Australasian Journal of Logic

Logical Truth*

Paal Fjeldvig Antonsen

Abstract

On the model-theoretic account, a sentence is logically true just in case it is true on all possible semantic interpretations. We differentiate four ways one can interpret the modality 'possible' in this definition, and argue that one of these readings is not subject to the criticism levelled against the model-theoretic account by Etchemendy. By explicating the four readings we also draw some consequences for what linguistic evidence a selection of logical theories should be sensitive to.

Keywords Logical truth • truth in a model • anti-exceptionalism about logic • possible semantic interpretation • analyticity

1 Introduction

It is widely thought that logical truths constitute a distinctive subclass of truths, distinguished by some characteristic features.[1] When it comes which sentences happen to belong to this group, however, logical theories sharply disagree. There are, for instance, diverging verdicts on the following:

(1) If Bucephalus is a horse, something is a horse.

(2) If Bucephalus both is and isn't a horse, Alexander is a manatee.

(3) Everything is such that either it is a manatee or not.

According to classical logic, (1) – (3) are all examples of logical truths. But for each one of them there is another theory that will rule it out. An advocate of intuitionistic logic, for example, won't accept that (3) is logically true, even if they were to assent to a particular assertion of it. Similarly, a paraconsistent logician won't accept that (2) is logically true, although they might perfectly well agree that it is contingently so. The classical logician will, of course, readily acknowledge that the sentences above may not be logical truths in other theories, but that doesn't amount to more than a recognition of the fact that there are disparate technical notions labelled 'logical truth'. If there is genuine

*The work behind this paper was funded by the Research Council of Norway, as part of the 'Anti-exceptionalism about logic' project located at the University of Bergen, Norway.

[1] Tarski (1983, 414), (1996, 654, 672), Beall and Restall (2006, Ch. 2), Sagi (2014, 945) and Glanzberg (2015, 76) emphasize necessity and formality; Gómez-Torrente (2008, 341) necessity and aprioricity; Sanchez-Miguel (1993, 108, 112) necessity and analyticity; Hanson (1997, 365) (2002, 244) necessity, generality and aprioricity.

disagreement between logical theories concerning the status of examples such as (1) – (3), it is ultimately about which of their respective technical notions deserve to be called logical truth *proper*.

One strategy for selecting among competing logical theories is to employ a set of abductive criteria, similar to how we select for theories in other sciences.[2] The most recognisable criterion among them would be fitness with evidence. It is well known, for example, that the relativistic mechanics enjoys a better fit with experimental data — such as gravitational time dilation and delay — than classical mechanics. So, by that light it is the better theory. The general theory of relativity also offers more explanatory power by treating gravity as a consequence of spacetime curvature rather than a fundamental force. If we follow this strategy in logic, we should also compare logical theories by how well they score with respect to such theoretical virtues.[3] If classical logic comes out better than its rivals by these parameters, then (1) – (3) should be classified as logical truths. This abductive methodology for theory selection falls under the heading *logical anti-exceptionalism*.

It is worth noting that when we talk about logical theories in this way we do not merely have the formal systems in mind. As mentioned, if what we are seeking is an account of logical truth proper, we are pursuing something that goes beyond the definitions in the metalogic. It is helpful, therefore, to think of a logical theory as containing two components: a formal logic L that characterizes concepts such as 'valid in L' and 'logically true in L', and a conceptual component that contains descriptions of the logical concepts using terms not internal to L. An upshot of using the term in this extended sense will be that the theory entails that 'logically true in L' represents logical truth proper.

Anti-exceptionalism, so conceived, deviates from the common strategy of appealing to our 'pretheoretic' or 'intuitive' logical concepts.[4] We do not start with a theory-independent grasp of the logical concepts and then try to find the formal logic that most adequately encapsulates those. Instead, the conceptual component is but one part of the unit of selection. Given that the conceptual component is included in the theory we are trying to establish, the logical concepts can't be set aside as somehow 'pretheoretic'. Fleshing out a logical theory involves taking sides on substantial philosophical issues. If we aim to carry out the anti-exceptionalist programme consistently, then, we also need to garner abductive support for our preferred conceptual accounts.

The model-theoretic account has a lot initial appeal by this light. It has earned its influence by yielding results that accord with common practice, draws on linguistic facts as evidence, and integrates logic with the dominant tradition

[2] See e.g. (Priest, 2006), (Williamson, 2013), (Maddy, 2014), (Hjortland, 2017), and (Blake-Turner & Russell, 2018).

[3] Other salient one might be strength, simplicity, parsimony and unificiation. It's not obvious how any of those should be precisified in the context of logic, but I'm going to ignore that concern for the moment.

[4] See Etchemendy (1988, 91), (1990, 6), (2008, 265), Garcia-Carpintero (1996, 309), Ray (1996, 622 - 3), Sher (1996, 654), Hanson (1997, 366), (2002, 244), Beall and Restall (2000, 476), (2006, 36-7), Blanchette (2001, 120), Gómez-Torrente (1998, 340), (2008, 340), Shapiro (1998, 132-3), (2014, Ch. 2), and Sagi (2014, 945) for such appeals.

in natural language semantics. At the same time it throws some light on what evidence might settle questions about whether (1) – (3) are logical truths. In particular, the evidence will be drawn from the semantics features that truth in a model is aiming to represent. In what follows I will outline a version of the model-theoretic account that makes the relationship between logic and natural language semantics more explicit. The presented version is further motivated by its ability to avoid a central objection to the model-theoretic account due to Etchemendy (1988, 1990, 2008).

2 Possible Semantic Interpretation

The model-theoretic account isn't primarily concerned with truth simpliciter. Instead, it is centred around a relativised notion: truth in a model. Fundamental logical properties, such as logical truth and logical consequence, are defined in terms of this relativised truth predicate. In particular, a sentence ϕ is a logical truth just in case ϕ is true in every model. This programmatic statement is, of course, not very illuminating when read as a stand alone definition. For 'truth in all models' to throw some light on the concept of logical truth we need to flesh out what it means to say that something is true in a model.

The formal and conceptual components of a logical theory are both present in the model-theoretic account. The term 'true in a model' is first given a technical definition relative to some predefined language, but this must then be followed up with an explanation of what truth in a model represents. In the case of a first order extensional language, our description of what counts as a model runs along these lines:

(4) A model is a pair $\mathcal{M} = \langle \mathcal{D}, \mathcal{I} \rangle$, where $\mathcal{D}$ is a non-empty set ('domain'), and $\mathcal{I}$ is a function ('interpretation function') that:

i. maps each individual constant to a member of $\mathcal{D}$;

ii. maps each n-ary predicate P (except identity) to n-ary sequences whose elements are members of $\mathcal{D}$;

iii. maps the predicate for identity to the set $\{\langle x, x \rangle : x \in \mathcal{D}\}$.

The next step is to outline a semantic theory that describes how truth values are determined relative to a model and a variable assignment, where the latter is a function that maps each individual variable to a member of $\mathcal{D}$. This task consists of giving a recursive definition of the doubly relativized truth predicate truth of a formula A at a model $\mathcal{M}$ and a variable assignment g, written as '$\mathcal{V}^{\mathcal{M},g}(A) = 1$'. If our logic is classical, it may include the following clauses:

(5) a. $\mathcal{V}^{\mathcal{M},g}(Pt_1 \dots t_n) = 1$ iff $\langle [t_1]^{\mathcal{M},g}, \dots, [t_n]^{\mathcal{M},g} \rangle \in \mathcal{I}(F)$[5]
b. $\mathcal{V}^{\mathcal{M},g}(\neg A) = 1$ iff $\mathcal{V}^{\mathcal{M},g}(A) = 0$

[5]The extension of a term t — i.e. an individual constant or an individual variable — in a model $\mathcal{M}$ is written as '$[t]^{\mathcal{M},g}$' and defined as follows: $[t]^{\mathcal{M},g} = \mathcal{I}(t)$ if t is an individual constant and $[t]^{\mathcal{M},g} = g(t)$ if t is a variable.

c. $\mathcal{V}^{\mathcal{M},g}(A \vee B) = 1$ iff $\mathcal{V}^{\mathcal{M},g}(A) = 1$ or $\mathcal{V}^{\mathcal{M},g}(B) = 1$
d. $\mathcal{V}^{\mathcal{M},g}(\forall x A) = 1$ iff for every $d \in \mathcal{D}$, $\mathcal{V}^{\mathcal{M},g^{[d/\alpha]}}(A) = 1$

The definition of 'true in a model' is then added as a supplementation to the semantic theory. We recover it from the doubly relativized truth predicate by generalising over assignment functions:

(6) ϕ is true in a model $\mathcal{M}$ iff for any g, $\mathcal{V}^{\mathcal{M},g}(\phi) = 1$.

The notion we've defined in (6) is just the theory-internal notion classical logical truth, or a special case of $\mathcal{M}$-validity. It can make a claim of being a formal representation of logical truth proper only if (6) is complemented with an account of what truth in a model is successfully a model of.

The guiding principle for such an account is, I think, nicely summarized by Shapiro's formulation of one of the aims of model-theoretic semantics: 'The conditions for truth in a model match the ways that that meaningful (mathematical) statements get their truth values.' (Shapiro, 1997, 47)[6] Although Shapiro is focusing on mathematical statements, the point holds in general. If the model theory is intended to characterize logical truth, the models must represent something that reflects the truth conditional profiles of the sentences the formulas are stand-ins for. It is a sentence's truth conditions that determines its truth value on any assignment of semantic values to its constituent expressions. Presenting a sentence as true in a model, then, is to represent it as true on some possible semantic interpretation. This leads immediately to the general model-theoretic analysis of logical truth:

(7) **Logical Truth** A sentence ϕ is a logical truth iff ϕ is true on every possible semantic interpretation (PSI).

Logical Truth can be seen as a rephrasing of the familiar criterion that the model-theoretic account considers both sufficient and necessary: logical truths are true on every possible reinterpretation of their non-logical expressions. Although correct as it stands, Logical Truth is still too vague. It allows for several precisifications, resulting in different versions of the model-theoretic account. The contentious term that sticks out is the modality 'possible' in PSI. To remedy the situation, we have to specify some range of semantic values that count as possible for different classes of expressions.

Because the model-theoretic account analyses logical truth as truth on all PSIs, questions of logic and questions of natural language semantics are intimately linked. The range of semantic interpretations that are in fact possible depends on the resources of the language under considerations. To give an example, the classical logician assumes that truth and falsity are the only possible truth values. A paraconsistent logician, on the other hand, can take a more liberal attitude, allowing for a third truth value: both true and false. They might then say that (2) doesn't meet the standard of logical truth as it is false on the interpretation where 'Bucephalus is a horse' is both true and false, and

[6] See also (Shapiro, 1991, 6-7) and (Shapiro, 1999, 44).

'Alexander is a manatee' is false. Whether the paraconsistent or the classical logician is correct about the logical status of (2) doesn't depend on questions about horses or manatees, but on the question whether both true and false is a semantic value that a sentence could in principle have. Taking this point to a more general conclusion, Blake-Turner and Russell (2018) claim that some linguistic facts are immediately relevant for what we should consider as evidence for a logical theory. In their words,

> the key things we would have to determine in order to determine the accuracy of [a logical theory] are (i) whether [it] accurately reflect the structure of the language, (ii) whether [it] accurately reflects the possible interpretations for the simplest, non-logical, parts of the language, and (iii) whether the truth-statuses of complex sentences are indeed functions of the truth-statuses of their parts... (Blake-Turner & Russell, 2018, 7-8)

Linguistic facts of the kind mentioned here are facts we would expect to be able to read off from theories in natural language semantics to some degree. My concern in this essay, however, is not whether we should opt for classical logic or some alternative. Rather, my interest lies in explicating the kind of linguistic facts that the model-theoretic account should consider salient. So the focus will remain throughout on the general question of what it means to say that something is true in a model.

Suppose we are working with a language fragment rich enough to express (1) – (3), consisting of singular terms and sentences as basic grammatical types. What semantic values are possible for any expression of this language depends on its assigned meaning. Following Kaplan (1989), it is helpful at this point to differentiate two kinds of meaning we can associate with an expression: *content* and *character*. While characters are properties of expression types, contents are properties of their tokenings. The content of an expression is the determinate value it has relative to a particular context. The character of an expression, on the other hand, is a context-independent value. It is a rule that in an idealized sense explicates what content an expression can be used to express in different contexts. Two expressions with the same content at some context of utterance may still have different characters. That is, they may come to be associated with the same content, but by virtue of different linguistic rules. The difference between these two kinds of meanings is most apparent in the case of context dependent expressions. While an expression like 'today' expresses different contents at different occasions of use, the character remains invariant. It is, roughly, the rule that tells you to pick out the day of utterance.

Kaplan's distinction between content and character will play an important role in the rest of the paper. The reason is that the range of semantic interpretations that are *possible* depend crucially on whether we are talking about the semantic values that are possible relative to an assignment of content or the semantic values that that are possible relative to an assignment of character. In the next section, I spell out the three readings of Logical Truth that are available to us if we think of semantic interpretations in terms of contents. I then go on and agree with Etchemendy that, so understood, truth on all possible semantic interpretations does not yield a plausible account of logical truth. Finally, I aim to show how using models to represent assignments of characters instead avoid these problems.

3 The Standard Options

The familiar versions of the model-theoretic account are primarily focused on linguistic content. Furthermore, they are developed with a semantic framework in the background that makes use of the following semantic domains:

(8) **Semantic Domains**
a. $D_o = \{x : x \text{ is an object}\}$
b. $D_t = \{x : x \text{ is a truth-value}\}$
c. If α and β are semantic types, then $D_{\langle \alpha, \beta \rangle}$ is the set of all functions from α to β.

The extension of a singular term belongs to the semantic type o, one-place predicates to type $\langle o, t \rangle$, and binary connectives to type $\langle t, \langle t, t \rangle \rangle$. If this was the only part of the semantic theory we intended our models to represent, we would end up with an unsuitably limited range of available interpretations. So to achieve a wider range of possible assignments of values, we add a definition of content as functions from worlds to extensions: if α is a semantic type, then $D_{\langle w, \alpha \rangle}$ is the set of functions from the set of possible worlds W to α. Functions of schematic type $\langle w, \ldots \rangle$ represent *contents*. The noteworthy property of contents is that they determine semantic values relative to worlds. For example, the content of 'Bucephalus' is a function in $D_{\langle w, o \rangle}$ that maps every world to Bucephalus, and the content of 'is a manatee' is a function in $\mathcal{D}_{\langle w, \langle o, t \rangle \rangle}$, that maps every world to a function g in $\mathcal{D}_{\langle o, t \rangle}$, such that $g(x) = 1$ iff $x \in \{y : y \text{ is a manatee}\}$. More intuitively, we can say that the content of 'is a manatee' maps every every world to the set of manatees. We'll also say that the intended interpretation of a sentence is an assignment of the content it actually expresses evaluated relative to the actual world.

When we associate expressions with contents, we restrict the range of possible semantic values an expression could have. In fact, there are only three ways in which a semantic interpretation, besides the intended one, can now be called possible: (i) we can vary the world of evaluation, (ii) we can vary the content expressed, and (iii) we can vary both the world and the content. The reason we are limited to these three options is that that the semantic value of an expression is exhaustively determined by its content and the world at which it is evaluated.[7] Given our background framework, there are no other factors that will influence how the semantic values of constituent expressions are de-

[7] This is only true when we ignore indexical and other context dependent expressions. But context dependence won't play a role in the versions of the model-theoretic account considered in this section

termined. On the assumptions that the logical expressions are truth functional, the same holds for a sentence's truth value.

Let's go through each of the three options, starting with (i). We're assuming that a semantic interpretation assigns semantic values relative to some possible world. In general, expressions tend to be sensitive to which world they're evaluated at and so they get different values relative to different worlds. For example, the content of 'is a manatee' can be construed as picking out the set of manatees that exist in whatever world the expression is evaluated at. As long as there could have been other manatees than there actually are, the semantic value of 'is a manatee' will change as we move from one world to the next. This motivates the idea that a semantic interpretation can be called possible in the sense that we can evaluate sentences relative to different possible worlds, leading to the following precisification of Logical Truth:

(9) **Possibility 1** A semantic interpretation is 1-*possible* (1-PSI) iff it is an assignment of semantic values to the non-logical expressions that they could have had in some *possible* world given their actual meaning.

ϕ is a logical truth iff ϕ is true on every 1-PSI.

If truth in a model is a model of truth on a 1-PSI, the domain must represent the set of objects that exist in the world the 1-PSI is considering. It also follows that the interpretation function represents the assignment of semantic values relative to that world. Now, if you find the idea of worlds with only one object too odd (e.g. just Bucephauls or the number 4), you might alternatively think of 1-PSIs in terms of 'subworlds or situations that appear as parts of possible worlds' (Hanson, 1997, 388). Note also that Possibility 1 is simply a more careful way of stating what Etchemendy calls 'representational semantics' (1990, 20).[8] What I prefer about the description in Possibility 1 is that the background assumption that we are modelling contents is made explicit.

Possibility 1 requires that we make a small adjustment to our basic semantic domains. When considering different PSIs we are only supposed make variations among the nonlogical expressions. Logical expressions, on the other hand, should remain constant throughout. But given the definitions above, the quantifier $\forall$ is associated with a function that maps every world to an invariant value of type $\langle\langle o, t, \rangle, t\rangle$, namely a function that maps every function g of type $\langle o, t\rangle$ to truth just in case g maps all members of D_o to truth. But if that's the case, the domain would have to remain the same throughout. So to allow for domain variation across models, Possibility 1 will recast D_o as a world-dependent type, that is a function from worlds to sets of objects. For every world w, the set $D_o(w)$ intuitively corresponds to the objects that exists in w.

Unfortunately for Possibility 1, truth across 1-PSIs is usually considered inappropriate as a way to cash out logical concepts. The most immediate objection is that it runs together two notions that should be kept apart: logically true and necessarily true. Notice that truth on every 1-PSI is equivalent to truth

[8]For defences of Possibility 1, see (Read, 1994), (Hale, 1996), (Baumgartner, 2014), and (Zinke, 2015).

in every possible world, which is just another way of saying that being logically true amounts to the same as being necessarily true. So on Possibility 1, even sentences such as 'if Alexander is a manatee, then he is not a horse' comes out as logically true. Should one still wish to hold on to the idea truth in a model is a model of truth on a 1-PSI, one would have to introduce severe restrictions on the kinds of models that we are permitted to use. Furthermore, one would have have to give up on the idea that logical truths are formal, since the assignment of truth values on 1-PSIs depend on facts such that being a manatee and and being a horse are mutually exclusive.

The second option has, I think, a lot more initial appeal. This one comes from the observation that another way a semantic interpretation might be considered possible is that it can associate expressions with different contents than the ones they happen to express. We might, for example, consider the semantic value that 'is a manatee' would have had it expressed something else than it does, such as the property of being a horse. On that interpretation, the semantic value of 'is a manatee' would be the set of horses. This observation takes us to the second version of Logical Truth:

(10) **Possibility 2** A semantic interpretation is 2-*possible* (2-PSI) iff it is an assignment of semantic values to the non-logical expressions that they could have had in the actual world given some *possible* meaning.

ϕ is a logical truth iff ϕ is true on every 2-PSI.

Since a 2-PSI assigns semantic values relative to the actual world, domain variation across models must be explained as variations of the sets of objects that the quantifiers range over. This means, however, that the quantifiers can no longer be associated with a constant semantic value across 2-PSIs. In particular, $\forall$ and $\exists$ are read as 'every *thing*' and 'some *thing*', respectively, where 'thing' can be associated with different contents on different interpretations. On a 2-PSI where 'thing' means being a manatee, for instance, the domain will be the set of all manatees. An immediate theoretical cost of Possibility 2, then, is that the quantifiers are no longer considered as logical expressions.

Possibility 2 corresponds to what Etchemendy calls 'interpretational semantics' (1990, 20), and is the target of most of his criticism. Central to his case against the model-theoretic account is the argument that Possibility 2 makes the set of logical truths depend on extralogical assumptions. To illustrate this, Etchemendy (1990, 118) uses examples of sentences, such as (11), whose truth value depends on whether the domain is finite or infinite.

(11) $(\forall x \forall y \forall z(Rxy \land Rxz \to Rxz) \land \forall x \neg Rxx) \to \exists x \forall y \neg Ryx$

Intuitively, (11) says that any transitive and irreflexive relation has a minimal element. There are models with infinite domains where (11) is false, but (11) is true in every model with a finite domain. So (11) is logically true, according to Possibility 2, if there are no 2-PSIs where the quantifiers can be interpreted as ranging over infinitely many objects. Given the semantics of the quantifiers, they are restricted in such a way that they can only range over possible values for singular terms. Something is a possible value for a singular term on a 2-PSI, in turn, only if it is drawn from the set of actually existing objects. So a consequence of Possibility 2, Etchemendy observes, is that if (12) is true simpliciter then (11) is logically true.

(12) There are only finitely many objects.

This shows, according to Etchemendy, that truth on all 2-PSIs is not a good analysis of logical truth. A finitist should, on his view, be able to consistently hold both that (12) and that (11) is not a logical truth. But if Possibility 2 is correct this is an inconsistent position. The status of (11) as logically true is an immediate consequence of the finitist metaphysical position.

A proposed solution to this argument is to move to our third option for what it means to say that a semantic interpretation is possible. On this version of the model-theoretic account, interpretations are allowed to vary both the world of evaluation and the contents expressed simultaneously.[9]

(13) **Possibility 3** A semantic interpretation is 3-*possible* (3-PSI) iff it is an assignment of semantic values to the non-logical expressions that they could have had in some *possible* world given some *possible* meaning.

ϕ is a logical truth iff ϕ is true on every 3-PSI.

Just as Possibility 1, Possibility 3 treats the domain of the model as a representation of objects that exist in some possible (sub-)world. Since the semantic values of singular terms on a 3-PSI can be drawn from nonactual worlds, the quantifiers in (11) can be interpreted as ranging over objects that do not actually exist. So Etchemenedy's 'finist' argument doesn't through, as it no longer follows from (12) being true simpliciter that (11) is logically true. Etchemendy's finitist no longer finds themselves in an inconsistent position.

This solution doesn't offer much comfort in the long run, though. As MacFarlane (2000) points out, the model-theoretic analysis is still open to a strengthened version of Etchemendy's argument. Because even on Possibility 3, the finitist occupies an inconsistent position if they hold both the modal claim (14) and that (11) is not a logical truth.

(14) Necessarily, there are only finitely many objects.

By varying the world of evaluation, a 3-PSI allows singular terms to draw their semantic values from nonactual worlds. This means, though, that the objects ranged over by the quantifiers must still be ones that could have existed. So as long as the finitist contends that there could not have been infinitely many objects, they are committed to (11) being true on all 3-PSIs. If the original argument was sufficient to show that Possibility 2 is inadequate, little is gained by moving to Possibility 3. We are still dependent on extralogical assumptions about the existence of objects. To circumvent the argument one could, of course,

[9]See (Hanson, 1997), (Shapiro, 1998) and (Sagi, 2014).

say that the modality in play in Possibility 3 is not metaphysical, but logical.[10] If truth on a 3-PSI is understood as truth relative to some merely logically possible world, it doesn't follow from (14) that (11) is logically true. But not much recommends going down this route. Since our goal is to give an account of what it means to be logically true, we probably shouldn't be helping ourselves to other logical concepts that are perhaps even less understood. I think MacFarlane (2000, 9) gets it right when he says that '[i]f the modality is logical, then the analysis threatens to be circular, or at best unilluminating'.[11]

Possibility 1–3 are all designed with a simple extensional language in mind. The model-theoretic analysis, though, should be able to account for logical truths in languages with more expressive powers. In particular, it should include logical truths that involve modal operators. But the models we use in modal logic seems to contain features not corresponding to anything we have in a 3-PSI. To illustrate, the simplest model for a modal language is a triple $\langle \mathcal{W}, \mathcal{D}, \mathcal{I} \rangle$, where $\mathcal{W}$ is a non-empty set ('possible worlds'), $\mathcal{D}$ is non-empty set ('domain'), and $\mathcal{I}$ is a function ('interpretation function'). The only detail that matters for the current point is that the sets of possible worlds varies across models, just as the domain varied in models for the extensional language. Possibility 3 explained the domain variation as a variation in what objects that existed in different possible worlds. There is, however, no corresponding explanation available for the variation of possible worlds. Truth on a 3-PSI is truth on some assignment of contents to the non-logical expressions relative to some possible world. The set of worlds that are "possible" in some context is most likely a feature of the how the modal operators are interpreted, but that is not explained by varying the meanings of nonlogical expressions or varying the world of evaluation. So truth on all (modal) models seems unsuited to represent truth on all 3-PSIs.

4 Semantic Categories

Although I agree with Etchemendy's criticism of Possibility 1–3, I want to resist the conclusion that the model-theoretic account rests on 'a simple, conceptual mistake' (2008, 264). So in this section, I will outline a version of the model-theoretic account that do not share the flaws discussed in the previous section. The proposal outlined here is greatly indebted to Sanchez-Miguel (1993), and can be seen as a development of his ideas. The central move is to read 'possible semantic interpretation' as an assignment of characters rather than contents. This gives us a fourth option for understanding Logical Truth.

[10] (Shapiro, 1998, 147-8) and (Hanson, 1997, 83). Sher (1996, 668) introduces something she calls 'formal modality', but it is unclear how this notion doesn't just amount to the same as logical or metaphysical possibility; see (Sagi, 2014).

[11] Biting the bullet, Shapiro (1998, 151) says that '[t]he only conclusion to draw... is that model theory is not available to a finitist '. I think his pessimistic conclusion is premature, because it assumes that Possibility 1–3 are the only options for understanding 'possible semantic interpretation'.

By way of motivation, let's start with the idea that logical truths are supposed to be special cases of analytic truths. By 'analytical truth' I have in mind the traditional conception of a sentence being true by virtue of its meaning. The thought is, then, that a sentence should count as logically true just in case it is true by virtue of the meaning of the logical expressions. The analyticity conception of logical truth is one that Etchemendy himself is happy to endorse:

> I think we all have a perfectly good, pretheoretic more fuzzy or objectionable than most others pursuits, pursuits like cooking or biochemical engineering. Similarly, I have no complaints about logical truth conceived as a form of analytic truth: a logically true sentence is one that's true thanks to the meaning or semantic functioning of certain special terms, those we've traditionally singled out as the logical constants of the language. (Etchemendy, 1988, 98)

It's not difficult to see, though, that Possibility 1–3 are unsuited for making good on the claim that logical truths are a subset of the analytic truths. This is reflected in their vulnerability to Etchemendy's 'finitist' argument. As long as the logical status of sentences such as (11) depends on the number of objects that exist in the actual or some possible world, their truth values are not determined by facts concerning the meaning of its constituent expressions. It is not surprising, then, that Etchemendy comes down against the model-theoretic account. If Possibility 1–3 are indeed the only options we have for precisifying Logical Truth, Etchemedy's conclusion seems warranted. The weak link in the argument is, of course, the assumption that Possibility 1–3 exhaust the ways that a semantic interpretation can be considered possible. Because the only reason we were limited to those three options was a prior decision to represent contents, and so extensions relative to possible worlds.

When we talk about analytic truths, on the other hand, we are focusing on the expressions' characters. In fact, one motivation for Kaplan's distinction between content and character was precisely to capture the notion of meaning involved in 'true by virtue of meaning' (Kaplan, 1979, 85). Recall that the character of an expression was a function from contexts to contents. On the analyticity conception, then, a sentence should count as true on any assignment of characters to their nonlogical expressions. So instead of considering what semantic values the nonlogical expressions could have had by varying their contents, we should instead consider what values they could have had one some assignment of characters. That take us to a fourth way Logical Truth can be understood:

(15) **Possibility 4** A semantic interpretation is *4-possible* (4-PSI) iff it is an assignment of semantic values to the non-logical expressions that they could have had given their semantic category.

ϕ is a logical truth iff ϕ is true on every 4-PSI.

Notice that Possibility 4 corresponds to the analytic conception of logical truth. That is, a sentence ϕ is true on all 4-PSIs just in case ϕ is true in every context, on every reinterpretation of the characters of the non-logical expressions. When we say 'every reinterpretation', though, we mean any assignment of characters that respect the expressions' semantic categories. A semantic category is a specification of the contribution that expressions belonging to that category makes to the semantic value of sentences in which they can occur. We can think them as functions from contexts to sets of contents, representing all the ways in which an expression could contribute to the semantic value of more complex ones. In Possibility 2 and 3 we assumed a similar restriction on permissible reinterpretations of contents. We did not, for example, allow 'is a manatee' to be interpreted as 'believes that'. Similarly, a 4-PSI may reinterpret the characters associated with the nonlogical expressions only as long as it selects alternatives from within the same semantic category.

We can give a more precise description of semantic categories if we make some changes to our semantic domains. Recall that on Possibility 1 and 3 we treated D_o as a world-sensitive semantic type, a function from possible worlds to sets of objects. Moving to Possibility 4 we make a similar refinement. The difference is that we now treat the semantic domains as context-sensitive, making D_o a function from contexts to sets of objects. Similarly, we can treat D_t as a function from contexts to sets of truth values, but for simplicity we'll assume this is a constant function that maps every context to {true, false}. An expression with character f is said to belong to a semantic category g just in case for every context, the content determined by f is a member of the set of contents permitted by g (that is, for every context c, $f(c) \in g(c)$). For example, both 'is a manatee' and 'is a horse' belong to the semantic category $D_{\langle o,t \rangle}$, and so on a 4-PSI can be assigned any character in that category. The only difference between their characters is that, relative to a context, the one determines a function that maps manatees to truth, while the other maps horses to truth. This difference is irrelevant, however, with respect to a sentence's logical status.

One difference between Possibility 1–3 and 4 that is worth noting is the way they understand the domain of the model. Since Possibility 1–3 assigned contents to expressions relative to worlds, the domain naturally had to represent their extensions. In Possibility 4, on the other hand, the domain is not a representation of objects existing in the actual or some possible world. It is only sentences-in-context, or utterances, whose referential expressions have references, and characters are context independent. So the extension of expressions in actual or possible worlds is not represented by a model that represents a 4-PSI. Instead models are exclusively representing features of the expressions' meanings, abstracting away from metaphysical considerations about what objects do or do not exist in possible worlds. Possibility 4 also affords, I think, a more natural way to understand domain variation across models, namely as a representation of the context-sensitivity of quantifiers. The character associated with $\forall$ remains constant across all 4-PSIs, for example as a function that belongs to the semantic category $D_{\langle o,t, \rangle}$ mapping every context c to a function that maps all functions of type $\langle o, t \rangle$ that are true on every object in $D_o(c)$ to truth. But to represent this character we have to include that the set of objects can vary from one context to the next. We therefore explain domain variation by considering the meaning of the quantifiers, without having to include differences among possible worlds in our model-theoretic representation.

A final attractive feature of Possibility 4 is that it contains sufficient resources to handle modal languages. Recall that the problem for Possibility 3 was making sense of the variation of possible worlds across models. In Possibility 4 this is described as a feature of the context dependence of modal expressions, in agreement with the standard treatment of modals due to Kratzer (1977). On her account, the modality 'Necessarily' operates like a universal quantifier over a contextually determined set of worlds, such that ⌜Necessarily A⌝ is true just in case A is true relative to all of those worlds. We arrive at the same semantic treatment by employing our framework. To handle modal expressions we first let D_w be a function that maps contexts to sets of worlds, allowing for different sets of worlds being determined relative to different contexts. 'Necessarily' is then assigned a value of type $D_{\langle\langle w,t\rangle,\langle w,t\rangle\rangle}$ that relative to a context c determines a function f, such that for all w in $D_w(c)$, and all g of type $\langle w,t\rangle$, $f(g)(w) =$ truth, just in case $g(w') =$ truth for all worlds accessible from w. In other words, ⌜Necessarily A⌝ is true at some context just in case A is true in all the worlds determined by that context. So on Possibility 4, the variation of worlds across models can be seen as a feature of the meaning of the modal expressions. In this way, truth in a model for a modal language fits better with truth on a 4-PSI than truth on a 3-PSI.

5 Conclusion

In this essay, I have defended the model-theoretic analysis of logical truth as truth on all possible semantic interpretations. I differentiated four versions of this definition, opting for the one according to which a possible semantic interpretation is an assignment of character permitted by the non-logical expressions' semantic categories. This version of the model-theoretic account is motivated by its ability to avoid Etchemendy's 'finitist' argument and its ability to make sense of modal languages. A consequence of the analysis is that the selection of logical theories depends on facts recovered from natural language semantics about what types of semantic values are possible.

References

Baumgartner, M. (2014). Exhibiting interpretational and representational validity. *Synthese*, *191*(7), 1349–1373.

Beall, J., & Restall, G. (2000). Logical pluralism. *Australasian journal of philosophy*, *78*(4), 475–493.

Beall, J., & Restall, G. (2006). *Logical pluralism*. Oxford: Oxford University Press.

Blake-Turner, C., & Russell, G. (2018). Logical pluralism without the normativity. *Synthese*, 1–19.

Blanchette, P. (2001). Logical consequence. In *The blackwell guide to philosophical logic* (Vol. 2001, pp. 115–135). Blackwell Publishers.

Etchemendy, J. (1988). Models, semantics and logical truth. *Linguistics and Philosophy*, *11*(1), 91–106.

Etchemendy, J. (1990). *The concept of logical consequence*. Cambridge: Harvard university press.

Etchemendy, J. (2008). Reflections on consequence. In D. Patterson (Ed.), *New essays on tarski and philosophy* (pp. 263–299). Oxford: Oxford University Press.

Garcia-Carpintero, M. (1996). The model-theoretic argument: Another turn of the screw. *Erkenntnis*, *44*(3), 305–316.

Glanzberg, M. (2015). Logical consequence and natural language. In C. Caret & O. Hjortland (Eds.), *Foundations of logical consequence* (p. 71-120). Oxford: Oxford University Press.

Gómez-Torrente, M. (1998). Logical truth and tarskian logical truth. *Synthese*, *117*(3), 375–408.

Gómez-Torrente, M. (2008). Are there model-theoretic logical truths that are not logically true. In D. Patterson (Ed.), *New essays on tarski and philosophy* (p. 340).

Hale, B. (1996). Absolute necessities. *Philosophical perspectives*, *10*, 93–117.

Hanson, W. H. (1997). The concept of logical consequence. *The philosophical review*, *106*(3), 365–409.

Hanson, W. H. (2002). The formal-structural view of logical consequence: A reply to gila sher. *The Philosophical Review*, *111*(2), 243–258.

Hjortland, O. T. (2017). Anti-exceptionalism about logic. *Philosophical Studies*, 1–28.

Kaplan, D. (1979). On the logic of demonstratives. *Journal of philosophical logic*, *8*(1), 81–98.

Kaplan, D. (1989). Demonstratives. In J. P. J. Almog & H. Wettstein (Eds.), *Themes from kaplan* (pp. 481–563). New York: Oxford University Press.

Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and philosophy*, *1*(3), 337–355.

MacFarlane, J. (2000). What is modeled by truth in all models? In *Conferencia dictada en la pacific apa, albuquerque, nm* (Vol. 8).

Maddy, P. (2014). *The logical must: Wittgenstein on logic*. Oxford University Press.

Priest, G. (2006). *Doubt truth to be a liar*. Oxford University Press.

Ray, G. (1996). Logical consequence: A defense of tarski. *Journal of Philosophical Logic*, *25*(6), 617–677.

Read, S. (1994). Formal and material consequence. *Journal of Philosophical Logic*, *23*(3), 247–265.

Sagi, G. (2014). Models and logical consequence. *Journal of Philosophical Logic*, *43*(5), 943–964.

Sanchez-Miguel, M. G.-C. (1993). The grounds for the model-theoretic account of the logical properties. *Notre Dame Journal of Formal Logic*, *34*(1), 107–131.

Shapiro, S. (1991). *Foundations without foundationalism: A case for second-order logic* (Vol. 17). Clarendon Press.

Shapiro, S. (1997). *Philosophy of mathematics: Structure and ontology.* Oxford University Press on Demand.

Shapiro, S. (1998). Logical consequence: Models and modality. In M. Schirn (Ed.), *The philosophy of mathematics today* (pp. 131–156). Oxford: Oxford University Press.

Shapiro, S. (1999). Do not claim too much: Second-order logic and first-order logic. *Philosophia Mathematica*, *7*(1), 42–64.

Shapiro, S. (2014). *Varieties of logic.* Oxford: Oxford University Press.

Sher, G. (1996). Did tarski commit "tarski's fallacy"? *The Journal of Symbolic Logic*, *61*(2), 653–686.

Tarski, A. (1983). On the concept of logical consequence. *Logic, semantics, metamathematics*, *2*, 409-420.

Williamson, T. (2013). *Modal logic as metaphysics.* Oxford University Press.

Zinke, A. (2015). On exhibiting representational validity. *Synthese*, *192*(4), 1157–1171.