian trick works well for indexicals like 'I' precisely because the effects of context on their contributions are well-defined.[11] When the influence on truth-conditions becomes as unconstrained as it seems to be with terms like 'green', this strategy leads to more and more complex theories. In order to remain moderate, this strategy requires there to be a relatively small number of contextual influences, each of which is specifically called-for by the semantic value of the expression. The power of the radical contextualist line consists in showing just how difficult it is to non-arbitrarily draw a boundary around the genuinely semantically significant features of context.

Another way to put the difficulty is that truth-conditional approaches to semantics seem forced to reduce *polysemy* to either *ambiguity* or *indexicality*. The context-sensitivity characteristic of ambiguity and indexicality can be handled truth-conditionally in finite ways: simply list all the distinct meanings of the ambiguous expression, or list all the contextual parameters to which the extension of an uttered expression is sensitive. That the 'axes of variation' for both ambiguity and indexicality are both determinate and specified antecedently to the determination of semantic content in a particular context is crucial for their integration into standard semantic theories. Non-finite lists, either of distinct meanings or contextual parameters, pose insuperable difficulties for theories of language use and acquisition. However, polysemy seems unsuited for such finite treatment. Of course, this has been questioned. One standard such approach involves identifying some finite list of 'axes of variation' as in Rabagliati, Marcus, and

[10]See Travis (1985) and Hansen (2011) for examples of the wide variety of influences on such ascriptions.
[11]For the most part, although see e.g. Cohen and Michaelson (2013) for a discussion of some of the complications that arise even for these better behaved expressions.

Pylkkänen (2011) and Vicente (2015). If an account of polysemy along these lines is successful, this would go a long way towards rebutting the radical contextualist challenge. However, it is worth noting that even defenders of this approach typically deny that it will be sufficient to cover all apparent cases of semantic variability. If the radical contextualists are right, then the range of factors that can influence the meaning of expressions in particular contexts is both indefinite and unspecified prior to actual usage. The moderate contextualist can *approximate* this phenomenon by adding contextual parameters when observed, but this will lead to more and more complex models which will not, in principle, capture the full range of semantic behaviour.

The minimalist is beset by a much more basic worry. The truth-conditions they posit just don't fit with the data. Given the flexible way in which we apply the term 'green', what grounds could we have for singling out a particular set of objects as the green objects? Are Pia's utterances both true or both false? We seem to have no way of settling this question. It just isn't clear what it means for something to be green *simpliciter*. This problem is exacerbated when the minimalist considers other expressions like 'tall' and 'ready'. The minimalist is thus faced with a dilemma: either semantic theories are properly responsive to our linguistic intuitions, in which case their predictions seem false, or they are not, in which case it becomes unclear how to empirically assess our theories at all.

The radical contextualist thus puts truth-conditional semantics in a bind. The data don't seem amenable to an invariantist analysis. However, in order to provide a contextualist analysis, we seem forced to accept more and more contextual parameters, with no clear stopping-point. The proposed alternative is to reject the claim that we are aiming to account for truth-conditions in the first place.

4. An alternative approach

What is shared by both the proponents of the argument and those responding to it on behalf of truth-conditional semantics is the assumption that if there is a mismatch between the target expressions' genuine semantic properties and those attributed to them by the semantic theory, e.g. if the former is context-sensitive in ways that the latter fails to incorporate, this is sufficient to show that the theory is in need of replacement. The advocates of the argument claim that there is such a disparity: natural language sentences do not determine truth-conditions, but truth-conditional semantics assumes that they do. Responses to the argument

aim to deny this and show that natural language sentences do in fact determine truth-conditions, either by denying that natural language sentences are as context-sensitive as has been supposed or by incorporating this context-sensitivity into the semantic theory.

My approach differs from both the advocates of and respondents to this argument in rejecting the assumption that the disparity between theory and target must be eliminated. Even if we cannot isolate all the contextual features relevant to determining the truth-conditions of an utterance in a context, truth-conditional semantics may still be a perfectly legitimate enterprise. The methodology of much good science centrally involves intentionally creating disparities between scientific representations and their targets.

Models provide an indirect way of investigating the world. By producing a model, we create a surrogate for its real-world target. We can then study the properties of the model and, under certain conditions, make inferences about the properties of the target.[12] Of course, no model will perfectly resemble its target. If it did, there would be no benefit to studying the model in the first place. So this means that theorists must make a choice about which properties of the target to include in the model. In this way, idealisation and abstraction are central to model building. A model may thus include entities known not to exist, or attribute properties to target entities that they are known not to have, or leave out entities or properties that are known to be found in the target system. For example, many of the models of population genetics assume that the likelihood of a parent passing on one allele to an offspring is independent of the likelihood of its passing on any others. However, due to genetic linkage, this assumption is known to be, in general, false. This falsehood is intentionally introduced in these models in order to facilitate explanation and prediction.[13]

Crucially, models call for a different mode of evaluation than theories. Theories may be true or false. However, such evaluations are unhelpful in discussing models. Because models invariably misrepresent, or at least only partially represent, their targets, they do not really aim at

[12]For a good overview of the role of models in science, see Frigg and Hartmann (2006).

[13]It is important to stress, so as to avoid confusion, that my use of the term 'model' is not the use familiar from mathematics, wherein a model is a set-theoretical structure (or set of such structures) used to interpret a formal language. I am using the term in the sense familiar from the empirical sciences, wherein a model is an entity which can be studied in order to find out about some other phenomenon. A model organism, e.g. a lab rat, will be a model in the latter sense (my sense), but not the former. This differentiates my proposal from that of Barba (2007) who argues that we should treat the formal languages of natural language semantics as interpretable by multiple models in the mathematical sense.

truth. Instead, models are best evaluated on the basis of how useful they are. More specifically, they are evaluated on how useful they are in solving some particular problem. Modelling is thus an alternative to describing, not a particular way of describing. The distinction can be seen clearly in cases in which a theory which is known to be false is helpfully utilised in a particular model. For example, consider a computational model of, say, a bridge which is developed to determine whether it will withstand the forces enacted upon it by passing traffic. Such a model will be useful to the extent that it accurately reproduces the behaviour of the real bridge, and is not rightly criticised for utilising the falsified assumptions of Newtonian mechanics.

My thesis is that the assignment of truth-conditions as sentential semantic values in semantic theory should likewise be treated as a modelling idealisation. Therefore, the fact that natural language expressions are context-sensitive in ways unlike that of the semantic values assigned to them by semantic theory need not be seen as indicating that the semantic theory needs to change any more than genetic linkage shows that the models used in population genetics must be replaced. In both cases, we should think of the properties of the model as approximations introduced in order to facilitate theoretical goals. A semantic analysis, on this proposal, is thus not a description of the semantic properties of natural languages, but instead a language in its own right, which is in certain ways simpler than the natural language it was designed to illuminate. Semanticists can then study the properties of this system as a proxy for studying the complex behaviour of the natural language. A good semantic model will be one which, *for the purposes of the current inquiry*, adequately reproduces the behaviour of the natural language phenomenon it targets, and suggests as yet unobserved behaviour of such a phenomenon which can then be tested.

The claim that semantic theory should be viewed as a model-building science is the first part of my proposal. The second is the closely related, but logically independent, idea that semantics should proceed via 'multiple-modelling'. According to the 'multiple-models approach', scientific inquiry aims at the production of a battery of different models, each of which makes different idealisations in the service of achieving different goals.[14] An idealisation is justified for use in a particular model to the extent that it enables the modeller to fulfil these goals. As different

[14]For the origins of this strategy see Levins (1966). For more contemporary defence and elaboration, see Godfrey-Smith (2006), Weisberg (2007), and Wimsatt (2007b).

modellers may have different goals, they may make different idealising assumptions in constructing their models.

The central motivation for such a picture is the existence of *trade-offs*.[15] Trade-offs occur when different theoretical desiderata cannot be simultaneously maximised in the same model. One common source of trade-offs is the existence of causally heterogeneous systems. That is, systems the behaviours of which are influenced by a wide variety of different kinds of factors.[16] When studying such systems, there is typically a choice between providing a relatively simple model, which may be able to capture general and relatively context-insensitive properties of the system; or providing a more complex model, with more parameters, which may more accurately predict the system's behaviour in particular cases. Crucially, the existence of trade-offs distinguishes the multiple-models approach, according to which a plurality of models is *required* to achieve our various scientific goals, from the universally accepted claim that at a given point in scientific inquiry we are likely to be best served by creating and comparing distinct scientific representations.[17]

Levins (2006) provides a beautiful example of the multiple-models approach in population ecology.[18] He considers how one would calculate the dynamics of a population of ants, and shows how to begin with fairly simple dynamical models, and gradually and systematically de-idealise them. Consider the full range of proposed models, from the simplest containing just one variable ($dx/dy = -mx$, i.e. the change in the population (x)) as a function just of the mortality rate (m)), up through the series of more complex models incorporating influences on population size like predator/prey interactions, abundance of food, temperature, etc. He argues that the simpler models are more likely to be general, in that they are abstracted away from the specific details of the target system enough that they can apply to many distinct systems. The simple equation relating population and mortality rate will be appropriately applicable to

[15]See Matthewson and Weisberg (2009).

[16]This is, of course, a rough and ready description of the kind of system I have in mind. For a more detailed account, see Ladyman, Lambert, and Wiesner (2013).

[17]These distinct views of the role of multiple models in science are sometimes conflated. Wimsatt (2007a), for example, in an often-cited discussion of the role of false models in scientific inquiry, discusses modelling practices which support both uses. For example, he mentions the way in which an oversimplified model may serve as the starting point in a series of more realistic models. This is of course true, but does not suffice to motivate the multiple-models approach in the stronger sense I am advocating in this paper. For this, these simplified models must remain useful to scientists as alternative, more realistic models are introduced, not merely act as stepping stones. I believe it is important to keep these two views separate, as their conflation can undermine the interest of the multiple-models picture.

[18]See also Wilson (2008, 180–181), for a similar case from classical mechanics, involving the range of parameters that could be invoked in accounting for the behaviour of billiard ball collisions.

populations of many different kinds of creatures. However, target-specific influences of extra factors will mean that such simple models are likely to give only qualitatively accurate predictions. To generate any quantitative predictions, these extra factors must be incorporated. But these will vary from target to target. As Levins points out, the movements of social insects like ants, unlike flies, depend on community-level properties such as the food levels of the nest (746). This will call for a different modelling strategy.

Levins case study shows both the need for idealisation, and for a multiple-models approach. Firstly, it is clear that, in discussion of a particular target, say a population of insects, there is simply no well-defined 'cut-off point' at which to stop including information in our models. An indefinite number of factors may influence the properties we, as modellers, are interested in. For this reason, our models will, of necessity, be somewhat idealised. While there is no cut-off in nature, there must be in our representations of it. Secondly, the choice of where exactly to make this cut is a pragmatic one. The question of what to include is dependent on the question of what we want the model to do. Do we want a simple model of how animal populations *in general* fluctuate? If so, a useful model will be highly abstract, but in virtue of this abstractness will fail to precisely capture the behaviour of specific targets. On the other hand, if we want to accurately predict how a particular system will behave, more information is better. But the inclusion of such information will thereby ensure that this model has limited scope: all those systems in which these parameters are not significant, or behave differently, will not be suitable targets for this particular model. Levins makes the point: 'But the kinds of observations and the specific questions we can ask [about different models] are different. Our work depends both on generalisation and respect for specificity' (747).

Just such a trade-off appears to confront the semanticist. The aim of semantic theory is to understand the system within which the meanings of natural language expressions are determined.[19] However, given the range of contextual influences on such meanings, it seems that empirical

[19]Note that what I say should apply equally to approaches which view semantics as aimed at understanding *sentence meaning* (Kaplanian *character*) as to those that view the target as *utterance meaning* (Kaplanian *content*), or even more fine-grained notions like *assertoric content* (Rabern 2012). While I am focusing on the semantic properties of utterances, it is an essential desideratum of a theory of sentence meanings that it be able to predict the meanings of these sentences as uttered in particular contexts. One reason for this is that judgements about utterances are what provide empirical evidence for or against all semantic theories. Standard approaches forge this connection by making sentence meaning a function from context to utterance meaning.

adequacy is likely to come at the cost of highly complex, and therefore highly specific, accounts of the behaviour of particular expressions in particular contexts. On the other hand, general models, capturing the central behaviour of these expressions, as found across a wide range of situations, are liable to make inaccurate predictions when applied to particular cases. This situation, in which generality can be bought only at the cost of empirical adequacy, and vice versa, is precisely a trade-off, and thus suitable for a multiple-models approach.

For this reason, I believe that a multiple-models approach can and should be adopted in the field of semantics.[20] Natural language meaning emerges from a causally heterogeneous system. *Prima facie* what an utterance of a sentence means in a particular context depends at least on the speaker's language-specific knowledge, general facts about that speaker's cognitive system, external features of the context of utterance, and the cognitive states of the interlocutors.[21] Because of this complexity, linguistic modellers face trade-offs. In order to account for the complex context-sensitive behaviour of natural language expressions, we are faced with a choice: incorporate this context-sensitivity into our models, or (partly) idealise it away. The former approach will lead to complex models that can accurately describe specific systems, while the latter will produce simple, general, and tractable models, which will fail in certain cases. If we imagine plotting models on a graph, with models ranked with respect to their empirical adequacy along the *Y*-axis, and generality along the *X*-axis, the existence of a trade-off is exemplified by there being no models plotted in the top right-hand corner. Contextualist models will cluster around the top left-hand corner, with minimalist models around the bottom right. Note that there is still room for evaluation here: better models are further away from the origin. What is denied is that there is a unique best model.

The argument against truth-conditional semantics can now be re-evaluated in light of the multiple-models strategy. The radical contextualist argument suggests that there is no finite set of contextual features that can be invoked to determine sentential truth-conditions. The truth-conditional semanticist is thus faced with a dilemma. The moderate contextualist aims to account for these kinds of data, but seems to lack a non-arbitrary

[20]Yalcin (2017) argues for a model-based approach to semantics, but does not extend this to a multiple-models approach.

[21]For an argument that the mental states of both the speaker and hearer are involved in determining reference, and therefore truth-conditions, see Heck (2014). For a discussion of the role of facts about context, see Stojnić, Stone, and Lepore (2013).

place to stop adding complexity to her theory, as has been stressed both by minimalists and radical contextualists. The minimalist denies that her theory is responsive to such data, at the risk of apparent empirical inadequacy. The unpalatability of these two options drives some to reject truth-conditional approaches in general. However, viewing semantics as a multiple-models discipline undercuts this dilemma. Finite semantic models can be viewed as approximations of the linguistic phenomena they target. These approximations can be more or less drastic, with the minimalist models at the most idealised end of the spectrum, and more complex moderate contextualist models at the other end. As there is no need to select *the* semantic analysis of a given phenomenon, there is no need to arbitrarily find a cut-off point. This provides an alternative to the usual 'elimination by counter-example' approach: failures of a model may indicate that the model ought be replaced, but they may alternatively reflect the idealisation which enables the model to maximise generality, or perhaps some other theoretical virtue.

This approach is radically contextualist in that it denies the existence of any privileged set of contextual parameters. However, the problem with the radical contextualist approach seems to be that it precludes semantic theory from making predictions. If we cannot say, in advance, what aspects of context are determinants of content, then it seems any content is, at least potentially, consistent with any utterance. The multiple-models approach does not have this defect. Despite its radical contextualist overall outlook, each model contains a finite set of contextual parameters. Some (minimalist) have very few, whereas others (moderate contextualist) have more, but all are finite. This finitude thus enables them to make specific, testable predictions. From the other direction, this approach is no longer susceptible to the central worries with the minimalist and moderate contextualist approaches. The worry with the minimalist approach was that it failed to cohere with the observations. On the new, multiple-models approach, while some models do indeed make bad predictions, this is not (necessarily) grounds for rejecting them, as it would be if semantic analyses were viewed as (truth-apt) theories. A model, remember, is evaluated with respect to its *usefulness*, not its truth. And so, as in the case of Levins' simple but general ecological models, if our model can be shown to be useful, despite making some bad predictions, it can/should be retained. I shall argue in the next section that minimalist models should be viewed exactly in this way: these simple models earn their keep by capturing generalisations about language, even though, in virtue of abstracting away from context, these generalisations will often fail in particular cases.

The objection to moderate contextualism centred around the *instability* of this approach. The worry was that the same kind of argument which motivated those parameters the moderate contextualists were willing to let in could be given by the radical contextualist for indefinitely many more parameters. Consistency then seemed to collapse the former position into the latter. Again, this is not a problem for the multiple-models approach. On this approach, there is no requirement to select some privileged stopping point along this series of contextual influences. Some models would include very few parameters, while others would include lots. Which such models theorists adopted would then be a pragmatic decision based on what they wanted to capture, not an empirical question of which such parameters genuinely influenced truth-conditions.

According to this picture, then, a semantic model is an artificial language which bears some resemblance to a particular fragment of a natural language. In the natural language case, utterances of expressions have extensions, but which extension they have is context-sensitive, and which features of the context will be relevant is highly open-ended. That is, natural language expressions do not on their own determine extensions, or even functions from context to extensions. The overall context of an utterance is needed for this, including the mental states of conversational participants and extra-mental features of the context. The formal model works by *stipulating* that some, and only some, contextual features matter. We can then study the behaviour of such a language, and compare it to the behaviour of the natural language fragment it aimed to illuminate. This comparison serves to generate predictions. A good model will make accurate predictions for at least some such cases. However, due to the indefinite nature of the contextual determinants of meaning, no model will work for every case. This creates the need for multiple models, making different idealisations to cover these different cases. These different models can now also be compared with respect to other theoretical virtues, such as simplicity or generality. As suggested above, minimalist models can be quite general. Contextual-sensitivity is widespread in that for large parts of the lexicon, context is *able* to influence an expression's meaning. But in many cases it does not. This is the core insight of the minimalist, and explains why minimalist models can have a wide coverage.[22] However, because context can always play a role, in many cases a contextualist approach will fit better with our observations.

[22]That, in many cases, the content of an utterance can be extracted in relative ignorance of the context of utterance provides strong reason for believing in just this sort of 'default' semantic value posited by minimalist models. This point is stressed by Cappelen and Lepore (2008). It seems likely that the 'first

Perhaps a mathematical analogy is helpful here. Certain functions can be represented as an infinite sum of powers of their variables. For a very simple example, $1/(1-x)$ can, when $-1 < x < 1$, be represented as $1 + x + x^2 + x^3 + x^4 + x^5 \ldots$. These sequences often involve more manageable mathematics than the functions they represent, and so can be used as helpful proxies. But, as these sequences are infinite, the precise values of the function cannot be finitely computed in this way. So applied mathematicians truncate these sequences, summing the values of the first n members of the sequence and ignoring the rest. Under certain conditions, this procedure can be used to generate an arbitrarily accurate estimate of the value of the function, with proximity to the true value increasing with n. We can think of the relationship between semantic values as posited in semantic theory and the real semantic properties of natural language expressions as standing in the same sort of relationship that truncated series stand in to the functions they represent. If the radical contextualist is right, real-world intensions, incorporating all the ways in which extensions depend on context, are not finitely statable, just like these infinite sequences. In the mathematical case, there is no question of what the right point for truncation is, in general. This will depend on how accurate we need to be. My claim is that semantic analyses should likewise be viewed as approximations, and so the question of how many contextual parameters the true semantic theory should posit likewise doesn't arise. Just as in the mathematics case, this depends on what we want our semantic theories to do.

The dispute between the minimalist and the contextualist (radical and moderate) stems from a tension between two ostensible facts about language. On the one hand, language is amazingly systematic.[23] With little explicit instruction, speakers within a linguistic community coordinate on the semantic properties of a huge range of expressions. This systematicity seems necessary to account for the ability to correctly interpret and produce indefinitely many utterances and to communicate successfully across different contexts. On the other hand, one can only be amazed by the subtlety with which language is put to use. What people

pass' lexical entries one finds in linguistics and philosophy papers dealing with meanings of expressions will often serve such a default role. For example, while '$[\![\text{green}]\!] = \lambda x.x$ appears green' to normal observers in normal lighting conditions' fails to predict our intuitions about ink, lamps, (some) painted leaves and other favourite subject matters of the radical contextualists, such an entry would probably cover many, if not most, uses of the term 'green'.

[23]N.B. While it is closely related, by 'systematicity', I do not mean to refer to the property of languages which enables permutation of syntactically similar expressions without undermining the meaningfulness of the sentence. I mean the more general, and vaguer, idea that natural languages are not mere collections of facts but rule-governed systems.

are able to communicate with a particular utterance can vary significantly depending on the conversational context. In proposing general pictures of semantics, the minimalist seems to focus on systematicity while contextualists are more impressed by subtlety.

This tension is dissolved by locating these phenomena at different levels of *granularity*. At a high level of abstraction, we see natural language's systematicity. Broadly accurate, but exceptional, generalisations express dependencies between different linguistic expressions. The existence of such generalisations is a precondition for the acquisition and use of language. However, when we shift our attention to the lower, more fine-grained, level(s) of actual communicative behaviour, the exceptions to the rules become apparent. Subtle contextually driven deviations from the general rules start to appear as a result of the fact that communication is a rational activity, and thus sensitive to all the complex intentions of human speakers. The minimalist and the moderate contextualist assume that there is some unique level of granularity at which semantic theory applies, and for which any causally more complex levels involve strictly pragmatic effects. They differ in where they draw this line, with the minimalist claiming that semantics proper is fairly abstract, and almost all contextual influences are properly speaking pragmatic. The challenge to the moderate contextualist is finding a principled way to draw the line so as to let in more factors than the minimalist, but without opening the floodgates and letting in just any contextual factors, *à la* the radical contextualists. The radical contextualist agrees that this can't be done, and so rejects the existence of this privileged level, claiming that there is no finite set of factors relevant to genuinely semantic, truth-conditional, properties. I argue that this does not undermine the truth-conditional methodology, but calls for a re-evaluation of its assumptions. Crucially, by denying that there is a privileged level at which truth-conditional content is determined, we must view descriptions at each level of granularity (i.e. with varying degrees of contextual-sensitivity) as each providing imperfect approximations of the target.

This meshes nicely with certain radical contextualist pictures such as that proposed in Recanati (2010) and Neale (2004). According to these views, interpretation is guided by a relatively stable compositional system operating on default lexical entries, but which is susceptible to top-down influence (modulation) along the way. The multiple-models approach is best suited for describing such a system. Some (minimalist) models capture the properties of these default meanings and their interactions. These models will capture the default operations of the system.

However, they will be confronted with counter-exemplifying observations whenever more than these default meanings are in effect. Rather than rejecting these simple models, and thus the generalisations that they capture, the multiple-model approach also makes room for more particularised (moderate contextualist) models, which can explain why the predictions of these default models failed, by pointing to the specific contextual feature that made the difference.

Models in line with the minimalist approach capture the structural properties of language, which enables them to account for its systematicity. At a coarse-grained level, we can see monadic predicates, such as 'is green', and 'is a prime number', as playing a particular role in the determination of truth-conditions. Minimalist-inspired approaches explicate just what this role is. However, once we pay attention not just to the structural facts, but also the influence of context, we notice differences in these expressions' behaviours. The exceptions to this systematicity involve the influence of contextual factors that are excluded from simpler models. Contextualist-inspired models incorporate these influences, making better predictions in specific cases. However, by incorporating this context-sensitivity, we lose the ability to capture some of these structural dependencies. The multiple-models approach enables us to switch between focusing on the high-level invariances and the lower-level complexities.

5. Entailment: a case study

As an example of the different strengths and weaknesses of these strategies, consider the explanation and prediction of entailment patterns. One central task of semantics is discovering the structural properties of sentences that account for relations of entailment between them. As entailment is essentially a matter of truth-preservation, it is far from clear that radical contextualists will be able to make use of this notion. If natural language sentences don't have truth-conditions, it is difficult to see what is supposed to be preserved in valid arguments.[24] My aim in this section is to show how we can throw out the bathwater of a simplistic view of natural language meaning, while retaining the baby of a theory of natural language entailment, and with it the ability to utilise entailment data in evaluating theories. I shall do so by arguing that entailment itself

[24]This is not to say that making sense of entailment is impossible within such an approach. An account of natural language entailment given strictly in terms of provability could be offered. But, this is against at least the spirit of many radical contextualist approaches.

should be viewed as context-sensitive,[25] and that a multiple-models approach can capture some basic facts about this context-sensitivity in ways that standard defences of truth-conditional semantics cannot. To give a very simple case, consider the following sentences:

4. This apple is green.
5. This is a green apple.[26]

Intuitively, the inference from (4) to (5) and vice versa is valid. A semantic theory should capture such facts. However, if 'green' is context-sensitive, this entailment will not be predicted. If 'green' is context-sensitive, then it may contribute different things to these different sentences. This has the same effect as positing an ambiguity: the inference is no longer valid. Unless the semantics of 'green' *guarantees* that its occurrences contribute to truth-conditions in the same way, then we cannot be sure that (4) and (5) will be true in the same contexts. The minimalist model, by treating 'green' as invariant, ensures this, thus enabling the explanation and prediction of such entailment patterns.

Of course, as with all simplifying idealisations, the minimalist strategy breaks down when the factors idealised away play a significant role. By assuming that 'green' contributes the same property to any sentences in which it occurs, we fail to predict the apparent fact that sentence (1) can express different propositions in different contexts. By focusing on the general, structural properties of language, we ignore the ways in which context does matter. By developing the relatively simple idealised models of natural language provided by the minimalist approach, we can understand compositional structures and dependencies. We can then use these as scaffolds onto which we can attach more complex

[25]It is worth distinguishing between two ways that entailment could be viewed as context-sensitive. A traditional picture of entailment in a context-sensitive language has it that sentences containing context-sensitive expressions may entail something only on particular evaluations of their context-sensitive expressions. For example, 'I am happy' will entail 'David Kaplan is happy' if and only if David Kaplan is the speaker of the former utterance. This traditional picture thus distinguishes between context-sensitive sentences, for which the sentences they entail will vary according to their contexts, and non-context-sensitive sentences for which entailment patterns can be identified independent of context. My claim goes beyond this traditional view in saying that the boundary between those sentences we ought treat as context-sensitive and those we ought treat as context-invariant is itself variable. Thus, one and the same sentence may sometimes be treated as invariantly entailing some other sentence, and sometimes be viewed as having entailments only given contextual evaluation.

[26]I am assuming that 'green' here is acting as a purely intersective adjective in both cases. If 'green' does not work this way (i.e. being *green-for-an-apple* does not entail being green *simpliciter*) then I would need to change my example. But the issue of whether 'green' is intersective is orthogonal to the issue of whether it is context-sensitive, and the point could be restated using another predicate which is the latter but not the former.

(e.g. context-sensitive) machineries in order to account for deviations from these general patterns.

Consider sentence (6):

6. This is a green apple, but it isn't green.

Sentence (6) is predicted, if the above inference pattern is indeed valid, to be contradictory. And indeed, *prima facie* this prediction is correct. However, there are contexts in which it can be truly uttered. Namely, those in which the two uses of 'green' express different contextually determined properties. Consider a situation in which we are sorting through a barrel of Granny Smith and Red Delicious apples. We are cutting them open to see if they have been infected with a fungus that turns their flesh green in order to test the hypothesis that this fungus affects all and only the green (Granny Smith) apples. I can utter (6) to note that I have found a counter-example to this hypothesis: a green(-skinned) apple which isn't green (inside).

We can see the contextualist strategy as that of *de-idealising* the minimalist models. While generalisations are found by abstracting away from some kinds of complexity, in cases where this complexity 'makes a difference', we must re-incorporate it into our models. By introducing contextual parameters, we can precisify our models so as to handle the context-sensitivity evinced by the earlier examples.

I anticipate that people may balk at the idea that entailment patterns have counter-examples. However, it is worth keeping in mind that I am claiming only that certain inference patterns *in natural language* are not perfectly general. There are two ways of interpreting this. Firstly, one could think that the logic that *best* captures natural language inference is a traditional system wherein entailment is defined strictly so as to have no exceptions. If we maintain this picture of logic, my claim would be that natural language is only imperfectly captured by such a logic. This view was basically universally accepted by philosophers of language in the early to mid twentieth century (e.g. by Strawson 1952 and Carnap 1962).[27] It is less popular now although the optimism sparked by

[27]In some sense, my view is an inversion of some of these traditional views. On one reading, Carnap viewed natural language as an approximation of an idealised formal language. This way of putting it makes it seem like we are expressing opposite views, but in fact I believe they are based on the same sets of facts. My view, like this traditional view, stresses the gap between natural and formal languages. The difference depends on what we aim to achieve. My goal, or the goal of natural language linguistics, is descriptive: to investigate and explain the properties of natural languages. From this perspective, disparities between natural and formal language can only be viewed as cases where the latter

Montague, Chomsky, Grice, and others seems now to be wearing off, and this proposal is looking more plausible. Alternatively, the best logic for natural language could be one in which entailments do have counter-examples, as in Default Logic (e.g. Reiter 1980), wherein derivations are not necessarily truth-preserving. These proposals show that the idea that linguistic entailments have counter-examples is not untenable, although of course more development of the picture, and how it could be integrated into linguistic theory, is needed.

The standard way of ensuring that inferences like that between (4) and (5) are valid treats entailment as defined only when the context remains constant.[28] However, this will not, in general, account for our intuitions about when entailments do and do not hold in natural language. Consider the following utterance, made by Pia to the artist:

7. This leaf is green, and this chair is green too.

This utterance appears to entail the following:

8. There are two green things.

However, if this restriction on our definition of entailment is enforced, this will not be licenced by a Szabó-style contextualist semantics. Presumably, the salient contrast classes for leaves and chairs are different, and they have different salient parts, and so context must change in order for an utterance of (7) to be true. If this is so, semantic theories which restrict the definition of entailment in the way just described will simply be silent on the relation between (7) and (8). But this inference is intuitively a good one. Relatedly, as context must shift mid-sentence in utterances of (6), theories of entailment which adopt this restriction will be silent on the semantic properties (e.g. whether it is trivial or contingent) of this sentence. A minimalist model, by abstracting away from the context-sensitivity of 'green' is able to apply standard logical inference rules to account for these inferential phenomema. In our idealised minimalist model, (7) says of two distinct objects that they are within the (invariant) extension of the predicate 'green', and so existential generalisation

deviates *from* the former, and so I describe formal languages as approximating natural language. Carnap, on the other hand, was interested in the normative project of working out how people, especially scientists, ought reason. From this perspective, the formal languages of mathematics and logic were viewed as ideal systems, which were imperfectly mimicked by natural languages, which ought therefore be improved by explication.

[28]E.g. Soames (2010, 101).

licences (8). By excluding context-change scenarios, the standard approach makes the explanation of such cases impossible.

Here we have exactly the kind of trade-off that motivates a multiple-models approach. The contextualist cannot account for the above general inference patterns.[29] If 'green' is context sensitive, inferring from (4) to (5) or (7) to (8) is illegitimate, little more than equivocation. If the extension of 'green' depends on the context of utterance, and this context has shifted between utterances, there is no guarantee that this extension has not also shifted, undermining the validity of the inference. Restricting our notion of entailment so as to apply only in stable contexts simply ignores the problem. But the minimalist can't account for the truth of certain utterances of sentence (6). The multiple-models approach, however, can make perfect sense of this. By ignoring the context-sensitivity of 'green', one kind of model (simple, general, minimalist) can validate these intuitive inference patterns. By re-incorporating this context-sensitivity, another (complex, specific, contextualist) model can account for this inference's failures. At a certain level, entailment can be given a purely formal analysis: the same symbol (e.g. 'green') in one sentence contributes in the same way to truth-conditions as it does elsewhere. Minimalist models provide a framework for such an approach. However, given that natural language is context-sensitive in more complex ways, we oughtn't focus only on this level. Moderate contextualist models serve to de-idealise in order to explain the cases in which these more abstract models fail, but do so at the cost of no longer explaining the original entailment pattern.

This reflects a popular view of the relationship between theories of a higher and lower granularity.[30] In the special sciences, generalisations can be found. However, these generalisations have exceptions. These exceptions are accounted for by moving to a more fine-grained level of description. At this level, exceptions to generalisations can be explained. But, for precisely this reason, at this level the generalisation can no longer be truly stated.

Now, while I hope to have shown that neither the minimalist nor the moderate contextualist approach can satisfactorily handle both the

[29]Of course, this is not to say that there are not moves they could make, just that the basic machinery of their system does not naturally handle such cases. For example, if we viewed 'green' as not merely context-sensitive, but also ambiguous such that 'green' as found in the conclusion of the above argument means 'green in some way or other', while in the premises 'green' means 'green in some specific, context-sensitive way', the above inference would be semantically valid. Whatever the merits of this and similar moves, it is clear that the minimalist account, treating 'green' as semantically invariant, is cleaner.

[30]E.g. Fodor (1974).

general patterns of entailment we observe in natural language and the specific ways in which context can undermine or influence such patterns, what about the radical contextualist approach? *Prima facie* this approach does no worse than the moderate contextualist approach. While it may fail to capture intuitive entailments between sentences, it should be able to explain those cases in which such inferences fail. In fact, I think, as an approach to semantic theorizing, as opposed to philosophical inquiry, radical contextualism is not even this good. The problem is, as alluded to earlier, that radical contextualism, while likely true, is highly unexplanatory. The very flexibility which motivates radical contextualist analyses seems likewise to undermine its ability to predict any linguistic phenomena at all.

Minimalist and moderate contextualist accounts of entailment, given in terms of *sentences*, or sentences relative to a finite set of contextual parameters, may face counter-examples. Intuitive entailments may be predicted to be invalid, intuitive contradictions may be predicted to be contingent, predicted entailments may have intuitive counter-instances, etc. However, my contention is that it is better for a theory of natural language entailment to cover some range of the target, while facing counter-examples, than to fail to connect with the target in the first place. Radical contextualist approaches typically view entailment as holding between propositions, rather than sentences, and view the relationship between sentences and the propositions they express as highly variable. The problem is that this makes it unclear how to retain the view of natural language semantics as a genuinely empirical discipline. Without constraints on the relationship between sentences and propositions, we lose the connection between semantic relations such as entailment and empirical observations such as speaker judgements.

Overall, then, I believe the multiple-models approach provides the best way to illuminate and understand natural language properties such as entailment. In particular, it exhibits sufficient flexibility, in allowing theorists to propose more or less complex models as required, that both generalisations and exceptions can be understood. In this way it improves upon the minimalist and moderate contextualist approaches. But, due to each individual model in this set being constrained, it retains the connection between prediction and observation that the flexibility of radical contextualism precludes. A set of models, each capturing some aspect of the target system, but with certain characteristic flaws, is better than an approach to studying language which precludes making empirical predictions at all.

I believe that examples like sentence (1) do indicate a gap between meaning and truth-conditions. Because context is able to influence the semantic content of an expression in subtle ways, it is unlikely that many general claims about sentential truth-conditions or inference patterns will be exceptionless. What is novel about my position is the acceptance of the truth of radical contextualism as a philosophical thesis, but the not as an approach to empirical science. The multiple-models approach shows how to make sense of the combination of these two positions. While aspects of our semantic theories are not perfectly aligned with the natural language expressions that are our targets, this does not motivate the rejection of truth-conditional (minimalist or moderate contextualist) approaches. These disparities are the result of idealisations introduced to bring some order to the messy phenomena. By abstracting away from the complexities of natural language meaning, we can produce simpler models which capture broad generalisations. When we wish to focus more on the subtleties themselves, we are able to de-idealise, producing more complex models in which counter-examples and anomalies can themselves be explained. Semantics involves the proliferation of such models, each of which should enable us to understand some different aspect of natural language meaning, viewed at various different levels of granularity.

6. What is semantic theory about?

This approach suggests a re-analysis of what semantic theory is about. Linguistic theory in general is usually described as aiming to discover what one knows when one knows a language.[31] Semantic theory then aims at uncovering the meaning-related aspect of this knowledge: what does one know when knows what linguistic expressions, simple or complex, mean? The usual assumption is that this knowledge can be fully stated in finite form. That is, one's semantic competence involves specifying, in advance, something like an axiom system which determines how each lexical item contributes to the meaning of utterances in which it occurs, including identifying the set of contextual parameters to which each item is sensitive. Moderate contextualists and minimalists alike accept this basic picture, but differ on the nature of this knowledge. Some critics of truth-conditional semantics (e.g. the generativists mentioned in fn. 1) accept this picture as well, but accept the radical

[31]Or, if the intentional idiom is deemed inappropriate, what states one is in when one knows a language.

contextualist argument that there is no stable, antecedently known, set of facts which determine truth-conditions, and so infer that natural language meaning does not determine truth-conditions in contexts.

I believe that the strength of the radical contextualist argument is in showing that the question (what does one know when they know what expressions mean?) exhibits presupposition failure. That is, there is no well-defined thing that corresponds to the traditional notion of semantic meaning. The traditional picture relies on there being a principled way to separate *endogenous* from *exogenous* influences on meaning: i.e. distinguish those influences of context which semantic competence 'anticipates', from those merely pragmatic factors which do not target an element in the lexical entry for the expression and therefore cannot modify literal meaning. For example, in Szabó's entry for 'green', comparison class and part play a role endogenous to the semantic system, and so utterances of 'green' can express different contents in contexts in which different object parts or contrast classes are salient. However, this complete lexical entry rules out any other possible influence. For example, Hansen (2011) notes the ways in which observation conditions influence our willingness to apply colour terms (e.g. something may look one colour when seen from afar, but a different colour when up close). On the traditional view, at least one of these proposals must be incorrect. There is some fact of the matter as to whether our knowledge of the term 'green' literally expresses something different when applied to distant and nearby objects, and either Szabó or Hansen (or, according to the minimalist, both) has misdiagnosed this fact. The radical contextualist argument suggests that this whole debate is mistaken. There simply is no principled exogenous/endogenous distinction to be drawn. Many things can influence what we mean when we speak, and our knowledge of language does not involve cleanly determining which influence literal meaning and which influence only what is conveyed, in advance of actually producing or encountering the particular utterances we make. In such a circumstance, all we can hope for is a collection of more or less complex models, each of which takes some subset of the possible influences on meaning into account. These models include some of the features of an utterance in a context which determine truth-conditions, but as these features may vary indefinitely, no model will include all such features.

Note that this is not the same as claiming that there is no semantics/pragmatics distinction. The radical contextualist picture just developed denies that there is any antecedently determined set of influences on what is said (strictly, literally, etc.), but it doesn't follow from this that

there is no distinction between what is said and what is implicated, implied, conveyed etc. When I say 'Those leaves are green'. there is one question of what property I have attributed to the leaves, i.e. what is the set of entities such that I am claiming that these leaves are in this set, and there is another question of what I was doing by making such a claim, i.e. was I offering the leaves as an example, or suggesting that the leaves ought be removed, etc. We can get at this distinction in the normal ways, by asking whether what is said could be true while what was conveyed was false, and so on. On this picture, utterances have literal truth-conditions, and these truth-conditions are what they literally mean, but what determines these truth-conditions is not itself specified in advance by the meanings of the constituent expressions.

Nothing I have said implies that there is no distinction between influences on content triggered by expressions which, as a matter of their lexical semantics, require contextual valuation, and those which are motivated instead by more 'pragmatic' concerns such as reasoning about context. One could call the former 'semantics' and the latter 'pragmatics' if one wanted, but the radical contextualist challenge shows that this would be to preclude truth-conditional approaches to semantics. Given the success such approaches have had, and the ways in which they enable connections between linguistic theorizing and observables such as truth-value judgements, I believe we are better off retaining the truth-conditional methodology by treating all such influences *as if* they were of the former kind.[32]

The question 'What does one know when one knows a language?' may, if this picture is correct, have no unique answer. Semantic competence may be describable at multiple levels of abstraction, and each level of description may be apt in some ways and inapt in others. To know a language would be to *both* understand the high-level regularities and systematicities *and* to know how and when such generalities may fail. When we describe our linguistic competence we will then have to choose either to highlight one or the other of these aspects of knowing a language, but no model will incorporate both the regularities as well as all of the exceptions.

This picture of semantics is more in line with the ontology of other special sciences dealing with complex objects. Models of ecosystems, weather systems, economies, etc. can be produced which are tractable

[32]Note that this 'mixing' of semantic and pragmatic mechanisms is standard in, e.g. discussions of demonstrative reference assignment, wherein clearly pragmatic phenomena such as conversational salience are used to figure out which values to assign semantic variables.

in virtue of keeping causal parameters to a minimum. There will be various advantages to abstracting and idealising in this way, but there will also be costs. Because of these costs, especially in predictive accuracy, other models can be produced which make explicit reference to these other causal factors. However, to ask, for example, 'Which of these *really* represents the ecosystem?' would be simply confused. They all do, at different levels of abstraction. If the above picture of semantic competence is correct, then such an answer may similarly be plausible in the semantics case. It is a contingent empirical proposal that knowing a language is having a fully modular faculty consisting of an axiomatic description of the truth-conditions of this language. If this assumption is incorrect, as the radical contextualists claim, and the semantic values of expressions can be influenced in this top-down way, then there will arise the question of how much of this influence should be included in our models. The multiple-models approach answers this question: it depends on what we want our models to do.

7. Conclusion

In this paper, I have argued that semantics is best served by the multiple-models approach. Because of the wide range of factors that can be relevant to meaning, semantics should proceed by producing simple and general minimalist models, as well as moderate contextualist models that account for the exceptions to generalisations stated at this coarse-grained level. This approach enables us to accept the radical contextualists' claim that the truth-conditions of an utterance are not determined by the sentence uttered, while also using the powerful tools of truth-conditional semantics. By treating the truth-conditions assigned in semantic theory as approximations of the messy semantic properties of natural language, we are able to incorporate the advantages of both styles of response to the argument against truth-conditional semantics, without paying the high costs associated with complete adoption of either. This approach shows how to account for both the systematicity and subtlety of language.

Acknowledgments

This paper benefited greatly from discussion with many people over the years, especially: Josh Armstrong, Sam Cumming, Katie Elliott, Gabe Greenberg, Gabby Johnson, Kim Johnston, David Kaplan, Kevin Lande, Eliot Michaelson, Torsten Odland, Guillermo Del Pinal, Jessica Rett, Sheldon Smith, Tim Stowell, Michael

Weisberg, and Howie Wettstein. Earlier versions of this paper were presented to audiences at UCLA, CU Boulder, Institut D'Études Avancées in Paris, and the Northern Italy Philosophy Consortium in Genoa, and benefited from the discussion of attendees there. Finally, I would like to thank an anonymous reviewer for this journal for helpful commentary.

Disclosure statement

No potential conflict of interest was reported by the author.

ORCID

Gabe Dupre http://orcid.org/0000-0001-9312-4691

References

Austin, J. L. 1950. "Truth." *Proceedings of the Aristotelian Society* 24, 111–128.

Bach, K. 1994. "Conversational Impliciture." *Mind and Language* 9 (2): 124–162.

Barba, J. 2007. "Formal Semantics in the Age of Pragmatics." *Linguistics and Philosophy* 30 (6): 637–668.

Borg, E. 2004. *Minimal Semantics*. Oxford: Oxford University Press.

Cappelen, H., and E. Lepore. 2008. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Hoboken, NY: Wiley-Blackwell.

Carnap, R. 1962. *Logical Foundations of Probability*. Chicago: University of Chicago Press.

Carston, R. 2013. "Implicature, Explicature, and Truth-theoretic Semantics." In *The Semantics – Pragmatics Boundary in Philosophy*, edited by M. Ezcurdia and R. J. Stainton, 261–283. Peterborough: Broadview Press.

Chomsky, N. 2000. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.

Cohen, J., and E. Michaelson. 2013. "Indexicality and the Answering Machine Paradox." *Philosophy Compass* 8 (6): 580–592.

Fodor, J. A. 1974. "Special Sciences." *Synthese* 28 (2): 97–115.

Frigg, R., and S. Hartmann. 2006. "Models in Science." In *Stanford Encyclopedia of Philosophy*, edited by E. Zalta.

Godfrey-Smith, P. 2006. "The Strategy of Model-based Science." *Biology and Philosophy* 21 (5): 725–740.

Gross, S. 2005. "Context-sensitive Truth-theoretic Accounts of Semantic Competence." *Mind & Language* 20 (1): 68–102.

Hansen, N. 2011. "Color Adjectives and Radical Contextualism." *Linguistics and Philosophy* 34 (3): 201–221.

Hansen, N., and E. Chemla. 2013. "Experimenting on Contextualism." *Mind and Language* 28 (3): 286–321.

Heck, R. 2014. "Semantics and Context-dependence: Towards a Strawsonian Account." In *Metasemantics: New essays on the foundations of meaning*, edited by A. Burgess and B. Sherman, 327–364. Oxford: Oxford University Press.

Hornstein, N. 1984. *Logic As Grammar: An Approach to Meaning in Natural Language*. Cambridge, MA: MIT Press.
Hornstein, N. 1989. "Meaning and the Mental: The Problem of Semantics After Chomsky." In *Reflections on Chomsky*, edited by A. George, 23–40. Oxford: Blackwell.
Kaplan, D. 1989. "Demonstratives." In *Themes From Kaplan*, edited by J. Almog, J. Perry, and H. Wettstein, 481–563. Oxford: Oxford University Press.
Kennedy, C., and L. Mcnally. 2010. "Color, Context, and Compositionality." *Synthese* 174 (1): 79–98.
Ladyman, J., J. Lambert, and K. Wiesner. 2013. "What is a Complex System?" *European Journal for Philosophy of Science* 3 (1): 33–67.
Lepore, E., and M. Stone. 2014. *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford: Oxford University Press.
Levins, R. 1966. "The Strategy of Model Building in Population Biology." *American Scientist* 54 (4): 421–431.
Levins, R. 2006. "Strategies of Abstraction." *Biology and Philosophy* 21 (5): 741–755.
MacFarlane, J. 2014. *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.
Matthewson, J., and M. Weisberg. 2009. "The Structure of Tradeoffs in Model Building." *Synthese* 170 (1): 169–190.
McGilvray, J. A. 1998. "Meanings are Syntactically Individuated and Found in the Head." *Mind and Language* 13 (2): 225–280.
Neale, S. 2004. "This, That, and the Other." In *Descriptions and Beyond*, edited by A. Bezuidenhout and M. Reimer, 68–182. Oxford: Oxford University Press.
Pietroski, P. M. 2005. "Meaning Before Truth." In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, edited by G. Preyer and G. Peter, 255–302. Oxford: Oxford University Press.
Pietroski, P. M. 2006. "Character Before Content." In *Content and Modality: Themes From the Philosophy of Robert Stalnaker*, edited by J. J. Thomson and A. Byrne, 34–60. Oxford: Oxford University Press.
Pietroski, P. M. 2018. *Conjoining Meanings: Semantics Without Truth-Values*. Oxford: Oxford University Press.
Predelli, S. 2004. "Painted Leaves, Context, and Semantic Analysis." *Linguistics and Philosophy* 28 (3): 351–374.
Rabagliati, H., G. F. Marcus, and L. Pylkkänen. 2011. "Rules, Radical Pragmatics and Restrictions on Regular Polysemy." *Journal of Semantics* 28 (4): 485–512.
Rabern, B. 2012. "Against the Identification of Assertoric Content with Compositional Value." *Synthese* 189 (1): 75–96.
Rayo, A. 2013. "A Plea for Semantic Localism." *Noûs* 47 (4): 647–679.
Recanati, F. 2010. *Truth-Conditional Pragmatics*. Oxford: Oxford University Press.
Reiter, R. 1980. "A Logic for Default Reasoning." *Artificial Intelligence* 13 (1–2): 81–132.
Rothschild, D., and G. Segal. 2009. "Indexical Predicates." *Mind and Language* 24 (4): 467–493.
Searle, J. R. 1980. "The Background of Meaning." In *Speech Act Theory and Pragmatics*, edited by M. Bierwisch, J. Searle, and F. Kiefer, 221–232. Berlin: Springer.
Soames, S. 2010. *Philosophy of Language*. Princeton, NJ: Princeton University Press.
Sperber, D., and D. Wilson. 1995. *Relevance*. Oxford: Blackwell.

Stanley, J. 2007. *Language in Context: Selected Essays*. Oxford: Oxford University Press.
Stojnić, U., M. Stone, and E. Lepore. 2013. "Deixis (Even Without Pointing)." *Philosophical Perspectives* 27 (1): 502–525.
Strawson, P. F. 1950. "On Referring." *Mind* 59 (235): 320–344.
Strawson, P. F. 1952. *Introduction to Logical Theory*. London: Methuen.
Szabó, Z. G. 2001. "Adjectives in Context." In *Perspectives on Semantics, Pragmatics, and Discourse*, edited by R. M. Harrish and I. Kenesei, 119–146. Amsterdam: John Benjamins.
Travis, C. 1985. "On what is Strictly Speaking True." *Canadian Journal of Philosophy* 15 (2): 187–229.
Travis, C. 1996. "Meaning's Role in Truth." *Mind* 105 (419): 451–466.
Travis, C. 1997. "Pragmatics." In *A Companion to the Philosophy of Language*, edited by B. Hale and C. Wright, 87–107. Oxford: Blackwell.
Vicente, A. 2015. "The Green Leaves and the Expert: Polysemy and Truth-conditional Variability." *Lingua* 157: 54–65.
Weisberg, M. 2007. "Who is a Modeler?" *British Journal for the Philosophy of Science* 58 (2): 207–233.
Wilson, M. 2008. *Wandering Significance: An Essay on Conceptual Behavior*. Oxford: Oxford University Press.
Wimsatt, W. C. 2007a. "False Models as Means to Truer Theories." In *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*, 94–132. Cambridge, MA: Harvard University Press.
Wimsatt, W. C. 2007b. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, MA: Harvard University Press.
Yalcin, S. 2017. "Semantics as Model-based Science." In *The Science of Meaning: Essays on the Meta-Theory of Natural Language Semantics*, edited by D. Ball and B. Rabern, 334–360. Oxford: Oxford University Press.