Semantics and Truth Relative to a World*

Michael Glanzberg

University of California, Davis

My concern in this paper shall be with the idea that truth is in some way relative to a possible world. There is no doubt something right about this sort of slogan, but nonetheless, most of what I have to say about it will be negative. In particular, I shall argue that relativity of truth to a world plays no substantial role in the empirical semantics of natural language as it is standardly done. This is so, I shall argue, even when we focus on truth-conditional semantics carried out in the model-theoretic tradition. Work in this tradition relies on some formal languages whose model theory is typically done using a model-theoretic notion of truth at a world; but I shall argue, this is not an empirically significant aspect of semantic theory.

Empirical work is never done in a vacuum. Empirical work in semantics is done not only against a background of the metatheory of formal languages, but also of philosophical ideas about content. Some ways of thinking about content do, I shall grant, show us a notion of truth relative to a world. But, I shall argue, they do not do so in a way that directly impacts empirical work in semantics itself. They offer us important ways of understanding the basic notion of content, but do not indicate any data, or any other sorts of demands on empirical

*Thanks to Kent Bach, Brit Brogaard, Dave Chalmers, Andy Egan, Cody Gilmore, Carrie Jenkins, Bernhard Nickel, Agustín Rayo, Adam Sennet, Paul Teller, and members of the contextualism reading group at the ANU for helpful discussions and comments on earlier drafts of this paper. The final preparation of this paper was done while I was a program visitor at the Philosophy Program in the Research School of Social Sciences of the Australian National University. Thanks to all the people there for gracious hospitality and a wonderful research environment.

theories, which require them to make essential use of a notion of truth at a world. Thus, I shall not say that there is no role for any notion of truth at a world in thinking about language, but it is far more limited than a brief glance at model-theoretic semantics might make one think.

The plan of this paper is as follows. In section I, I shall focus on the main tool of model-theoretic semantics: a rich and strong intensional theory of types. I shall point out there that though the metatheory of intensional logic is standardly done with possible worlds, intensional type theory itself is more cautious and more flexible in its treatment of intensionality. On the basis of this, in section II, I shall argue that the notion of truth at a world does not play an empirically significant role in semantic theory, even when it is done using intensional type theory. In section III, I shall step back from the machinery of empirical semantic theory, and ask where the notion of truth at a world fits into our notion of content. There, I shall grant that at least one widespread view about the nature of content offers us a clear notion of truth at a world. Even so, I shall argue, it does not offer one which constrains the empirical domain of semantic theory in a way that makes the notion of truth at a world empirically significant.

Finally, as an application of these conclusions, I shall turn to the issue of relativism in section IV. It is tempting to think that the question of whether truth is relative to a world is significant to relativism. If truth is already supposed to be relative to a possible world, it might be thought, there is nothing remarkable about the idea that truth may be relative to more than a world. The relativity of truth to something beyond a possible world is one of the hallmarks of relativism. Thus, we seem to have, if not an argument for relativism, at least an argument that it is a less radical thesis than it is often taken to be. I shall show in section IV that no such argument is available. Once we see what role the relativity of truth to a world plays in philosophical analyses of content, and how little role it plays in empirical semantic theory, we will also see that it at least fails to motivate relativism, and can even

render relativist theses even more mysterious. This will not amount to an argument against relativism; rather, it will show that what might have seemed to be an easy road to relativism leads nowhere.

I Semantics and Intensional Logic

I claim that standard semantic theory makes no substantial use of a notion of truth at a world. To substantiate this, I shall mainly focus on the role of possible worlds, and truth at a world, in the empirical discipline of truth-conditional semantics. In particular, I shall focus on the model-theoretic variant of truth-conditional semantics, of the sort represented in such textbooks as Chierchia and McConnell-Ginet (1990) and Heim and Kratzer (1998), or the slightly older Dowty *et al.* (1981), as well as many many research papers.[1] Model-theoretic semantics may seem as if it is permeated by the notion of truth relative to a world. After all, this would seem to be the characteristic mark of the treatment of intensionality in model-theoretic semantics. More generally, if you pick up a textbook on model-theoretic semantics, especially one with a logical bent, you will see symbols like $w \models \phi$ or $[\![\phi]\!]^w = \top$, which certainly appear to be notations for truth at a world in some form. The former tells us that sentence ϕ comes out true at world w, the latter that the semantic value $[\![\phi]\!]$, evaluated at w, has value $\top$ = true. In the absence of context dependence, these amount to much the same thing. Our main task is to see what roles these notions play in the theory, both empirical and metatheoretic, and why they are not as significant as the widespread appearance of such notion might make it seem.

[1] Model-theoretic semantics is in the tradition of Montague (e.g. Montague, 1973), and is sometimes called 'Montagovian'. This points out a contrast with the 'Davidsonian' variant of truth-conditional semantics, in the tradition of Davidson (1967), and as represented in the textbook of Larson and Segal (1995). The Davidsonian approach avoids relativity to a possible world in a much more pronounced way than the model-theoretic approach, so focusing on the latter if anything makes my case stronger. In the long run, I think some of the points I shall raise here help to minimize the apparent differences between these two variants of truth-conditional semantics. I shall return to this point, very briefly, in section III.

To do this, we should start by looking at the background apparatus used by model-theoretic semantics. The background for model-theoretic semantics is a rich theory of types. This theory is rich both in its full intensional form, and its extensional fragment. To see how possible worlds and the notion of truth at a world function in semantics, it will be important to first see how they function in intensional type theory. It is striking just how specific a role they play, and where they do not play that role. The way the role of worlds is limited in intensional type theory, in turn, will help us to see how it is limited in the enterprise of semantics as a whole.

Ultimately, the types provided by type theory are used as the semantic values of expressions in natural languages. Just what counts as a semantic value, and what semantic values do, will become important in a moment. For now, we may operate with a loose idea that some appropriate object is assigned to each expression. For instance, if we focus on the extensional fragment, it is natural enough to think that semantic value of proper name or pronoun is the individual it denotes, and the semantic value of a predicate is, roughly, the set of individuals that bear it.

Type theory generalizes this idea. We begin with two basic types. The type e is the type of values of proper names and pronouns, i.e. the type of individuals. The type t is the type of truth values, which is typically the type of the semantic values of whole sentences or clauses. We then build up more complex types as functions. The type $\langle e, t\rangle$ is the type of functions from type e to type t. This is in effect the type of sets of individuals. These are taken to be the semantic values of predicates or intransitive verbs. More generally, for any types ν and τ, there is a type of functions $\langle \nu, \tau\rangle$ from ν to τ.

Included in type theory is the operation of λ-abstraction. If ϕ is an expression of type τ and x is a variable of type ν, then $\lambda x.\phi$ is of type $\langle \nu, \tau\rangle$. If ψ is of type $\langle \nu, \tau\rangle$ and γ is of type ν then $\psi(\gamma)$ is of type τ.

A great deal of work in empirical semantics, especially compositional semantics, runs at

the level of associating expressions of natural language with the right types, and working out how those types can combine to get the right types of larger phrases. To mention one well-known example, working out that the type of quantified noun phrases is $\langle\langle e, t\rangle, t\rangle$ amounted to a major advance.

Type theory, as I have been describing it, is a *formal language*, much like other logical languages. It is a term language, built upon constants and variables (terms) of various types, plus logical constants and operators applied to terms of the right types.

This formal language has a natural class of standard models. To each type, we associate a domain of values for expressions of that type. Standardly, for type t, the domain D_t is the set of truth values $\{\top, \bot\}$. (There are actually a number of other options for interpreting type t, but this one is standard.) For e, we have a domain D_e, typically taken to be a domain of individuals. Terms of type e will then have as values individuals drawn from D_e. Function types have domains of functions. For types ν and τ, $D_{\langle\nu,\tau\rangle} = D_\tau^{D_\nu}$. The collection of domains D_τ forms a standard model. Each such model is uniquely determined by the domain D_e. Thus, for a given set D, there is a standard model based on D generated by setting $D_e = D$.

So far, we have looked only at an extensional type theory. A richer, intensional theory may be built up by making a few modifications. The fundamental move is to add new derived types. For any type τ, there will be a type $\langle s, \tau\rangle$. We are accustomed to thinking of s as the type of possible worlds (or more generally, indices of some sort), and this is indeed how the model theory is standardly developed. But as far as the type theory itself is concerned, what we have added are a wide range of new types of form $\langle s, \tau\rangle$. These types are *intensions*, and the expanded theory is the intensional theory of types. One of the key features of this intensional type theory is that though it has a great many intensional types, it does not have a basic type s. The only way s appears in types is in building intensions—derived types of the form $\langle s, \tau\rangle$. There is no such thing as expressions of type s in the language of type theory.

Of course, the language of type theory must have some way of working with intensions. In the traditional Montagovian settings, intensions are handled by a pair of operators $^\wedge$ and $^\vee$. For α of type τ, $^\wedge\alpha$ is the intension of α, of type $\langle s, \tau\rangle$. If ι is an intension, of type $\langle s, \tau\rangle$, then $^\vee\iota$ is its extension, of type τ. Thus, we move around the intensional side of the theory by intension and extension-forming operators. We do not ever talk about the possible worlds (which wold be of type s) themselves.

Standard models for intensional type theory can be built using an additional set W, which intuitively, we think of as a set of possible worlds. Given sets D and W, we build a standard model as before, with the addition of the rule that expressions of type $\langle s, \tau\rangle$ have values in D_τ^W. We thus have, for a set of worlds and a set of individuals, a standard model of intensional type theory. This approach to the model theory of intensional type theory does indeed use sets W, which we are intended to think of as sets of worlds. It allows us to define notions of truth at a world. For a given model M and sentence ϕ of type t, we can ask if $w \models_M \phi$, or in other notation, if $[\![\phi]\!]_M^w = \top$. However, in keeping with the observation that there is no type s in intensional type theory, there is no domain $W = D_s$ in a model of intensional type theory. W alone is not properly part of a model.

We thus have something that looks like truth at a world, as expected. But it is worth stressing that we have it only in our metatheory. The theory of types is a formal language, in which we can formulate axiom systems, and about which we can investigate model-theoretic properties, completeness, decidability, etc. We can ask about it whatever we care to ask in the metatheory of formal languages. Doing model theory is part of that, and it is in model theory that we first encountered our notion resembling truth at a world.

I have just described the basics of an intensional type theory, known as IL, essentially due to Montague (1973).[2] The language of IL contains no expressions of the type of worlds.

[2]Montague's own presentation of this formalism includes temporal operators. Given Montague's preferred approach to tense, this is a practically trivial difference. The version sketched here is presented in detail in Gallin (1975), and reviewed in Gamut (1991).

Instead, it contains operators $^\wedge$ and $^\vee$, and terms of types $\langle s, \tau \rangle$. It contains lots of terms for intensions, and operators for working with intensions, but has no way to talk about worlds directly. It does have a natural model theory, providing models built on possible worlds, but we only encounter worlds when we turn to the metatheory of IL. We do not find them in using IL itself.[3]

The fact that this sort of intensional type theory does not have a type of worlds is not an accident. It is squarely in keeping with the tradition of intensional logic. This tradition, most prominently represented by Church (1951), takes it cue from Frege's (1892) idea that in some environments, terms refer to their senses rather than their customary referents. But in a departure from Frege's own views, intensional logic develops logics with terms for intensions, which play the role of Frege's senses, as well as terms for extensions. IL does this by having terms of types $\langle s, \tau \rangle$. It is not a basic part of this idea that intensions are functions from worlds to extensions, or otherwise involve worlds. They can just as well be understood as coarse-grained variants of Fregean senses. Because they area coarse-grained, treating them as functions from worlds to extensions does prove to be a highly useful tool in the metatheory of intensional logic, but it is not the starting point for intensional logic itself. We can certainly think about senses or intensions without thinking about possible worlds, and the intensional theory of types helps us to do just that.

This does relate to the question of temporalism: the question of whether truth conditions are time relative. In formal terms, this appears as the question of whether W should be simply a set of worlds, or of world-time pairs. There have been a number of arguments against the temporalist view that truth conditions are time-relative; from Enç (1986, 1987), from Evans (1985), from King (2003), and from Richard (1981, 1982). I am inclined to accept these arguments, and I believe some of the considerations I shall raise here offer additional reasons to be skeptical of temporalism. But, my main arguments will go through on a temporalist interpretation with 'world' simply replaced by 'world-time pair'. I shall hence systematically ignore questions of tense and temporalism here. A survey of some recent empirical work on tense can be found in Kusumoto (1999).

[3]Gallin (1975) also formulates an alternative type theory, called $Ty2$, which has a basic type s of worlds. Extending work of Gallin, Zimmermann (1989) shows that IL and $Ty2$ are equivalent. $Ty2$ is an extensional theory, which makes is convenient to work with logically. But, as we will see more fully in the next section, the intensional variant IL more directly corresponds to the structure of natural language, as it is standardly understood.

Though possible worlds models are a useful tool for the metatheory of intensional logic, they are not required. For instance, Church (1951) himself developed an axiomatic approach to intensional logic. Axiomatic formulations of systems related to IL are likewise formulated by Gallin (1975). Though we may use possible worlds to build models of intensional logic, there are many other ways to do the semantics of intensional logic as well. Algebraic and topological semantics for modal logics have been well-studied, and already Gallin (1975) applied them to logics related to IL.[4] Anderson (1984) essentially shows how to construct a term model for a form of Church's logic from equivalence classes of terms.[5]

Though we do not need possible worlds to do the semantics of intensional logic, possible worlds models do provide an excellent tool. They allow for generalized completeness and compactness results for IL, and have been instrumental in a great number of metatheoretic results in modal logic.[6] They also give us a nice intuitive way to think about intensions. Just as we do formally in possible worlds models, we may think of intensions as functions from worlds to extensions (essentially following Carnap, 1947). Possible worlds thus offer a nice metatheory for intensional logic; nice both in leading to a wide range of metatheoretic results, and in offering an intuitive picture of the semantics of intensional logic. ('Semantics' here in the sense of the semantics of a formal language—a branch of mathematics—not the empirical discipline of natural-language semantics.)

I have dwelt at length on the nature of intensional type theory and its formal semantics, as type theory is the main tool of model-theoretic semantics of natural language. I have pointed out that intensional type theory itself does not require worlds. It has no type of worlds, and thus no terms that pick out worlds. It is not the logic of possible worlds, or the logic of truth at a world. What it has, in contrast, is a rich array of intensional structures. It

[4]Some nice examples of algebraic and topological approaches to modality may be found in Goldblatt (1993).

[5]Relations between Church's and Montague's approaches to intensional logic are explored by Parsons (1982).

[6]See, for instance, Blackburn *et al.* (2001) for a survey.

is, as its name suggests, the logic of intensions. Possible worlds enter the study of intensional logic as one way to do formal semantics. It is a very good way, both technically, and for the intuitive picture it gives, but it is one among many ways to think about how intensional logic works. It is not an essential part of intensional logic. When we use intensional logic as a tool, as the semantics of natural language does, we do not need to pay very much attention to the metatheory of intensional logic. Perhaps more importantly, when we do pay attention to it, we should remember that the metatheory is a way to help us understand our tool and how it works, not an essential part of the tool's function itself.

II Semantic Values

So far, we have looked at the intensional theory of types, which is the main logical tool of model-theoretic semantics. We have seen that it has a model theory making use of the notion of truth at a world, but that this is not essential to intensional logic. Though this model theory is elegant and useful, intensional logic can be, and has been, developed in many ways not relying on possible worlds. With this background in hand, we may now turn to the empirical study of semantics of natural language, and see where possible worlds, and truth relative to a world, fits into empirical semantic theory.

Type theory is important to semantics because it provides a source of semantic values. One of the main tasks of a semantic theory is the assigning of semantic values to expressions—lexical items and phrases—of a language. This is not all there is to semantics, of course. Among things, the theory also needs to show how semantic values combine compositionally to compute values of larger phrases. More is required of a good semantic theory in the long run, but assignment of semantic values and rules for compositional derivation of semantic values comprise a great deal of the work of semantics as an empirical discipline.

In this sort of work, we see the Fregean tradition in full view. Frege held that the

sense of an expression is a mode of presentation of its referent. Referents—concepts and objects—combine compositionally to determine the referents of complex phrases. As Frege is often read, one displays the sense of an expression by giving its reference, in the right way. Contemporary semantic theory focuses on the truth conditions of sentences, rather than Frege's notion of sense. But it follows in Frege's footsteps in seeking to determine the truth conditions of sentences by showing how the referents of their parts are determined, and how they combine compositionally. Semantics typically works at the level of reference, in order to determine truth conditions compositionally.[7]

In its modern model-theoretic form, semantics pursues this program by assigning to expressions semantic values drawn from the theory of types.[8] We may, if we like, think of these semantic values as the referents of expressions. Certainly for a few clear cases, they are. Proper names, for instance, might be assigned elements of type e, i.e. assigned the individuals they refer to, as their semantic values. But semantics assigns semantic values much more widely than the intuitive notion of reference suggests. Again, following Frege, the semantic values of sentences are taken to be elements of type t, i.e. truth values. With Frege, we could say that sentences refer to their truth values, if we like. But the theory has no real need of this pronouncement. What is important is that compositionally determining the truth value of a sentence is a way to display its truth conditions. Generally, assigning the right objects from the theory of types to expressions is a way to display their semantic properties relevant to compositional semantics. In some cases, this is intuitively what we think of as reference, but semantic value is a theory-internal idea that can go beyond the intuitive notion of reference.

We must pause to remind ourselves that the semantic value of an expression is not gen-

[7]For a recent discussion of Frege's views, see Heck and May (2006).

[8]This is one of the characteristic marks of the model-theoretic approach to semantics. Davidsonian approaches can be understood as developing the Fregean idea, but they do not assign higher-type semantic values.

erally its meaning, or the content that might be asserted with the expression. For instance, in the standard approach there are only two semantic values for sentences, $\top$ and $\bot$, but clearly nobody thinks that there are only two meanings for sentences, or two contents that can be asserted. What semantic values do is enable the compositional determination of truth conditions. Truth conditions give us a representation of the contents expressed by sentences, and we think of meanings as mappings from contexts to these contents, determined by compositional semantics. Though there are many arguments about whether truth conditions, or truth conditions with the right additional structure superimposed, exhaust the notion of meaning, it is those features of meaning that the empirical study of truth-conditional semantics is after. Thus, as far as semantic theory is concerned, semantic values are not themselves meanings or contents, but theoretical tools for working out meanings.

As I have briefly described the project of semantics, there is yet no need for us to introduce intensional notions into the theory itself. We want to compute truth conditions, not make them objects mentioned by the theory. For the most part, this allows us to leave the semantic values of sentences as type t, and work with purely extensional semantic values throughout. These extensional semantic values suffice to compute the truth conditions for a wide range of natural-language phenomena, as is evident in the textbook of Heim and Kratzer (1998), for instance, but also in a large number of research papers in any semantics journal.

It is thus no accident that a huge amount of work in empirical semantics is carried out using the purely extensional fragment of type theory. In many cases of interest to semantics, compositional derivations themselves are insensitive to issues of intensionality, and so extensional semantic values suffice. Typically, intensional semantic values are introduced into model-theoretic semantics because they are needed to analyze certain intensional constructions in natural language. When we encounter intensional constructions like modal auxiliaries, intensional verbs, and a host of others, our extensional compositional semantics cannot get the truth conditions right. Intensional semantic values are needed to get the

truth conditions right for these cases.

The intensional theory of types, from which we take our semantic values, offers a number of ways to make room for intensionality. I shall start with one which is, I think, a natural starting place for philosophical concerns, and useful for discussion. I am doubtful that it is really standard in semantics, so I shall then consider a different option.

The natural move, in the face of intensional constructions in natural language, is to change the semantic value of sentences. If we need more fine-grained compositional inputs, especially in clausal complements, we may treat the semantic values of sentences as their *intensions*, of type $\langle s, t\rangle$, rather than their extensions, of type t. This is more or less what many philosophers have thought, from Carnap (1947) and C. I. Lewis (1943) to Cresswell (1973) and David Lewis (1970). Of course, we cannot just change the semantic values of sentences, as these values need to be the result of compositional derivations. But this idea does suggest a systematic way to move to intensional semantic values. At least to a first approximation, simply replace every t in a type assignment with $\langle s, t\rangle$. Type assignments are thus intensionalized. It is still necessary to work out the semantics of intensional constructions with these new-found intensional resources, and some delicate questions can arise which were not significant in the extensional setting. Intensionalizing type assignments is the beginning of intensional semantics, not the end. But the shift to an intensional setting is essentially made by this strategy.[9]

Putting all these delicate matters aside, we can draw some general conclusions. Let us write $[\![\alpha]\!]^c$ for the semantic value of an expression α in context c.[10] For a sentence ϕ, we began

[9]This approach to intensionalization is discussed by van Benthem (1991), and more extensively by Ben-Avi and Winter (2007). It is worth noting that it produces some semantic values that are not strictly speaking intensions in the sense of intensional logic, i.e. of type $\langle s, \tau\rangle$. For instance, if intransitive verbs are of type $\langle e, t\rangle$ in the extensional fragment, they become type $\langle e, \langle s, t\rangle\rangle$. This approach thus uses what Lewis (1970) calls compositional intensions, though not exactly in Lewis's way.

[10]Unfortunately, notation like $[\![\bullet]\!]$ has two different meanings: at the moment, it is being used for the semantic value of an expression of natural language, as assigned by an empirical theory, but it is also used for the value of an expression of a formal language in a model theory or other formal semantics. For the most part, context will suffice to disambiguate these. We will see in a moment a point where this ambiguity

with the idea that $[\![\phi]\!]^c$ is an element of type t. This gave us a way to set up compositional derivations of truth conditions. In light of intensional constructions in language, we have moved to an intensional setting, where the semantic value of a sentence $[\![\phi]\!]^c$ is an intension, of type $\langle s, t\rangle$ (if you like, a proposition).

Now, we can return to the moral we drew about intensional types when looking at type theory. We started with an approach to truth-conditional semantics which relies upon extensional semantic values drawn from the extensional fragment of the theory of types, and then moved to one which assigns intensional semantic values, drawn from the full intensional theory of types. But we still have not made any genuine use of the notion of truth at a world. What we have used are intensions, as provided to us by a theory like IL. As I stressed in the last section, this theory itself makes no mention of worlds, or truth at a world. Rather, it hands us a rich palate of intensional objects. We find, in doing empirical semantics, that we need to use some of these intensional objects as semantic values for expressions of natural language. We need these intensional objects, for instance, to draw enough distinctions to get a workable compositional semantics for modal constructions. But these intensional objects do not include possible worlds *per se*, or the relation of truth at a world, and we have not yet seen any reason to appeal to those notions directly.

In discussing intensional logic in the last section, I noted that IL is in a tradition of intensional logic which is wider than possible worlds semantics, and so it is no surprise that it contains no type of worlds. But here we see a striking match-up with natural language (presumably, not an accident). As far as standard semantic theory tells us, natural languages likewise contain no expressions for possible worlds, and no way to express relations like truth at a world (except by developing model theory, of course). Our theory of types and our empirical semantic theories agree on needing intensional objects, but not specifically possible worlds or truth at a world.[11]

is exploited.

[11]My claim here is similar to one of Partee and Hendriks (1997, p. 27), "Possible worlds are taken to be

Of course, as we also saw in the last section, the model theory of theories like IL can be done with possible worlds, and normally is. That is important to model theory, but it is not really any part of the empirical work of semantics. If tomorrow some other approach to the semantics of IL were to look better to us, that would be big news in logic, but it would not be more than a curiosity for the work of empirical semantics. Semantics in the model-theoretic tradition uses tools like IL, not their metatheories. I thus maintain that we have yet to see any substantial role for the notions of possible world or truth at a world in empirical semantic theory.

The strategy for adding intensional semantic values in model-theoretic semantics we have just considered makes this point vivid, as its method is to simply replace extensional types with intensional ones. This strategy is certainly common, but it is not the only one. It may not be the standard one; at least, it is not the one originally suggested by Montague, and followed by a number of authors. One of the reasons this strategy is contentious is that it has intensionality permeating type assignments. Though terms of type e are immune to intensionality, most everything else is not. Sentences, predicates, quantifiers, etc., all come out with intensional semantic values. Perhaps most importantly, the semantic value of every sentence is its intension, whether that is called for by an intensional construction in language or not. In this way, the strategy is a kind of 'generalization to worst case' strategy for intensionalization, which makes practically everything intensional, to make sure we can account for intensional constructions.

Another, more conservative, strategy is to leave semantic values extensional as much as possible. Particularly, the more conservative strategy leaves semantic values of unembedded sentences at type t. The idea is that this assignment worked well before we encountered intensional constructions, and there is no general reason to change it just because we have

ingredients of model structures and tools for the definition of intensionality, but never themselves a domain of possible interpretations of any category of expressions."

some intensional constructions where embedded clausal elements will need intensional types. Of course, a great deal of work is needed to implement this strategy in such a way as to make intensional and non-intensional constructions all work out right. But for our concerns, we can ignore these issues.[12] The more pressing point, for our concerns, is that this more conservative strategy on intensional semantic values requires us to invoke a notion suspiciously like relatively to a world. In order to keep semantic values of unembedded sentences extensional, and use intensional values here appropriate, we have to assign semantic values only relative to a world w. Where before, we have $[\![\phi]\!]^c$ of type t, we now have $[\![\phi]\!]^{w,c}$ of type t. With the presence of the w superscript, have we now gotten truth at a world into our semantic theory?

The notation makes this suggestion inviting, but, I shall argue, we have not really gotten ourselves involved with truth at a world in any important way. Here is one reason to doubt the suggestion can be on the right track. If we are more thoroughly intensional, as in the first strategy we just considered, we have no need to put the w superscript on semantic values. But surely, trying to be more conservative with intensional analyses should not by itself introduce any sort of world relativity. Something thus seems off about taking the notation to indicate a more substantial use of the notion of truth at a world.

This is a reason to be skeptical about the suggestion, but to really see what is wrong with it, we have to look more carefully at what the superscript w is doing in $[\![\phi]\!]^{w,c}$. The answer is: pretty much nothing except simplifying the notation. The only thing it does of any significance is allow us to avoid operators like $^\wedge$ that are a fixture of IL, but not necessarily so easy to work with. It does so by allowing us to think of intensions metatheoretically.

[12] In classical Montagovian forms, such as Montague's own (1973) or Dowty *et al.* (1981), this strategy makes heavy use of the intension-forming operator $^\wedge$ and the extension-forming operator $^\vee$ to make intensional and extensional constructions work out right. In, for instance, the brief discussion of Heim and Kratzer (1998), similar results are achieved with an intensional functional application rule. Because Montague often took a 'generalize to worst case' approach to extensional aspects of type assignment, there is a way in which Montague's own version of the conservative strategy is not so conservative. Every lexical category will receive an intensional semantic value, on Montague's original approach to intensionalization.

Where official IL has $^{\wedge}\alpha$, our new system allows us to write $\lambda w.[\![\alpha]\!]^{w,c}$.[13] This is another way of denoting the intension of α. It displays such intensions in terms of the *model theory* of IL, rather than using IL itself.

To emphasize this, let us glance back to the formalism of IL once again. Purists will note that expressions like $\lambda w.[\![\alpha]\!]^{w,c}$ are not well-formed terms of IL. IL contains λ-abstraction, but $[\![\alpha]\!]^{w,c}$ is a term of the metalanguage in which we do the model theory of IL, not of IL itself. Indeed, as IL does not have a type of worlds, it contains nothing corresponding even to λw. I am not advocating purism: no one ever said that empirical semantics had to be couched in the language of IL alone. But this helps to emphasize what we have done in making our notational simplification. We have jumped right to the model theory of IL, and used a combination of it and IL itself to simplify our combination of intensional and extensional semantic values. In particular, rather than simply mentioning semantic values for natural language expressions by using terms from IL, we have made the semantic value of a natural language expression the same as the model-theoretic value of a term in IL. Thus, $[\![\alpha]\!]^{w,c}$ is playing both model-theoretic (formal) and empirical semantic roles. We are seeing truth at a world when we use this notation, precisely because that is a feature of the model theory.

Using this hybrid of IL and its model-theoretic interpretation, we manage to simplify our notation, and thereby simplify our combination of intensional and extensional semantic values. In doing so, we also build into our semantic values the model-theoretic way of thinking about intensions. This is certainly a nice way to present things, and it helps to make the more conservative intensionalization strategy clear. But, I claim, it is not otherwise a significant change. The reason is that other than sometimes λ-binding it, we do not do anything with the w superscript in our semantic values. Its presence makes our empirical

[13] This is crucial, for instance, for the formulation of the intensional functional application rule in Heim and Kratzer (1998).

semantic values match up with our model theory, but that is not something we make use of in doing the empirical work of semantics. All we do with the w superscript on the empirical side is occasionally λ-bind it, when we need intensions, and otherwise we do nothing with it except carry it through our computations. Yet λ-binding w is exactly a way to get an intension. Thus, the only empirically significant effect of the w superscript is providing us with intensions. As we have seen, that is something we can easily do with IL proper, without relying on the w superscript. One of the key features of IL is that it provides us with intensions, and it does not need any notion of truth at a world to do so. Putting these two points together,we see that whether or not we make use of the w superscript, and the attendant notion of truth at a world, is of no empirical significance to semantic theory. We get can easily get the same empirical results with it or without it. Of course, the purist IL approach would rob us of our useful and suggestive notation. But this is still a matter of notational convenience. In terms of the empirical content of semantic theory, the superscript w is not significant.

My claim so far has been one about what aspects of the rich apparatus used by empirical semantic theory (especially in the model-theoretic tradition) are empirically significant. I have argued that in spite of notation which suggests to the contrary, the notion of truth at a world is not empirically significant. We make no use of it in empirical semantics that could not be readily achieved by other means the apparatus already provides us. At most, it buys us notational simplicity and elegance, and a useful way to present things. I shall close this section by making three comments on the status of this claim

First, this claim is substantially one about where empirical semantic theory stands. We did a fair bit of conceptual work as part of my argument for it, in sorting out how intensional type theory, semantic values, and intensionalization strategies work. But the argument crucially relies on the claim that in empirical work, we do nothing with the w superscript except λ-bind it, and otherwise carry it along in computations as an appendage. This is an

observation about empirical semantic theory as it is standardly done, and so contingent on the empirical facts these theories seek to describe.[14] It is possible we could find some feature of natural language the analysis of which would require us to make some really significant use of a w parameter, and then the situation would have changed. As far as I can tell, this is not where the empirical situation stands at present.

To sound one note of caution here, I should mention that some of the key issues for the empirical situation are not fully resolved. We would need reasonably complete accounts of modality, conditionals, and a number of other difficult constructions before we could be absolutely sure that no more substantial role for a w parameter might emerge. I have trouble imagining why any such role would emerge that really would be empirically significant, and so it seems to me the situation in standard semantic theory is on the right track. I have trouble imagining this, as I find it hard to see what could be needed that could not be captured by yet more operators on intensions. The conceptual work we have done here shows that such operators do not commit us to any substantial use of worlds or truth at a world *per se*, so it is hard to see what in modality could commit us to such notions. It is hard to see what in modality would undercut the non-conservative intensionalization strategy, and it thoroughly avoids any notion of truth at a world. All the same, we should not underestimate natural language's ability to surprise us. I have put a great deal of weight on the empirical situation, so I shall have to follow the empirical facts wherever they turn out to lead.

[14]Though it is not quite textbook-standard, it is sometimes argued that there are variables over worlds in natural language. This is typically offered, e.g. by Elbourne (2005), Heim (1991), and Percus (2001), as a way to resolve certain issues of the scope of definites in intensional contexts. Intensional operators, on this approach, are treated as quantifiers over worlds, which can bind world variables in definites. According to this approach, there are thus at least variables of type s (or something to that effect), though I do not know of any proposal that there are lexical items of that type. This sort of theory hence uses a type structure more like that of $Ty2$, and offers an extensional treatment of intensional operators.

Though this sort of theory weakens the claim that there is no reference to worlds in natural language, it does not support the claim that truth is relative to a world. On this approach, world variables are on par with other variables. They are either bound by operators, or have their values fixed by context. Thus, as with the non-conservative intensionalization strategy, semantic values—both intensional and extensional—are *not* assigned relative to a world at all. We are back to $[\![\alpha]\!]^c$, with no separate w superscript.

Second, my conclusion is in no way that there is anything wrong with the conservative strategy, or any other use of possible worlds to present analyses of natural language phenomena. My conclusion does indicate that there is not much at stake in whether one goes in for possible worlds analyses or not. Yet this is no reason to avoid possible worlds analyses. If anything, it tells us they can be used 'guilt-free'. If possible worlds makes an analysis clear and vivid, it may be used without fear that it could make the analysis come out empirically false, for the worlds themselves will not indicate anything of empirical significance. If putting a superscript w on semantic values serves to make something clear, so much the better.

Third, this reminds us of the immense heuristic value of possible worlds model theory. As it is sometimes suggested to me, it is hard to imagine how sophisticated analyses of modal notions, like the analysis of counterfactuals, could have been discovered without using possible worlds. I do not want to suggest there is anything wrong with those analyses, or anything less than ideal about thinking about possible worlds in developing them. My point is only about what in those analyses turns out to be empirically significant, versus what turns out to be heuristically useful parts of the metatheory of intensional logic. An analogy might be helpful here. My friends in physics tell me that like semantics, quantum mechanics can be developed with different mathematical tools, and that these tools are not mathematically equivalent, they lead to empirically equivalent formulations of quantum mechanics. Suppose this is right. It would not change if we discovered that some aspect of quantum theory was just intractable to us without the use of one of those mathematical tools. Likewise, the fact that we might find certain empirical analyses of modality intractable without appeal to model theory does not by itself show that there is any empirical significance to the notion of truth at a world.

So far, we have looked closely at some of the machinery of intensional semantics, both from a logical perspective, and an empirical one. We have come to the conclusion that there is no empirical work done in semantics by the notion of truth relative to a world. What

we need empirically is rich intensional structure. It is only in the model theory—part of the logical metatheory—that we encounter possible worlds and truth at a world. There is nothing in the languages we speak, or the empirical work of studying the semantics of these languages, which makes any significant use of the notion of world or truth at a world *per se*.

III Truth Conditions and Semantics

In the previous sections, I argued that we do not find any empirically significant role for the notion of truth at a world in the empirical semantics of natural language proper, even as it is done in the model-theoretic tradition. Yet a misgiving may remain. In section I, I was quick to gloss this sort of work in semantics as *truth-conditional* semantics. Surely in the notion of truth conditions lies something like a notion of truth at a world. Is that not enough to show that the notion of truth at a world is an important part of semantics after all?

There is something right about this observation, and I shall try in this section to articulate what I think it is. But I shall also argue here that what is right about it does not undermine the thesis argued in the last two sections, that there is nothing empirically significant about the notion of truth at a world, as it is used in model-theoretic semantics. The notion of truth conditions that clearly implicates an idea of truth at a world is part of a philosophical account of the nature of content, I shall argue, which does not sufficiently constrain the empirical domain of semantic theory to make the notion of truth at a world itself empirically significant.

To see this, let us first look again at semantics as an empirical discipline. As such, semantics is part of linguistics, and like other areas of linguistics such as syntax, its main goal is to explain aspects of speakers' linguistic competence. In particular, it seeks to explain aspects of linguistic competence related to our grasp of the meanings of expressions, and how those meanings combine to produce meaningful sentences. In doing this, semantics, like any

empirical investigation, takes itself to be responsible to a range of data revealing facts about speakers' competence, and seeks to formulate theories which well-explain that data.

What sort of data is semantics supposed to rely upon? The main data for linguistics is a range of speakers' judgments, and of special importance for semantics are judgments which reveal speakers' grasp of meaning. Typical lists of such judgments include entailment and contradiction judgments, synonymy and non-synonymy judgments, ambiguity judgments, truth-value judgments, and others.[15] Judgments of these sorts are especially important, as they carve out a domain of facts about linguistic competence specifically about meaning.

Where does the idea of truth conditions in truth-conditional semantics fit into all this? One way it fits in is simply that this kind of data turns out to be nicely explained and predicted by theories of the sort described in section II, that assign truth conditions to sentences in compositional ways. Such theories capture entailments, truth-value judgments, synonymy, etc., as relations among truth conditions, which are derived compositionally by assigning the right semantic values to constituents. One of the key virtues of truth-conditional semantics is precisely that it can approach this data with theories which makes sense of how the meanings of expressions combine compositionally, and so it also helps explain the fact, much discussed by Davidson (1967) and others, that speakers are able to understand infinitely many novel sentences. The main point of the last two sections was simply that all this work can be done with theories which make no empirically significant use of the notion of truth at a world. Extensional theories clearly do so, and I argued that the same holds for intensional theories.

The picture I have presented so far has a starkly empiricist ring to it. There is data, and there are theories that fit or do not fit the data. Period. There are two reasons to think this

[15]For a discussion of the data for semantics, see Chierchia and McConnell-Ginet (1990). Other sorts of data used in semantics include more theoretically complex judgments like scope judgments, judgments of felicity or acceptability in context which are also basic to pragmatics, and often data more properly about syntax as well. Facts about language acquisition can be important too. Good experimentalists dream up all sorts of interesting ways to get data from speakers. I have heard of attempts to use eye-tracking devices to get data about reference, for instance. However, our shorter list will be enough to give us a picture of the kind of data at issue.

may not be all there is to the story. Still on the more empirical side, we might worry that we cannot make sense of what the data really should be, much less what a good theory would be, in the absence of some more robust idea of what phenomena we are exploring. We are trying to figure out important aspects of meaning, and we cannot do this unless we have a clear understanding of the notion of meaning with which we start. This is all the more important in the messy world of language, where facts about performance and on-line processing, facts about pragmatics and communication, and facts about social norms, all tend to mix with the facts about linguistic competence we are really after. Some conceptual ground-clearing is very much in order. The goes hand in hand with a straightforward philosophical interest in the basic notions of semantics. As a foundational or philosophical matter, we would like to understand better what meaning and related notions really are.

With this in mind, let me sketch one important philosophical way of thinking about notions related to meaning. We start with the idea of the content of an assertion or a belief. The content is what is conveyed by an assertion in successful communication, or the claim about which a belief is an attitude (or at least, a core part of it). Terminology is not terribly uniform here, but let us reserve the term 'content' for this notion.

This notion of content has its home in the philosophical analysis of communication or attitudes. These are ideas about what we as speaker do with sentences. Its relation to semantics proper—a part of linguistics—is never so simple. Certainly the semantic properties of a sentence, together with features of the context in which it is asserted or entertained, have a great deal to do with what content is conveyed or believed. Just how much is a long-standing point of debate; particularly, debate about the boundaries between semantics and the pragmatics of communication.[16] I shall not worry about that issue here. Nor shall I worry about keeping track of what is merely implicated by asserting a sentence, as opposed to what is more directly determined by its semantic properties. Noting these omissions, we

[16]Some contributions to this debate include Bach (1994) and King and Stanley (2004), among many others.

may suppose that somehow, a sentence plus a context determines a (part of a) content, but leave open exactly how that happens. I think it is safe to assume that the semantic facts about a sentence—its linguistic meaning—comprise one of the key factors in this process. Thus, empirical semantics has as its goal exploring one of the key factors that determines the content of an assertions.

So far, we have roughly identified a notion of content, and in part placed the subject matter of semantics with respect to it. But a philosophically rich account of the notion of content requires more. One way to build a richer account begins by glossing the notion of content in question as a species of *information*. The content of an assertion is the information that the world is some way or another. In a slogan, content is information that tells us what the world is like. So far, however, notions like truth conditions or truth at a world have yet to appear.

These notions come into play in further analysis of the idea of information that tells us what the world is like; or rather, they come into play in one way of analyzing this idea. According to this analysis, such information is to be thought of as a way of dividing up possibilities. The information that the world is some way sorts possibilities for whether they are that way or not. Information about what the world is like is thus coded in the collection of possibilities which are compatible with that information.

On one way of understanding this analysis, the possibilities in question are just whatever sorts of things suffice to capture information when grouped together. This is a highly theory-internal notion of possibility, which makes possibilities simply the individual components that compose a content. For the moment, let us suppose that truth conditions are these collections of possibilities, where individual conditions of truth are whatever suffices to make a set of truth conditions capture a content. Thus, I shall refer to this analysis as the truth-conditional analysis of content. (We will return in section IV to how this relates to other ideas about possibility.) I shall leave open whether the truth-conditional analysis captures

all there is to content, or if contents need further fine-grained individuation over and above truth conditions. I shall thus assume that sets of truth conditions can capture at least a core part of our basic idea of content.[17]

The truth-conditional picture of content is not the only one available, and not the only one which answers to the slogan that content is information that tells us what the world is like. Our main interest here is with the truth-conditional view; but, to fix ideas, I shall very briefly sketch one other option. One other way to think about information is in terms of structures of objects and properties or relations. The information that John is six feet tall may be thought of as the structure combining the individual John and the property of being six feet tall, for instance. This view of content is basically non-truth-conditional. Though in the long run, objects and properties will have modal features, we are starting with the idea of building contents from these things themselves, not the possibilities that go with them. If we like, we may think of objects and properties as in this world, though that is really more metaphysics than we need. To give it a label, let us call this the structured content analysis of content.[18]

So far, we have looked at some philosophical ideas about the notion of content. We have also assumed that semantics will identify aspects of linguistic competence that help determine what content is conveyed by the use of a sentence (and left just how rather vague). But what does this all tell us about our main question of whether the notion of truth at a world figures substantially in semantics?

On the philosophical side, it is clear that the truth-conditional approach to content builds in a notion we can certainly call truth at a world. If we adopt the truth-conditional approach, then contents are sets of truth conditions, and so are true or false at individual conditions of

[17]This picture of content is closely associated with Stalnaker (1978, 1984). Stalnaker's view is more developed than what I have sketched in some respects. For instance, he further articulates the notion of possibility in question, and defends the view that truth conditions alone suffice to individuate contents.

[18]This sort of view is associated with structured proposition theories of content, such as that of Soames (1987). Kaplan (1989) discusses structured proposition theories, but also works with truth-conditional ideas.

truth. If we take 'world' to have meant such a condition, then our account of content shows us how truth is relative to a world. Indeed, it does so trivially. On this view, the idea that truth is relative to a world is not a substantial claim by itself; but rather, a restatement of the idea that content is thought of as truth conditions. A truth condition is just one of the conditions making up some information content. That each one makes the content true tells us no more than they are part of that content. This is not, as I understand it, a substantial claim about truth or truth-making; rather, it is just part of the analysis of content.

The structured content view of content does not automatically invoke any notion of truth at a world, and may work with objects and properties without appealing to this notion. But nonetheless, we have at least one philosophical account of content which does provide a notion of truth at a world.

This much I am happy to grant. But so far, we have only observed that a philosophical account of the notion of content provides a notion of truth at a world. This does not directly tell whether the empirical discipline of semantics relies on such a notion. Of course, a philosophical account of content is a foundational analysis of a notion important to semantics, so it may tell us a great deal about the basic nature of the subject. Semantics helps to determine content, and so understanding the notion of content will help us better grasp what we are doing when we do semantics. In particular, we expect our semantic theories to display truth conditions (in contexts, modulo some of the messy questions about the interface between semantics and pragmatics I put aside). Our philosophical accounts of content—especially the truth-conditional account, of course—offer accounts of what is being displayed by such theories. An account of content thus helps us to understand what semantics is about, and what it does.

However, the question we face is really more narrow that this. Our question is whether the truth-conditional account of content will show us something about the empirical situation with semantic theory that uncovers an empirically significant role for a notion of truth

at a world. To do so, it is not enough to give us a philosophical take on the outputs of semantic theory. The truth-conditional account would need to uncover some sort of data, or something about how our explanations should go, which would impact the construction of semantic theories themselves. Here, I shall argue, it uncovers no such thing. Our choice of philosophical account of content (at least, among the sorts we have considered), is neutral with respect to the empirical situation in semantics. In fact, neither account uncovers either data or demands on explanation that were not already evident in the stark ('overly empiricist') picture I sketched a moment ago. Hence, in particular, the notion of truth at a world which goes with the truth-conditional account of content does not turn out to have any empirical ramifications for semantic theory. From this, in conjunction with the arguments of the last two sections, we may conclude that there is no empirically significant role for the notion of truth at a world to play.

The key reason for this is that empirical semantics, regardless of the philosophical understanding of content it may combined with, is still after the facts about linguistic competence that relate to content. Nothing in the accounts of content of the kind we have glanced at indicates any data about competence beyond what our theories were already posed to uncover. To see this, let us start by noting how our accounts of content make sense of the kinds of data we have already considered. Consider truth-value judgments, for instance. In truth-conditional terms, a truth-value judgment elicits from a speaker evidence of whether they take the actual world to be among the truth conditions of a given sentence used in a context. Truth-value judgments are thus good data about how speakers' competence relates to content. Similarly, on the structured content view, a truth-value judgment elicits evidence of which objects and properties the speaker associates with the sentence in context. We can say similar things about entailment and synonymy judgments, which show where speakers find relations between contents. Presumably the truth-conditional approach has a leg up in making sense of this data, as it has ready to hand an explanation of entailment relations

between contents in terms of containment relations between sets of truth conditions. But I assume any good structured content view will provide for entailment relations as well, so this difference is not substantial.[19]

This shows us some examples of how our views of content can help us to make sense of some data for semantics. It appears that both views make sense of the data, and more generally the empirical situation for semantics as I initially described it. Does either indicate any other data or empirical demands? I do not see how they can. I just do not see what sorts of judgments from speakers could be indicate by our accounts of content beyond those about relations between contents, and how contents obtain in this world. What should be probe for in speakers about information conveyed, beyond whether it obtains or not, and what its relations are to other such pieces of information? Yet judgments about this are precisely the truth value and entailment sorts of judgments we already have. Given this is the data, I do not see how either account of content would indicate any other sorts of explanations than our semantic theories are already set up to give on the basis of this data. For many reasons, we have already asked our explanations to flow from compositional theories, built presumably with an eye towards facts about syntax and about language acquisition. But what more would our understanding of content ask of an account of semantic competence than an account of the data we have identified that results in truth conditions for sentences used in contexts? Again, this is what we already have. I thus do not see any change to the empirical situation in semantics we already described on the basis of either account of content, and especially, none on the basis of the truth-conditional account.

Admittedly, my argument here amounts to reminding ourselves what semantic theory is about, and failing to see anything additional following from our accounts of content. Let me offer one further argument. It might be thought that the truth-conditional account does

[19]Some structured content views might indicate some extra role for structure in content as well. However, any compositional semantic theory will already rely on a great deal of syntactic structure to make the theory work. So, structure will be a feature of semantics, regardless of the view of content adopted.

offer one more sort of data for semantics: direct data about truth conditions, or if you like, about 'modal profiles' of sentences. But from an empirical perspective, this is not so. On the data side, we simply cannot elicit direct judgments about truth conditions beyond truth-value judgments. These are judgments made by speakers in this world. We can extend beyond the current facts of this world slightly, by creating scenarios and querying speakers about them, or by carefully asking speakers to imagine the scenarios. Even so, the speakers and their judgments remain essentially world-bound. We cannot directly ask for truth-value judgments in other worlds than the one we are in, and I take it that there are not really any such things as direct judgments of modal profile speakers can give as a reflection of semantic competence. Whatever we may be able to glean from honed philosophical intuition, there are no such judgments that reflect the competence every normal adult human achieves with their native language.

This is not say there is nothing in semantics to provide modal profiles. We work them out, from all the data we have at our disposal. Such data includes truth-value judgments, based on both the current facts, and any other scenarios with which we can confront speakers. Importantly, we can elicit truth-value judgments about modal sentences, which may be our best evidence about modal profiles. Entailment judgments also provide useful data. But this means that we get modal profiles as the result of doing compositional semantics with the sorts of data we already considered. It is the output of semantics that gives us the truth conditions of sentences, and so, in effect, their modal profiles. Even if we take content to be truth conditional, this does not make modal profile a special source of data of its own.

I have now argued that our philosophical accounts of content do help us to understand what semantics is about, but they do not have any direct effect on the empirical domain of semantics. From this, my main claim of this section follows, that the truth-conditional account of content does not reveal any empirically significant role in semantics for the notion of truth at a world. The data and explanations we posit for semantics remain the same, inde-

pendently of which way we understand what content is. In particular, the truth-conditional account of content does not reveal any special data for semantics. It is true that the truth-conditional account of content embodies, almost trivially, a notion of truth at a world. But in not indicating any distinctive data or sort of explanation, this leaves the claims of the previous two sections unaffected. I argued in sections I and II that semantic theories—including those developed in the model-theoretic tradition—do not make any empirically significant use of the notion of truth at a world. The only way the notion of truth at a world embodied in our truth-conditional account of content could undermine this claim is by indicating something empirical—something about the data or requirements of explanation—to which these theories are responsible and that we had overlooked in those sections. We have seen it does not.

It is worth noting that the arguments of sections I and II parallel those given in this section. In sections I and II, looking at the apparatus of semantic theory, I showed that what is required for viable semantic theories is a rich structure of intensions, but it is not empirically significant whether we think of those in terms of worlds and truth at a world, or in some other way. In parallel, here I have noted that we may seek a philosophical account of content, but it does not turn out to affect the empirical situation whether we analyze content as truth conditions or in some other way. Of course, our different approaches to content will also help us to understand the notion of intension that is at work in semantics. Truth conditional views of content will go naturally with the possible worlds view of intensions, while structured content views could help fill in an alternative account of intensions. Philosophically or conceptually, these differences may well be of great importance. But, I have tried to show, they are not important to semantics as an empirical discipline.

This is not to say that our analysis of content is irrelevant. We construe semantic theories as showing us truth conditions. We do so even though they do not explicitly mention truth conditions, nor even make crucial use of the notion of truth at a world which embodies the

truth-conditional account. Thus, our semantic theories, for various empirically grounded reasons, output truth conditions, but do not fully explain what truth conditions are like or what they do. Our philosophical analysis of content then steps in, to further explain what our empirically driven theories have produced.

The picture that has emerged is one on which philosophical accounts of content are important, but their empirical ramifications for semantics are less than it might have appeared. Philosophical accounts of content certainly help us to understand what semantics is about, and what the results of semantic theories are, but not in ways that directly impact the kinds of data, or kinds of theories, the empirical work produces. In particular, both truth-conditional and non-truth-conditional views of content are consonant with the same empirical situation.

This may help to explain one of the more puzzling features of semantic theory, and its relation to the foundational study of content. Many approaches to semantics seem to have two tiers. One is where the empirical work gets done, and the other where the philosophy of content does. The most notorious version of this is Frege view, which at least on some readings (as I mentioned in section II), envisages all the empirical work going on at the level of reference, and much of the philosophy at the level of sense. When we come to Davidsonian approaches, the empirical theories are couched in terms of a theory of truth *simpliciter*—or truth value—while the philosophy of language still talks about truth conditions. I have argued that up to some notational differences, model-theoretic semantics is in the same boat. Where it appears to be talking about truth conditions directly, it is not empirically significant. It shows you truth conditions without directly mentioning them in the theory in any significant way, much as Fregean and Davidsonian approaches do.[20]

[20]This does make me suspect that the differences between Davidsonian and model-theoretic semantics is less important than it has often been supposed (as I also hinted in note 1). Of course, I have not here dealt with many of the main arguments against model-theoretic semantics, such as those of Higginbotham (1988) and Lepore (1983). The differences that I suspect are not so great are those about the foundations of the frameworks themselves. When it comes to specific tools and analyses, such as those invoking type shifting

I suggest that we find this two-tier structure because much of what is really interesting in philosophical thinking about content goes beyond what we should ask of empirical semantic theory. Whether content is deeply truth-conditional or modal, as the truth-conditional dividing possibilities view would have it, or deeply world-bound, as some structured content views might suppose, just does not seem to be something the study of linguistic competence can decide, or of which the study of linguistic competence really takes note. This is not to cast aspersions on either tier. But I think we get a better handle on semantic theorizing if we clearly see what each tier does.

Can we then say that simply as an understanding of content, truth is relative to a truth condition or to a world? Though I reject the claim that the relativity of truth to a world plays any empirically significant role in semantics, I accept this claim. I have suggested that the picture of content in terms of dividing possibilities is not philosophically trivial, and so we should consider the arguments before agreeing to it (Stalnaker, 1984 offers many). But I have no brief against this view, and point of personal fact, I find it compelling. If what we mean by the slogan that truth is relative to a world is that this picture of content is right, then the slogan seems to me correct. But it is not correct to infer from this anything about the empirical significance of the notion of truth at a world in semantics.

IV Relativism and Truth at a World

I have now reached my main conclusions: the notion of truth at a world may be a part of some philosophical views of content, but it is not an empirically significant part of semantic theory. As an application of this, I shall turn to an issue that arises in connection with some current discussions of relativism. It is a natural thought for relativists that their thesis is not a radical step, because truth is already relative to a world. Thus, any further relativism

or event decomposition, there seem to be a number of important differences, many of them presumably empirical.

of truth is only an incremental change to a situation we already accepted. It should already be clear why I find such an appeal unpersuasive. As I see only a limited role for the notion of truth at a world, I am not moved to find relativism less radical on the basis of the appeal. In this section, I shall explain more fully what I take the appeal to be, and why I find it unpersuasive. Before doing so, I should note that this section will not constitute an argument against relativism *per se*. I shall be arguing that what might have seemed a particularly easy road to relativism fails, but that will leave the real question of whether relativism is viable quite open. Seeing that the easy road is closed will help to see what is at stake in relativism, and offer an application of the work we have done in sorting out the role of truth at a world in semantics.

Relativism comes in many sorts, and cases for relativism come from many sources. Sometimes, we reach conclusions about relativism by looking at the facts in some domain. The conclusion that motion is relative, or as my physics textbook puts it, that Newton's Laws describe motion relative to an inertial frame of reference, is an uncontroversial lesson from physics. More controversial, but familiar to philosophers, is the idea that moral assessment is also relative. As it is often glossed, this may be taken on the model of relativity of motion. As there are inertial frames of reference, relative to which motion is described, there are also moral frames of reference, relative to which things come out right or wrong.[21] Many of the most radical forms of relativism implicate something about semantics or about content. They require the contents expressed by certain utterances to be true or false only relative to something like a frame of reference.[22]

Recently, a number of authors such as Egan (2007), Egan *et al.* (2005), Kölbel (2002),

[21]Such a gloss may be found, for instance, in Harman's contribution to Harman and Thomson (1996).

[22]I suspect that some classical forms of moral relativism, including Harman's, are structurally quite different from these more radical forms, including the semantic forms of relativism I shall discuss. Classical forms of moral relativism may not count as relativism by the lights of these theories. I mention moral relativism because it gives us a familiar example of the kind of frame of reference a relativist might require, even if not the structure of the relativist theories on which we will focus.

Lasersohn (2005), and MacFarlane (2003, 2005), have made the semantic aspect of relativism especially prominent. These authors have focused on semantics, both as a source of arguments for relativism, and as a way to explain and defend relativist theses. Particularly, Lasersohn, Egan, and Egan *et al.* have offered arguments that considerations of the semantics of certain expressions, like epistemic modals and predicates of personal taste, lead to relativism. For instance, Lasersohn (2005, forthcoming) argues that the semantic properties of predicates of personal taste—predicates like *tasty* or *fun*— require the semantic values of sentences to be relativized to judges who judge what is fun or tasty. These authors differ in exactly how they formulate relativism, and in what they take to be significant about it. With an eye towards the discussions of Egan and Lasersohn, I shall take the important feature of relativism for our purposes to be that it requires contents to be true or false only in virtue of two independently varying factors, one like a possible world, and the other like an individual. The individual judges what is tasty or fun (for Lasersohn), or distinguishes the center of a world (for Egan). Regardless, an individual is sufficiently different from a world to make the independence of the two factors clear. Both Egan and Lasersohn take semantic values to be likewise relativized to individuals, to account for this feature of contents.[23]

I have discussed some of these semantic arguments for relativism, skeptically, elsewhere (forthcoming). Most of these arguments are what we might call *hard* arguments for relativism. They seek to show us that some phenomenon, be it metaphysical or semantic, cannot be understood without resorting to relativism, and seek to show us how the needed form of relativism makes sense. In this respect, the recent run of semantic approaches to relativism have much in common with their more traditional predecessors (though, there are some interesting differences in both the kinds of arguments and conclusions reached by

[23]As I mentioned in note 2, I am ignoring issues of tense and time for purposes of this discussion. As I said there, I am inclined to accept a number of arguments against temporalism. Regardless, the points I am making hold on a temporalist interpretation, with 'world' replaced by 'world-time pair', as they have throughout. Thus, ignoring tense is harmless for our purposes.

semantically motivated relativists and more traditional relativists). In contrast, the appeal to truth at a world seeks to offer an easy road to relativism. At least, it seeks to find an easy motivation for relativism, and an easy plea for why it is not so radical after all. For semantic approaches to relativism, the easy road would seek to show that standard operating procedure in semantics or philosophy of language already involves a form of relativity, and then conclude that a more drastic form of relativity, such as relativity to an individual, or a moral framework, is really a modest extensions of what standard thinking about language has already established.

It is important to this easy road that it starts with what we already take ourselves to know from semantics or philosophy of language. There are hard arguments for relativism from considerations of metaphysics, and hard arguments from considerations of semantics, but the easy road is supposed to be easy precisely because it starts with standard assumptions about language. The easy road can be formulated as an argument along the following lines:

Premise: The relativity of truth to a possible world is already a standard part of semantic theory or philosophy of language.

Conclusion: It is thus only a modest step to see truth as relative to something beyond a possible world, such as an individual.

Call this the *easy road argument*.

I should hasten to say that the easy road argument has an uneasy status in the relativism literature. I know of no one who puts much weight on it. At best, it is a softening-up argument, meant to make the conclusion of some other hard arguments for relativism seem more palatable. At worst, it is an argument some relativists have discussed critically (such as MacFarlane, 2005). Even so, I think the argument is appealing enough, and enough 'in the air' to be worth understanding clearly.[24]

[24]One prominent relativist described the easy road argument to me as the one that he "makes in the bar."

The easy road argument concludes that it is a modest step to see the truth of what we say as relative to a judge, or frame of reference, etc. I have left it vague in the conclusion just what counts as a 'modest step'. I have done this because it is the premise which I think is mistaken. The premise is also equivocal, and from our earlier discussion, we can see that the equivocation is important. It might have in mind semantics as I have discussed it here: part of the empirical study of natural language. Or, it might have in mind by 'philosophy of language' the kind of theory of content we discussed in section III. Either way, I shall argue, the premise fails to yield a sound argument. On the first understanding, the premise is outright false. On the second, it is too weak to support the conclusion, without further input from a hard argument. If there is any road here, it is not easy.

On the first understanding of the premise of the easy road argument, it says that the notion of truth at a world is a standard part of the empirical semantics of natural language. The arguments of sections I and II show this to be false. There I argued that in fact, standard semantic theory makes no empirically significant use of the notion of truth at a world, even in its model-theoretic formulations.

We did see there that possible worlds are a part of one way of doing the metatheory of intensional logic. It might be argued by relativists supporting the easy road argument that the mere availability of possible-worlds-based model theory for intensional type theory is enough to show that we have a notion of truth relative to a world in semantic theory. This would pretty clearly be a mistake. It is simply not a good argument to conclude from a metatheory making use of some notion that the object theory itself does. Here is a stark example. One way of doing the metatheory of arithmetic—particularly, doing the proof theory of arithmetic, including consistency results—involves using languages with infinitely

At least as a point of motivation, something like it appears in Kölbel (2003, p. 72), "In this respect, my proposal [for relativism] is similar to the world relativity of truth proposed in possible worlds semantics; the same proposition (content) can be evaluated differently in different possible worlds. The same for perspectives; the same proposition can be evaluated differently in different perspectives."

long formulas.[25] But it would be a clear mistake to conclude from this that the language of arithmetic involves sentences of infinite length. We may use all sorts of tools in doing metatheory, but we may not conclude they are in any way part of the object theory our metatheory is about. Like the use of infinitely long formulas in proof theory, the use of truth relative to a world in model theory is simply mathematics. We start with a set of worlds W which is nothing more than an arbitrarily chosen set (or perhaps a set with some useful properties), and we *define* a notion of truth at a world. Such a mathematical exercise in metatheory is not going to reveal any substantial notion of truth at a world at work in the object theory. It is thus not able to support the premise of the easy road argument, and is of no help to the relativist.

On the second (and perhaps more likely) understanding of the premise of the easy road argument, it says that the notion of truth at a world is a standard part of the philosophy of language more broadly. I take it what this has in mind is the truth-conditional analysis of content we explored in section III. This idea is standard enough, and I granted it does imply a notion of truth at a truth condition—a notion of truth at a world. On this understanding, the premise is true enough, but I shall argue, it does not lend any support to the conclusion of the easy road argument on its own. If we grant this understanding of the premise to the relativist, the conclusion cannot hold; or at least, it cannot be seen to hold without first independently showing the coherence and correctness of relativism by a hard argument.

Understood this way, the premise of the easy road argument invokes the truth-conditional analysis of content. It notes that this analysis provides us with a notion of truth at a truth condition. It then seeks to conclude that we can extend the analysis, to see a content as true relative to something beyond a possible world or truth condition, like an individual. In discussing the truth-conditional account in section III, I was happy to identify truth conditions with possible worlds, but I also noted that the notion of truth condition used

[25]See Pohlers (1989) for a textbook presentation of this theory.

by the account is primarily a theory-internal one. Truth conditions according to it are whatever is divided up to capture content. They are whatever make contents true. This reveals a problem for the relativist. If possible worlds understood this way are what the premise of the argument is about, then the conclusion makes no sense. It makes no sense to take whatever is the right kind of thing to make a content true, and propose that we then see contents as made true only by independent contributions of that thing and something distinct from it.

I believe the right move for the relativist at this point is to invoke a different notion of possible world. An independent characterization of possible worlds might be offered, perhaps based on modal metaphysics, or perhaps on some other ideas in philosophy of language. It might then be argued that the possibilities involved in the truth-conditional account of content are not just any old things the dividing up of which represents content, but those very possible worlds. This is a non-trivial step, that does not follow automatically from the truth-conditional view of content. I'm not sure if it should be counted as a 'standard part of philosophy of language' or not, but it is frequently made.[26]

This is a step in the right direction for the relativist, but it is not enough to make the conclusion of the easy road argument sensible. The same problem we saw a moment ago reappears in slightly different form. Though we are no longer assuming that possible worlds are whatever is divided by content, it is a standard part of most pictures of possible worlds that they are *complete* ways things could have been.[27] If the worlds in question are complete ways things could be, then the conclusion of the easy road argument becomes mysterious. Just as before, it is hard to make sense of a world both being a complete way things could have been, and there being more than a world that could be needed to make a content true.

[26]This step may be a very substantial for modal realists like Lewis (1986). It is much less substantial for views of possible worlds like that of Stalnaker (1984).

[27]This assumption has been challenged, e.g. by Barwise and Perry (1983). It is a fixture of possible worlds approaches to content, from Carnap (1947) to Cresswell (1973), Kripke (1980), Lewis (1973, 1986), and Stalnaker (1984), in spite of the differing views these authors hold on the metaphysics of possible worlds.

The completeness of possible worlds makes the relativist conclusion dubious.

When we took worlds to be whatever makes a content true, the relativist conclusion looked incoherent. Once we invoke a more substantial notion of world, the problem is less drastic. *Prima facie* it is hard to see how we can have both complete possible worlds, and something more than worlds involved in making contents true. But this is simply to call for a hard argument for relativism which would show how to solve the *prima facie* problem. Such arguments might come from two sources. Perhaps a more refined understanding of possible worlds might provide such a solution,[28] or perhaps a more refined understanding of content might.[29] I suspect both might be required to make such a hard case for relativism.

I conclude the easy road is not open. If we are to make the premise of the easy road argument come out true, it only allows us to make sense of the conclusion on the basis of a hard argument for relativism. The only road is the hard road. Having reached a clearer understanding of where the notion of truth at a world fits into semantics and into thinking about content, I suggest, helps us to see why the road is so hard.

References

Anderson, C. A., 1984. General intensional logic. In D. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, vol. II, pp. 355–385. Dordrecht: Kluwer.

Bach, K., 1994. Conversational impliciture. *Mind and Language* **9**:124–162.

Barwise, J. and J. Perry, 1983. *Situations and Attitudes*. Cambridge: MIT Press.

[28]It has been suggested to me that possibilists would have more resources for doing this than actualists. I find this suggestion intriguing, and it might help show how substantial claims about the metaphysic of modality could help a relativist, but I shall not pursue it here.

[29]For instance, Andy Egan has suggested to me a strategy for developing such an account of content. First, it would argue, in opposition to the claim I made in section II, that an adequate account of the semantic values of modals would make non-trivial use of a type s parameter, i.e. of worlds and truth at a world. But furthermore, it would argue that to get the facts right, we need to see the elements of type s as not simply possible worlds, but as worlds plus individuals (or as centered worlds). This in turn would support the claim that contents are true only relative to both a world and an individual. I mentioned why I am skeptical of the first step at the end of section II, but I did grant that the issue is not fully resolved.

Ben-Avi, G. and Y. Winter, 2007. A modular approach to intensionality. In E. Puig-Waldmüller (ed.), *Proceedings of Sinn und Bedeutung 11*, pp. 76–90. Barcelona: Universitat Pompeu Fabra.

van Benthem, J., 1991. *Language in Action*. Amsterdam: North-Holland.

Blackburn, P., M. de Rijke, and Y. Venema, 2001. *Modal Logic*. Cambridge: Cambridge University Press.

Carnap, R., 1947. *Meaning and Necessity*. Chicago: University of Chicago Press.

Chierchia, G. and S. McConnell-Ginet, 1990. *Meaning and Grammar: An Introduction to Semantics*. Cambridge: MIT Press.

Church, A., 1951. A formulation of the logic of sense and denotation. In P. Henle, H. M. Kallen, and S. K. Langer (eds.), *Structure, Method, and Meaning: Essays in Honor of H. M. Sheffer*, pp. 3–24. New York: Liberal Arts Press.

Cresswell, M. J., 1973. *Logics and Languages*. London: Methuen.

Davidson, D., 1967. Truth and meaning. *Synthese* **17**:304–323. Reprinted in Davidson (1984).

———, 1984. *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.

Dowty, D. R., R. E. Wall, and S. Peters, 1981. *Introduction to Montague Semantics*. Dordrecht: Reidel.

Egan, A., 2007. Epistemic modals, relativism, and assertion. *Philosophical Studies* **133**:1–22.

Egan, A., J. Hawthorne, and B. Weatherson, 2005. Epistemic modals in context. In G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*, pp. 131–168. Oxford: Oxford University Press.

Elbourne, P. D., 2005. *Situations and Individuals*. Cambridge: MIT Press.

Enç, M., 1986. Towards a referential analysis of temporal expressions. *Linguistics and Philosophy* **9**:405–426.

———, 1987. Anchoring conditions for tense. *Linguistic Inquiry* **18**:633–657.

Evans, G., 1985. Does tense logic rest upon a mistake? In *Collected Papers*, pp. 342–362. Oxford: Oxford University Press.

Frege, G., 1892. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* **100**:25–50. References are to the translation as "On Sense and Meaning" by M. Black, reprinted in Frege (1984).

———, 1984. *Collected Papers on Mathematics, Logic, and Philosophy*. Oxford: Basil Blackwell. Edited by B. McGuinness.

Gallin, D., 1975. *Intensional and Higher-Order Modal Logic*. Amsterdam: North-Holland.

Gamut, L. T. F., 1991. *Logic, Language, and Meaning*, vol. 2: Intensional Logic and Grammar. Chicago: University of Chicago Press. 'Gamut' is a pseudonym for van Benthem, Groenendijk, de Jongh, Stokhof, and Verkuyl.

Glanzberg, M., forthcoming. Context, content, and relativism. *Philosophical Studies* .

Goldblatt, R., 1993. *Mathematics of Modality*. Stanford: CSLI Publications.

Harman, G. and J. J. Thomson, 1996. *Moral Relativism and Moral Objectivity*. Oxford: Blackwell.

Heck, R. G. and R. May, 2006. Frege's contribution to philosophy of language. In E. Lepore and B. C. Smith (eds.), *Oxford Handbook of Philosophy of Language*, pp. 3–39. Oxford: Oxford University Press.

Heim, I., 1991. Artikel und definitheit. In A. von Stechow and D. Wunderlich (eds.), *Semantics: An International Handbook of Contemporary Research*, pp. 487–535. Berlin: de Gruyter.

Heim, I. and A. Kratzer, 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.

Higginbotham, J., 1988. Contexts, models, and meanings: A note on the data of semantics. In R. Kempson (ed.), *Mental Representations: The Interface Between Language and Reality*, pp. 29–48. Cambridge: Cambridge University Press.

Kaplan, D., 1989. Demonstratives. In J. Almog, J. Perry, and H. Wettstein (eds.), *Themes From Kaplan*, pp. 481–563. Oxford: Oxford University Press. First publication of a widely circulated manuscript dated 1977.

King, J. C., 2003. Tense, modality, and semantic value. *Philosophical Perspectives* **17**:195–245.

King, J. C. and J. Stanley, 2004. Semantics, pragmatics, and the role of semantic content. In Z. G. Szabó (ed.), *Semantics vs. Pragmatics*, pp. 111–164. Oxford: Oxford University Press.

Kölbel, M., 2002. *Truth Without Objectivity*. London: Routledge.

———, 2003. Faultless disagreement. *Proceedings of the Aristotelian Society* **104**:53–73.

Kripke, S., 1980. *Naming and Necessity*. Cambridge: Harvard University Press.

Kusumoto, K., 1999. *Tense in Embedded Contexts*. Ph.D. dissertation, University of Massachusetts at Amherst.

Larson, R. and G. Segal, 1995. *Knowledge of Meaning*. Cambridge: MIT Press.

Lasersohn, P., 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* **28**:643–686.

———, forthcoming. Relative truth, speaker commitment, and control of implicit arguments. *Proceedings of the North East Linguistics Society* **37**.

Lepore, E., 1983. What model-theoretic semantics cannot do. *Synthese* **54**:167–187.

Lewis, C. I., 1943. The modes of meaning. *Philosophy and Phenomenological Research* **4**:236–250.

Lewis, D., 1970. General semantics. *Synthese* **22**:18–67. Reprinted in Lewis (1983).

———, 1973. *Counterfactuals*. Cambridge: Harvard University Press.

———, 1983. *Philosophical Papers*, vol. 1. Oxford: Oxford University Press.

———, 1986. *On the Plurality of Worlds*. Oxford: Blackwell.

MacFarlane, J., 2003. Future contingents and relative truth. *Philosophical Quarterly* **53**:321–336.

———, 2005. Making sense of relative truth. *Proceedings of the Aristotelian Society* **105**:321–339.

Montague, R., 1973. The proper treatment of quantification in ordinary english. In J. Hintikka, J. Moravcsik, and P. Suppes (eds.), *Approaches to Natural Language*, pp. 221–242. Dordrecht: Reidel. Reprinted in Montague (1974).

———, 1974. *Formal Philosophy*. New Haven: Yale University Press. Edited by R. Thomason.

Parsons, C., 1982. Intensional logic in extensional language. *Journal of Symbolic Logic* **47**:289–328.

Partee, B. H. and H. L. W. Hendriks, 1997. Montague grammar. In J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language*, pp. 5–91. Amsterdam: Elsevier.

Percus, O., 2001. Constraints on some other variables in syntax. *Natural Language Semantics* **8**:173–229.

Pohlers, W., 1989. *Proof Theory*. Berlin: Springer-Verlag.

Richard, M., 1981. Temporalism and eternalism. *Philosophical Studies* **39**:1–13.

———, 1982. Tense, propositions, and meanings. *Philosophical Studies* **41**:337–351.

Soames, S., 1987. Direct reference, propositional attitudes, and semantic content. *Philosophical Topics* **15**:47–87.

Stalnaker, R. C., 1978. Assertion. In P. Cole (ed.), *Pragmatics*, vol. 9 of *Syntax and Semantics*, pp. 315–322. New York: Academic Press. Reprinted in Stalnaker (1999).

———, 1984. *Inquiry*. Cambridge: MIT Press.

———, 1999. *Context and Content*. Oxford: Oxford University Press.

Zimmermann, T. E., 1989. Intensional logic and two-sorted type theory. *Journal of Symbolic Logic* **54**:65–77.