JSL 66 (Dec. 2001): 1701–1726

On Tarski On Models *

Timothy Bays

Abstract

This paper concerns Tarski's use of the term "model" in his 1936 paper "On the Concept of Logical Consequence." Against several of Tarski's recent defenders, I argue that Tarski employed a non-standard conception of models in that paper. Against Tarski's detractors, I argue that this non-standard conception is more philosophically plausible than it may appear. Finally, I make a few comments concerning the traditionally puzzling case of Tarski's ω-rule example.

In his 1936 paper, "On the Concept of Logical Consequence," Alfred Tarski provides an analysis of logical consequence which anticipates the modern, model-theoretic analysis of that notion. Tarski argues that a sentence X should be regarded as a logical consequence of a collection of sentences Γ just in case every model of Γ is also a model of X. Despite the seeming familiarity of this analysis, Tarski's paper poses two interpretive problems. First, it's hard to determine just what Tarski means by "model." The most natural reading of Tarski's paper commits him to a non-standard conception of models on which all models share the same domain; as a result, Tarski seems to face a collection of insurmountable philosophical and mathematical difficulties. Second, on any reading of the term "model," some of Tarski's examples of logical consequence don't seem to satisfy his own analysis. In particular, Tarski claims that ω-inferences should count as "logical"; but these inferences don't, at least on the surface, satisfy the model-theoretic condition mentioned above.

This paper attempts to resolve these interpretive problems. My principle goal is to show that Tarski's conception of models does, in fact, entail that all models share a single domain. However, I also argue that this "non-standardness" in Tarski's conception creates fewer difficulties than most previous commentators have thought. Near the end of the paper, I consider Tarski's ω-inference example. I argue that Tarski's reasons for regarding ω-inferences as logical are closely, albeit subtly, related to his fixed-domain conception of models (and, hence, lend credence to interpreting Tarski as advancing such a conception). Along the way, I show that the relationship between Tarski's views on ω-inferences and his undestanding of logical constants is both more complicated and more plausible than several influential commentators have supposed.[1]

*My work on this paper has benefitted from discussions with Mark Rubin, Gila Sher, and Joseph Almog. Special thanks are due to Tony Martin, John Etchemendy and an anonymous *JSL* reviewer for helpful comments on earlier drafts of the paper.

[1]Over the past decade, the literature on Tarski's conception of logical consequence has proliferated wildly. In the late 1980's, John Etchemendy published a series of seminal pieces claiming that Tarski (wrongly) advanced a fixed-domain conception of logical consequence (see, [3] and [4]). Since then, Gila Sher, Mario Gómez-Torrente, and Greg Rey have all argued—largely on grounds of charity—that Tarski *could not* have intended such a conception (see [6], [7], [11], [12], and [13]). I disagree with both of these positions: Tarski clearly *did* have a fixed-domain conception of logical consequence, but this conception causes none of the problems which Etchemendy and his critics think it causes. I discuss Etchemendy's views in sections 2, 4 and 5; I discuss his critics' views in section 3.

1 Background

Before examining Tarski's definition of "model," I want to sketch some of the material which precedes this definition in Tarski's paper. In particular, I want to examine two analyses of logical consequence which Tarski rejects and to explain his reasons for rejecting them. In doing so, I will lay the groundwork for several arguments concerning the interpretation of Tarski's definition of "model."

On the whole, Tarski's paper is concerned with providing a mathematical analysis of the notion of logical consequence. In theory, this project should involve two different kinds of investigation: first, a philosophical investigation which clarifies our intuitive conception of logical consequence, and second, a mathematical investigation which develops a formal definition corresponding to our intuitive conception. Tarski's paper, however, tends to proceed as though the first of these investigations has already been completed; Tarski simply runs through a series of candidate definitions, eliminating those which conflict with his intuitions concerning "what counts" as an instance of logical consequence. His final analysis, therefore, has a provisional character: although it does not conflict with the particular intuitions considered in earlier portions of the paper, and although it seems to "agree quite well with common usage," it may have hidden inadequacies which would be revealed if we considered additional intuitions.

The first definition which Tarski considers ties logical consequence to derivability in some formalized deductive theory (more precisely, this "derivational" analysis provides a whole collection of definitions corresponding to the collection of possible deductive theories). Tarski gives two reasons for rejecting this type of definition. First, he considers the case in which the deductive theories in question are limited to using "the normal rules of inference" ([21], p. 410). In this case, Tarski provides direct counterexamples to the claim that logical consequence is captured by derivability:

> Some years ago I gave a quite elementary example of a theory which shows the following peculiarity: among its theorems there occur such sentences as:
>
> A_0. 0 possesses the given property P,
> A_1. 1 possesses the given property P,
>
> and, in general, all particular sentences of the form
>
> A_n. n possesses the given property P,
>
> where 'n' represents any symbol which denotes a natural number in a given (e.g. decimal) number system. On the other hand the universal sentence:
>
> A. Every natural number possesses the given property P,
>
> cannot be proved on the basis of the theory in question by means of the normal rules of inference. . . . Yet intuitively it seems certain that the universal sentence A follows in the usual sense from the totality of particular sentences A_0, A_1, . . . , A_n, ... Provided all these sentences are true, the sentence A must also be true. ([21], p. 410–11)

Thus, any "derivational" analysis of logical consequence must allow inference rules which go beyond "the ordinary rules of inference."

Second, Tarski argues that adding new inference rules to a derivational system will not allow us to evade the sort of problems mentioned in the last paragraph. Citing Gödel's incompleteness theorem, he argues that

> in every theory (apart from theories of a particularly elementary nature), however much we supplement the ordinary rules of inference by new, purely structural rules, it is possible to construct sentences which follow, in the usual sense, from the theorems of this theory, but which nevertheless cannot be proved in this theory on the basis of the accepted rules of inference. ([21], 412)

Thus, no matter how we modify our conception of derivation, we will be unable to construct a derivational system which completely captures our intuitive understanding of logical consequence.[2]

In his second analysis of the notion of logical consequence, Tarski abandons the project of grounding logic on derivation. Instead, he argues that if a sentence X follows logically from a collection of sentences K, then this "consequence relation cannot be affected by replacing the designations of the objects referred to in these sentences by the designations of any other objects" ([21], p. 415).[3] As a result, he proposes that we determine the bounds of logical consequence by checking to see which consequences retain their validity under the systematic substitution of new terms for all "non-logical" terms in the sentences expressing the original consequence. Consequences which can withstand such substitution will count as logical; consequences which cannot withstand such substitution will count as non-logical. Tarski formulates this analysis as follows:

> (F) If, in the sentences of the class K and in the sentence X, the constants—apart from purely logical constants—are replaced by any other constants (like signs being everywhere replaced by like signs), and if we denote the class of sentences thus obtained from K by 'K'', and the sentence obtained from X by 'X'', then the sentence X' must be true provided only that all the sentences of the class K' are true. ([21], p. 415)

On Tarski's analysis, this condition is necessary for a particular consequence to be counted as an instance of logical consequence: i.e. for X to be a logical consequence of K, condition (F) must hold between K and X.

Tarski argues, however, that although satisfying condition (F) is necessary for a particular consequence to count as logical, it is not sufficient. In particular, the condition fails to adequately capture the intuition that logical consequence "cannot be affected by replacing the designations of the objects referred to in [the sentences in question] by the designations of any other objects" ([21], p. 415). The problem here stems from the way condition (F) uses language. The intuition would have us check to see whether *any* systematic replacement of objects could affect the status of the consequence in question; condition (F) only checks for replacements which can be captured using the expressive resources of the language in which we are working.

To illustrate this problem, consider the following inference:

[2]Here, Tarski is making the natural assumption that any reasonable extension of the "ordinary rules of inference" will be recursive (this is, perhaps, implicit in the phrase "purely structural rules"). Without this assumption, Gödel's theorem is inapplicable.

[3]It is important to notice that Tarski is working in some variant of second (or higher) order logic. In particular, the "objects" referred to in this passage include the properties and relations "designated" by the predicates of the language in question. Throughout this paper, I will follow Tarski in this use of "object."

	1.	Tim is a philosopher.	(A)
So,	2.	Tim is a graduate student.	

Intuitively, this inference should come out as an instance of non-logical consequence. If we replace the designation of "Tim" with Alfred Tarski (rather than Tim Bays), then the premise winds up true while the conclusion winds up false. However, if we are working with a particularly impoverished language—say, one which contains only the name "Tim" and the predicates "is a philosopher" and "is a graduate student"—then the inference will satisfy Tarski's condition (F). For, in this language, there is but one way to systematically substitute new "constants" into (A). Doing so gives the following:

	1.	Tim is a graduate student.	(B)
So,	2.	Tim is a philosopher.	

And, in this inference, both the premise and the conclusion are, once again, true. So, the mere fact that the language in which we formulate our inference cannot "get ahold" of facts like "Tarski had tenure" forces condition (F) to incorrectly characterize the inference as logical.[4]

Unfortunately, this kind of simple example cannot show that there is any *fundamental* problem with condition (F). If we generalize (F) by allowing complex terms to be substituted for non-logical constants, then trivial examples of the sort just mentioned can be avoided. So, for instance, the inference in example (A) would have the following as a substitution instance:

	1.	Tim is a philosopher.	(C)
So,	2.	Tim is not a graduate student.	

Since this inference tries to move from a true premise to a false conclusion, it is invalid; and on our new reading of condition (F), this invalidity ensures that the inference in (A) does not count as logical.[5]

However, this modification of condition (F) will not save the condition from more sophisticated counter-examples. Suppose that our language includes only two non-logical constants: the name "Tim" which designates Tim Bays and the unary predicate "is a philosopher." Using this language, we will be unable to define any subsets of the universe which contain exactly two elements.[6] As a result, condition (F) will wind up classifying the following as a logical inference:

[4]For a discussion of similar examples, see [11] or [13].

[5]I am grateful to Tony Martin for noticing this alternate reading of condition (F) and for pointing out the (resulting) inadequacy of my initial counter-example to that condition.

[6]This will be true even if the logical apparatus we have at our disposal is fairly sophisticated (certainly it will be true in second-order logic or in any sublanguage of $\mathcal{L}_{\infty,\infty}$). The argument for this claim is a relatively straightforward application of the fact that, in any standard system of logic, every definable subset of a model is fixed (setwise) by any automorphism of the model.

Using this fact, we argue as follows. Let X be any two-element subset of the universe; then we know that X must contain some object other than Tim Bays. Call this object Sam, and consider permutations of the universe which switch Sam with some object outside of X, while leaving everything else fixed. If Sam is a philosopher, for instance, consider a permutation which switches Sam with Alfred Tarski, Kurt Gödel, or Alonzo Church (one of whom must live outside of X). If Sam is not a philosopher, consider a permutation which switches Sam with one of three randomly chosen coffee cups (again, one of which

1. Tim is a philosopher. (D)
2. Tim is not the only philosopher.

So, 3. There are at least three philosophers.

Once again, (F) is forced to count this inference as logical, only because the language in question is so impoverished. If, for instance, the language were expanded to include a name for Alfred Tarski, then we could define the two-element set {Tim, Alfred}; and by substituting this set for the one picked out by "is a philosopher"—i.e. by substituting it throughout (D)—we would obtain an invalid inference.

Similarly, suppose that there happen to be two distinct kinds of things in the universe: material things and spiritual things. Suppose also that our language contains only two non-logical constants: the unary predicate "is material" and the unary predicate "is spiritual." Then condition (F)—even when modified to allow complex terms to be substituted for non-logical constants—will classify the following as a logical inference:[7]

1. Something is material but not spiritual. (E)
2. Something is spiritual but not material.

So, 3. Everything is either material or spiritual, but nothing is both material and spiritual.

Again, this result follows only because the language in question is so impoverished. If the language were expanded to include predicates which partially overlap—say, "is a philosopher" and "is a graduate student"—then neither of the above inferences would count as logical (i.e., neither would satisfy condition (F)). As a result, then, even our newly generalized formulation of (F) fails to capture the intuitive notion of logical consequence. What's more, it fails for the same reasons that Tarski's original formulation failed: given a sufficiently weak language, (F) can't get ahold of enough sets to recognize the non-logicality of certain inferences.

Tarski argues that we can only save condition (F)—i.e., we can only ensure that whenever an inference satisfies condition (F), that inference is also a logical inference—by assuming that "the designations of all possible objects [occur] in the language in question" ([21], p. 416). Since this assumption is unrealistic, we must find an alternative method of ensuring that our test for logical consequence takes into account all possible ways of "replacing the designations of the objects referred to in [the sentences in question] by the designations of any other objects" ([21], p. 415). And this alternative method will, of necessity, move us beyond the realm of linguistic substitution.[8]

must live outside of X). In either case, the resulting permutation gives an automorphism of the universe which does not fix X. Hence, by the fact noted above, X cannot be a definable subset of the universe.

[7]The argument for this claim is (another) straightforward application of the fact about automorphisms mentioned in the last footnote. Using that fact, we show that the only definable subsets of the universe are the set of material things, the set of spiritual things, the empty set, and the whole universe. From here, the above claim follows directly. As in the last example, this argument works even when our background logic is fairly sophisticated.

[8]This argument is probably mistaken. If the universe is infinite, then there are many languages which give the same "results" vis-a-vis logical consequence as do those languages in which there is a name for "every possible object." Proving this simply

This brings us (finally) to Tarski's definition of a model. Tarski proposes that, instead of permuting the bits of language which make up (the "non-logical" portion of) our sentences, we simply permute the objects to which our sentences refer. To accomplish this, we employ the technical notion of a "model":

> Let L be any class of sentences. We replace all extra-logical constants which occur in the sentences belonging to L by corresponding variables, like constants being replaced by like variables, and unlike by unlike. In this way we obtain a class L' of sentential functions. An arbitrary sequence of objects which satisfies every sentential function of the class L' will be called a *model* or *realization of the class L of sentences* (in just this sense one usually speaks of models of an axiom system of a deductive theory). ([21], pp. 416–17)

So, in order to build a model for a collection of sentences, we must do two things. First, we replace the original terms of our sentences with variables. Second, we look to find "sequences" of objects which satisfy the resulting sentential functions.

This definition gives Tarski the tools to ensure that all possible permutations of objects are accounted for when we assess the logicality of a particular inference. Using the notion of model, Tarski defines logical consequence as follows:

> the sentence X follows logically from the sentences of the class K if and only if every model of the class K is also a model of the sentence X. ([21], p. 417)

Since there will always be a model which "represents" any particular permutation of objects, this notion of logical consequence satisfies the intuitive requirement that "the consequence relation cannot be affected by replacing the designations of the objects referred to in these sentences by the designations of any other objects" ([21], p. 415).[9]

involves coding "possible objects" as individuals in some two-sorted extension of the universe and then applying the downward Löwenheim-Skolem theorem. Of course, this is simply a mathematical trick: it gives the right "results," but it doesn't capture the informal intuition behind Tarski's insistence that tests for logicality should examine "all possible" substitutions. Hence, to the extent that this intuition is really the motivating force behind Tarski's analysis of logical consequence, the analysis is (at least roughly) on the right track.

[9]Before moving on, we should note that Tarski's definition of "model" differs from modern definitions in ways that have little to do with the question of fixed domains versus variable domains. On the modern definition of "model," we keep the original language of our sentences—i.e. we do not follow Tarski in replacing terms with variables—but we regard this original language as *uninterpreted*. When we move from model to model, we simply vary the interpretation of our otherwise uninterpreted language. Mathematically, there's not much of a difference here, but from a philosophical standpoint, the two techniques reflect a different understanding of the relationship between languages and models.

On Tarski's analysis, the significance of our language does not change when we pass from model to model. In particular, to say that M is a model for the sentence X is not to say that X is "true in M." At most, it says that a completely different (though syntactically related) *sentential function* X' is satisfied by M. And this, even on the surface, seems philosophically innocuous.

On the modern notion of model, the significance of our language does change as we pass from model to model (assigning new sets to our predicates along the way). Hence, we (mathematicians) are forced to employ locutions like "X is true in M," and we (philosophers) are faced with the task of sorting out the significance of such talk.

2 Tarski's Models

I turn now to the question of whether Tarski's definition of "model" entails that all models share a single domain—i.e. share the entirety of the actual world as their domain. The difficulty here is caused by the brevity with which Tarski introduces his notion of a model. In particular, consider Tarski's phrase: "an arbitrary sequence of objects which satisfies every sentential function of the class L'." Does this simply mean a sequence of "objects" which correspond to the argument places of the sentential functions in the class L', or does it also include an "object" which makes up the domain on which those sentential functions are to be evaluated? In the former case, Tarski would have a "fixed-domain" account of models; in the latter case, Tarski would have a "variable-domain" account.

It seems to me that there are three reasons for thinking that Tarski intended to present a fixed-domain conception of models in his paper. First, Tarski never explicitly says that he is adopting a variable-domain conception of models. Of course, he never says that he is adopting a fixed-domain conception either; but in this particular case, the absence of explicit commentary on the issue has to favor the fixed-domain interpretation. This is because the introduction of a variable-domain conception of models *should* involve the description of a fair bit of technical apparatus. To give a variable-domain definition of "model," we must introduce conventions regulating the interaction of the objects corresponding to the variables of our sentential functions: the objects corresponding to individual variables must live in the object constituting the model's domain, the objects corresponding to predicate variables must be subsets of the model's domain, etc. The fact that Tarski, who usually takes great pains to be technically precise, fails to describe any of this technical apparatus seems to indicate that he does not consider the apparatus necessary. And this, in turn, seems to indicate that Tarski intends to introduce a fixed-domain conception of models.

The second reason for adopting a fixed-domain reading of Tarski's definition is the fact that this reading seems to fit better with the *sequence* of arguments in Tarski's paper. At the point at which Tarski introduces his definition of "model," he is trying to avoid a particular problem—i.e. the problem concerning artificially impoverished languages which we examined above. And this problem is solved simply by adopting the apparatus of sentential functions and fixed-domain models. The introduction of variable-domain models does nothing to help solve this problem, or any other problem with which Tarski is concerned in his paper.

In connection with this point, note that a fixed-domain conception of models would give us an analysis of logical consequence which is mathematically equivalent to the analysis given by Tarski's condition (F) *under the assumption that our language contains terms referring to all the objects in the world*. In contrast, a variable-domain conception of models always diverges from the analysis given by condition (F). So, for instance, both the fixed-domain conception of models and the analysis given by condition (F) (for any language, no mater how rich) will count the following as an instance of logical consequence:

	1.	Something exists.
So,	2.	Two things exist.

A variable-domain conception of models, on the other hand, will not count this as an instance of logical consequence. Thus, if Tarski thinks that condition (F) comes close to giving the correct account of logical consequence—that is, if his only worry about condition (F) concerns the explicitly mentioned problem with linguistic poverty—then we have to conclude that Tarski intends to be introducing a fixed-domain definition of "model."

On the surface, Tarski does seem to think just this. He argues, for instance, that his account of models provides a "means of expressing the intentions of the condition (F) which [is] completely independent of fictitious assumptions [in particular, the assumption that our language contains designations for every object]" ([21], p. 416). However, the defender of a variable-domain reading of Tarski has some wriggle room here. Tarski's only explicit comment concerning the adequacy of a rich-language version of condition (F) reads as follows: "the condition (F) could be regarded as sufficient for the sentence X to follow from the class K only if the designations of all possible objects occurred in the language in question" ([21], p. 416). To conclusively prove that Tarski intended to endorse the sufficiency of a rich-language version of condition (F), we would need to have the "only if" in this passage replaced by an "if and only if." And while it seems plausible to think that Tarski would have endorsed such a reading, the text itself provides no conclusive confirmation of this.[10]

The third, and in my view most compelling, reason for adopting a fixed-domain reading of Tarski's definition involves Tarski's analysis of the difference between logical and material consequence. In the section of "On the Concept of Logical Consequence" which immediately follows his definition of "model," Tarski discusses the division of our language into "logical constants" and "extra-logical terms." In the course of this discussion, he considers the impact of ignoring this division and classifying all of the terms in our language as "logical constants":

> In the extreme case we could regard all terms of the language as logical. The concept of formal consequence would then coincide with that of material consequence. The sentence X would in this case follow from the class K of sentences if either X were true or at least one sentence of the class K were false. ([21], p. 419)

In examining this passage, we find that Tarski's analysis fits well with the fixed-domain conception of models (and its corresponding notion of logical consequence), but that it fits ill with a variable-domain conception of models (and *its* corresponding notion of logical consequence).

First, suppose that we accept the assumption that all terms in our language count as logical constants, and suppose that we are working with a fixed-domain conception of models. To form a model for a particular sentence X, we first associate to X the sentential function X' which is formed by replacing all the non-logical constants in X with variables. As there are no non-logical constants in X, this sentential function is simply equivalent to the sentence X itself. In particular, since X' has no variables, its truth value cannot depend on the particular "sequence of objects" at which we choose to evaluate it. We obtain, therefore, the following

[10] The German version of the paper does not help with this problem, as it also uses an "A is a necessary condition for B to be a sufficient condition . . . " type of construction.

equivalences:

$$\begin{aligned} X \text{ is true} &\iff \text{some sequence of objects is a model of } X. \\ &\iff \text{every sequence of objects is a model of } X. \end{aligned}$$

When these equivalences are filtered through the definition of logical consequence, we obtain the further equivalences:

$$\begin{aligned} X \text{ follows logically from } K &\iff \text{every model of } K \text{ is a model of } X. \\ &\iff \text{every sequence of objects is a model of } X \\ &\qquad \text{or no sequence of objects is a model of } K. \\ &\iff X \text{ is true or some sentence in } K \text{ is false.} \end{aligned}$$

Thus, on a fixed-domain conception of models, the assumption that all terms in our language count as logical constants is sufficient to ensure that logical consequence coincides with material consequence. And this is just what Tarski claims.

Second, suppose that we (once again) accept the assumption that all terms in our language count as logical constants, but suppose now that we are working with a variable-domain conception of model. On these assumptions, logical and material consequence no longer coincide. To see this, let 'n' represent the number of grains of sand in my friend Nicholas' sandbox and consider the following argument:

1. $\exists x$ (x is a grain of sand in young Nicholas' sandbox).

So, 2. $\exists_n x$ (x is a grain of sand in young Nicholas' sandbox).

Because both the premise and the conclusion of this argument are true, the inference counts as an instance of material consequence. However, the inference fails Tarski's test for counting as an instance of logical consequence. For, by varying the domains of our models, we can construct a model containing all of the objects in the world except one of the grains of sand in Nicholas' sandbox. This model will satisfy the premise of the above argument but not the conclusion; hence, the above inference will not be an instance of logical consequence. Since this reveals a difference in the extensions of logical and material consequence, and thereby contradicts Tarski's explicit claims, it should lead us to conclude that Tarski is not advancing a variable-domain conception of models.

3 Objections and Replies

The claim that Tarski was advancing a fixed-domain conception of models in [21] was originally advanced by Etchemendy in [3], and the arguments of the last section seem to support Etchemendy's claim.[11] Recently, however, this claim has come under sharp criticism in the literature. In this section, I examine two of the

[11] Etchemendy himself presents few arguments for reading Tarski the way he does, since his project is more analytical than interpretive. We can surmise, however, that considerations like those raised in the last section lie behind his reading.

more influential arguments for thinking that Tarski *was not* advancing a fixed-domain conception of models in [21]. Neither of these arguments involves the direct analysis of Tarski's 1936 paper. Instead, they each involve the claim that fixed-domain conceptions of models are *so* implausible that it would be uncharitable to attribute such conceptions to Tarski. Hence, to the extent that Tarski *seems* to be advancing a fixed-domain conception of models, we should attribute this to Tarski's failure to make himself clear, rather than to his having any real intention of advancing such a conception.[12]

The first reason one might think that Tarski *must* have been advancing a variable-domain conception of models rests on the claim that fixed-domain conceptions were (and still are) "non-standard." When we look to the writings of other model theorists who were working around the time "Logical Consequence" was published, we find that they often assume a variable-domain conception of models.[13] More importantly, Tarski himself assumes such a conception in several papers written around this time.[14] Finally, when we look to the years following the publication of Tarski's paper, we find that variable-domain conceptions of models have become the dominant conceptions among mainstream logicians (including Tarski himself).[15]

This "non-standardness" point can be strengthened by noting that many of the theorems of classical model theory actually *depend* on our ability to vary the domains of our models. So, for instance, the upwards and downwards Löwenheim-Skolem theorems involve our ability to construct models with different size domains.[16] Similarly, standard proofs that arithmetic and analysis can be given categorical (second-order) axiomatizations depend on the fact that we are allowed to vary the domains of our models (as the domains of our models of arithmetic need to be countable, while the domains of our models of analysis must be uncountable). Hence, the "non-standardness" of a fixed-domain conception of models seems to have

[12]In [9], for example, Wilfrid Hodges argues that Tarski simply "omits to mention that the quantifiers of a formal language can be relativised to range over the domain of a structure rather than over the whole universe of individuals." Hodges attributes this "omission" to the fact that Tarski was writing for philosophers and that Tarski "didn't think philosophers would be interested in relativisation of quantifiers" (see [9], p. 138).

[13]See, for instance, sections III.11 and IV.6 in Hilbert and Ackermann's "Principles of Mathematical Logic" [8]. Here, the authors examine a collection of axiom systems (for, e.g., arithmetic and geometry) and explain how to interpret these systems with respect to different "domains of individuals." Similarly, Skolem discusses techniques for interpreting axioms in different "domains" in [14].

[14]In [24], for instance, Tarski proves two theorems concerning the categoricity of several (second-order) systems of axioms. First, he proves that the axioms for second-order arithmetic are categorical *on the assumption that these axioms include an axiom stating that every individual is a number*. Second, he proves that there is a categorical set of axioms which characterize the real numbers. Clearly, these two theorems cannot be jointly accepted on a fixed-domain conception of models. For, on such a conception, the first result would show that the number of objects in the world is merely countable (since we can find some model of the natural numbers with the whole world as its domain), while the second result would show that the world is uncountable (since it contains enough individuals to construct a model for second-order analysis). Given this, Tarski *must* have had a variable-domain conception of models when he wrote [24]. Similar comments apply to Tarski's work in [23].

[15]The fact that Tarski himself often employed a variable-domain conception of models (both at the time he wrote [21] and in his later work) is emphasised by Gómez-Torrente in [6]. and by Ray in [7]. They both take this as evidence that Tarski intended to introduce such a conception in [21].

[16]This "Löwenheim-Skolem" argument is advanced by Sher in [11] and by Gómez-Torrente in [6].

serious mathematical consequences: if we adopt a fixed-domain conception of models, then we may have to give up certain basic theorems of classical model theory.

These circumstances make it odd to think that Tarski would propose a non-standard definition of models without, at any point, commenting on the fact that this definition *was* non-standard. Suppose, for the moment, that Tarski was trying to introduce a fixed-domain conception of models in "On the Concept of Logical Consequence." At the time he was writing this paper, he was clearly aware that other logicians were proving theorems which depended on a variable-domain conception of models. Hence, given Tarski's general concern for clarity, he should have mentioned the fact that he was proposing a significantly different account of models. Similarly, once Tarski eventually adopted the variable-domain conception of models, he should have mentioned that he was changing his mind. The fact that he mentioned neither of these things suggests that he was not, after all, trying to introduce a fixed-domain conception of models.

To reinforce this point, note that Tarski actually *claims* that the conception of models introduced in "On the Concept of Logical Consequence" is a standard conception. Immediately following his definition of "model," Tarski writes: "in just this sense one usually speaks of models of an axiom system of a deductive theory" ([21], p. 417). Hence, either Tarski intended to introduce a variable-domain conception of models, or Tarski was confused concerning the state of mathematical practice in the period in which he wrote his paper. Or so, at any rate, the defender of a variable-domain reading of Tarski's paper might argue.[17]

Unfortunately, however, this argument both oversimplifies the state of mathematical practice in the period in which Tarski was writing and overestimates the mathematical importance of the variable-domain conception of models. To see this, we should begin by noting that there is a relatively straightforward technical trick which allows the proponent of a fixed-domain conception of models to obtain all the mathematical advantages of a variable-domain conception. Given a collection of sentences Γ, he has only to introduce a new predicate D (for domain) and to explicitly relativize each of the quantifiers in Γ to the predicate D. Having done this, he will induce a natural correspondence between the collection of variable-domain models of the original Γ and the collection of fixed-domain models of the newly relativized Γ'.

As a result, every theorem concerning the collection of variable-domain models for Γ can be translated into an equally interesting (and, indeed, essentially *identical*) theorem concerning the collection of fixed-domain models for Γ'. The Löwenheim-Skolem theorems, for instance, translate into theorems concerning the possible cardinalities of the sets picked out by D (when Γ and Γ' are first-order). Similarly, theorems asserting the catagoricity of second-order number theory and/or analysis translate into theorems asserting the existence of structure-preserving bijections between any two sets picked out by the predicate D (when D

[17] Note that this argument gives the defender of a variable-domain reading of Tarski a reply to one of the arguments offered in the last section of this paper. In that section, I argued that if Tarski intended to introduce a variable-domain conception of models, then he should have spent more time filling in the details of such a conception. But, it might be replied, if Tarski knew that the variable-domain conception of models was standard, he might have felt that he could get away with skimping on the details. The phrase "in just this sense one usually speaks of models of an axiom system" might have meant "if you don't understand this, just see any standard reference in the area."

is the "domain predicate" for some relativized axiomatization of number theory or analysis). Hence, *even if* fixed-domain conceptions of models are "non-standard," this non-standardness need not have the disastrous consequences for mathematical practice suggested by the proponent of a variable-domain conception of models.

In fact, however, fixed-domain conceptions of models were relatively common (indeed, relatively "standard") at the time Tarski wrote [21]. At this time, logicians often worked in some form of type theory in which the domain of individuals was taken as fixed. To axiomatize some particular mathematical structure—i.e. some structure other than the underlying universe of logical individuals, classes of individuals, classes of classes of individuals, etc.—they proceeded in precisely the manner described above: they introduced a new predicate to specify the domain of their structure and relativised their axioms to that predicate. So, for instance, both Russell and Carnap formulate the axioms of arithmetic in a language which includes the predicate N (for "is a number") among its primitives. Using axioms, they then ensure that the object picked out by "0" lives in the set picked out by N, that this set is closed under the function picked out by "Succ," and that the whole system "looks like" a copy of the natural numbers.[18]

More significantly, Tarski himself adopts this approach to models in a number of works written around the time "Logical Consequence" was published. In [22] and [19], for instance, Tarski provides relativised axiomatizations for Boolean algebras. In [16] and [17], he does the same for arithmetic, analysis, and (fragments of) geometry.[19] And even in papers where Tarski does employ a variable-domain conception of models, he often remarks that it is possible to prove the same results on a fixed-domain conception.[20] Hence, Tarski himself cannot think that fixed-domain conceptions of models are "non-standard" in any objectionable sense of that term.[21]

This, then, provides a response to the first argument against attributing a fixed-domain conception of models to Tarski. Contra that argument, fixed-domain conceptions of models do not have any substantial mathematical disadvantages. Nor would they have been considered "non-standard" at the time Tarski

[18] See [10] and [2]. Note that none of this makes much sense if we assume a variable-domain conception of models. On such a conception, there is no need to use axioms in order to ensure that our constants get interpreted into the domains of our models or that these domains are closed under a model's functions. On a variable-domain conception of models, this is all *built into* the very definition of "model."

It is worth noting here that both Russell and Carnap argue that the predicate N can be eliminated from the axioms of number theory *because this predicate can be defined in terms of the function "Succ"* (i.e., $x \in N \longleftrightarrow \exists y\, y = \mathrm{Succ}(x)$). Again, this argument would make very little sense on a variable-domain conception of models. On such a conception, the domains of our models come for free; they don't need to be *defined* in terms of other primitives.

[19] Note that it is only on a fixed-domain conception of models that, e.g., Tarski's axiom "$0, 1 \in B$" could have any point (see Postulate $\mathfrak{A}_6$ (a) in Tarski's [22]). On a variable-domain conception of models, this fact follows from the very definition of "model." In general, relativisation to a predicate doesn't *do anything* for a variable-domain conception of models: it amounts to little more than a (gratuitously complicated) way of adding the axiom "$\forall x\, D(x)$" to our theory.

[20] See, here, the footnotes on pages 310–314 of [24] and the concluding remarks on page 392 of [23].

[21] In [6], Gómez-Torrente also notices Tarski's use of relativised axiomatizations but argues that Tarski still intended to present a variable-domain conception of models in [21]. For the reasons discussed in section 2, I disagree with this conclusion.

wrote "Logical Consequence." In particular, Tarski himself made sufficient use of fixed-domain conceptions that there cannot be a straightforward "charity" problem with reading his 1936 paper as advancing such a conception.

At this point, I want to turn from this "non-standardness" argument and consider a second reason for thinking that Tarski was advancing a variable-domain conception of models in [21]. Recall from the last section that, if we use a fixed-domain conception of models in formulating Tarski's notion of logical consequence, then facts about the cardinality of the universe will be logical consequences of *any* sentences whatsoever. So, for instance, if it happens to be true that there are at least 37 things in the universe, then the following will count as a logical inference:

1. Tim exists.

So, 2. At least 37 things exist.

Even if fixed-domain conceptions of models are legitimate on their own, this result might be taken to show that such conceptions ought not to be combined with Tarski's definition of logical consequence. How, after all, can logic alone determine the number of things which happen to exist? To the extent that this question seems right—i.e., to the extent that we think that logic *cannot* determine the number of things which exist—it will also seem right to refuse to attribute a fixed-domain conception of models to Tarski. In short, the odd consequences of combining a fixed-domain conception of models with Tarski's definition of logical consequence provide us with a reason for reading [21] as advancing a variable-domain conception of models.[22]

This is, however, not a reason which would appeal to Tarski. For one thing, Tarski often treated the axiom of infinity as a part of logic (see, for instance, [18], [23], [24], or [25]).[23] Hence, he cannot be *too* averse to treating the cardinality of the universe as a logical issue. Indeed, in both [18] and [25], Tarski actually discusses the fact that his acceptance of the axiom of infinity forces him to reject finite domains and (thereby) affects his conception of logical consequence (see [18], p. 423 and [25], pp. 293–295). Finally, in [25] Tarski explicitly responds to concerns about logical axioms which entail the existence of infinite sets (see [25], p. 282 n. 2).[24]

More significantly, note that the above "problem" with Tarski's definition of models actually plays a key role in making sense of the connections between Tarski's analysis of logical consequence and some of the other sections in Tarski's paper. As I argued in section 3, Tarski's connection between logical and material consequence requires viewing the cardinality of the universe as a logical issue. So too does Tarski's insistence that there exists a relationship between logical consequence and substitutional consequence (for sufficiently rich languages). Also, as I will argue in section 5 of this paper, Tarski's decision to count ω-inferences as logical requires that analogs of the natural numbers can be defined in purely logical terms. Hence, some of

[22] This argument against a fixed-domain reading of Tarski can be found in Sher's [12] and Ray's [7]. Etchemendy also discusses the argument in [3] and [4], although he uses it more as a *criticism* of Tarski than as a tool for *interpreting* Tarski.

[23] Tarski also mentions this axiom approvingly in [16] (see pp 81 and 130).

[24] For additional discussion of Tarski's attitude towards the axiom of infinity, see the recent [6].

JSL 66 (Dec. 2001): 1701–1726

Tarski's own examples of logical consequence depend on treating the cardinality of the universe as an issue of "logic" (since they require that logic alone be able to prove the existence of "natural numbers").

Given all this, it seems implausible to count arguments of the sort just sketched as definitive (or even reasonable) with respect to the interpretation of [21]. Tarski clearly knew that a fixed-domain conception of models would lead to the result that questions concerning the cardinality of the universe become issues of logic. Just as clearly, this result did not bother him. Indeed, some of the arguments he presents in "Logical Consequence" actually depend on this result. This fact, together with evidence discussed in section 2, points overwhelmingly to the conclusion that Tarski intended to advance a fixed-domain conception of models in his 1936 paper.

4 A Puzzle about ω-inferences

At this point, I want to turn away from Tarski's conception of models to examine more carefully the ω-inference example which we discussed earlier. Recall the context in which this example occurred. Tarski was trying to show that derivational systems cannot capture our intuitive understanding of logical consequence. To do this, he noted that standard derivational systems do not allow us to perform ω-rule style inferences, even though the consequence relation between the premises and conclusions of such inferences seems to be straightforwardly logical.[25] Thus, he concludes, logical consequence must outstrip derivability in any standard system of deduction.

This example poses a problem for the reader of Tarski's paper. To the extent that Tarski is simply trying to resist derivational analyses of consequence, the example serves him well. However, just as the ω-rule outstrips derivability in standard systems of deduction, it also outstrips model-theoretic consequence in standard systems of model theory. So, in this regard, Tarski's example seems to serve him poorly, for it seems to threaten his own (model-theoretic) analysis of logical consequence. The remainder of this paper will explain why this threat (and a related threat involving Tarski's discussion of Gödel sentences) does not, in the long run, prove fatal to Tarski's analysis of logical consequence.

To begin, let me say a little more concerning the reasons why Tarski's ω-inference example might prove threatening in the first place. Suppose that Tarski's conception of models is a standard, first-order, variable-domain conception. Then, it is straightforward to show that combining this conception of models with Tarski's definition of logical consequence leads to a conception of logical consequence which fails to count (at least some) ω-inferences as logical. So, for instance, let $\mathcal{L}$ be any first-order language sufficient for doing arithmetic, and let T be an axiomatization of arithmetic in $\mathcal{L}$. Expand $\mathcal{L}$ to $\mathcal{L} = \mathcal{L} \cup \{P\}$ where P is a new unary predicate. Then we can find a model M such that M satisfies each of the following sentences:

[25] The point, of course, is not merely that standard derivational systems fail to explicitly include the ω-rule among their axioms. Rather, it is that standard derivational systems cannot prove (even by way of other inference rules) that the conclusion of an ω-inference follows from its premises.

A_0. 0 possesses property P,
A_1. 1 possesses property P,
. . . .
A_n. n possesses property P.

At the same time, M also satisfies the sentence:

A. There exists a number x which does not possess the property P.

Therefore, if Tarski's model-theoretic analysis of logical consequence is fleshed out using this "standard" conception of models, the model M will witness the fact that Tarski's ω-inference is not an instance of logical consequence.[26]

Now, we might be tempted to think that this difficulty can be evaded simply by accepting the interpretation of Tarski presented in the first portion of this paper. After all, if Tarski's conception of models is not the "standard, first-order, variable-domain" conception, then results such as the one cited above will have little to do with Tarski's conception of logical consequence.[27] Unfortunately, however, nothing intrinsic to the results mentioned above rests on the fact that first-order logic allows its quantifiers to range over variable domains. Even if we accept Tarski's fixed-domain semantics, we can readily reconstruct all of the problems sketched in the last paragraph.

To see this, we need to consider a slight strengthening of the result mentioned two paragraphs ago. For the moment, assume that we are still working with a variable-domain conception of models, and let $\mathcal{L}$ and T be as given. Then, if N is *any* model for T, we can find an M with the properties previously mentioned such that M and N share the same domain—so, N and M interpret quantifiers the same way, but they give different interpretations to the predicates and relations.

Now, let us switch our attention to a fixed-domain conception of models. Let $\mathcal{L}$ be some "first-order" language, let T be some axiomatization of arithmetic in $\mathcal{L}$, and let N be some model of T—so, N is a sequence of "objects" which interpret the constants, predicates, relations and functions of $\mathcal{L}$. Combining this sequence with an explicit mention of the universal domain, we obtain a variable-domain model N'; further, N' satisfies

[26] Note that very little in this example rests on the fact that we needed to expand the language in order to obtain our chosen predicate P. As long as the theory T is recursively axiomatizable (and as strong as, say, Peano Arithmetic), we can forget about expanding the language and simply let P be picked out by a formula, $\phi(x)$, in our original language $\mathcal{L}$. Then we can still find a model M such that:

1. For every natural number n, M satisfies "$\phi(n)$."

and 2. M satisfies "there is a number x such that $\neg\phi(x)$."

So, once again the ω-rule does not constitute a (model-theoretically) valid form of inference.

[27] Indeed, I have sometimes heard people suggest that this line be used as an *argument* for the conclusion that Tarski intended to endorse a fixed-domain conception of models. The reasoning, presumably, goes as follows: variable-domain model theory cannot capture ω-inferences; Tarski's conception of logical consequence can capture ω-inferences; therefore Tarski must have had a fixed-domain conception of models when he defined logical consequence.

For the reasons mentioned in the text, I find this line of argument unpersuasive. If the only difference between Tarski's conception of models and more standard conceptions involves the variability of a model's domain, then Tarski has simply failed to put forward a definition of logical consequence which will count ω-inferences as logical.

exactly the same sentences as N satisfied. Next, we apply the result mentioned in the previous paragraph to obtain a (variable-domain) model M' such that M' and N' share the same domain, and M' has the type of ω-rule violation we are after. Since M' and N' share the same domain—i.e. the universal domain—the sequence M of constants, predicates, relations and functions associated with M' constitutes a fixed-domain model of T. And because this sequence still gives a model with an ω-rule violation, it witnesses the fact that ω-inferences are non-logical.

Therefore, Tarski's "ω-inference problem" cannot be solved simply by adopting an interpretation which makes all of Tarski's models share a single domain. To solve it, we need some stronger means of restricting the models considered when testing an inference for logicality. One suggestion—a suggestion broached by John Etchemendy—involves assuming that Tarski considered many of the terms associated with number theory to be logical constants. Etchemendy asks:

> How did Tarski see his own definition of consequence avoiding the very same objections he leveled against the competing definitions? Why shouldn't ω-incompleteness, or the Gödel sentences, also present problems for his analysis?... The answer ... lies in the flexibility Tarski allows in our choice of "logical" constants. Clearly, if we choose to treat the numerals '0', '1', '2',..., as logical constants, as well as the quantifier 'every natural number', then the sentence A will come out as a consequence of the infinite sentences A_1 , A_2 , A_3 ,... ([3], p. 73).

Etchemendy's idea is this: by considering all of the terminology associated with number theory to be logical, Tarski can ensure that logical consequence coincides with material consequence for sentences which use only number-theoretic terminology (plus, of course, the traditional logical constants). Thus, since the ω-rule is materially valid, this analysis of logical consequence makes it logically valid as well.

However, Etchemendy argues, this solution to the ω-inference problem borders on being arbitrary. By stipulating that a certain portion of our vocabulary counts as logical, we can always make material consequence and logical consequence coincide for sentences which use only this specified portion of our vocabulary. So, for instance, by counting all terms in our language as "logical," we can make all instances of material consequence count as instances of logical consequence. And this, Etchemendy argues, "involves a certain trivialization of Tarski's analysis. For with this choice of logical constants, a true sentence is a logical consequence of *any* set of sentences whatsoever" ([3], p. 73).[28]

To avoid this "trivialization," Tarski needs to produce a criterion for determining which terms are to count as logical and which are to count as non-logical. He needs, that is, to produce some grounds for distinguishing between number-theoretic terms (which should count as logical) and more ordinary terms like "cat," "mat," and "rat" (which should count as non-logical). If Tarski cannot produce such a criterion (or, at the very least, produce some positive argument for counting number-theoretic terms as logical constants), then his solution to the problem of logical constants (and the original, related problem of ω-inferences) will be inadequate, amounting to a "solution" by sheer stipulative fiat.[29]

[28] See chapters 6 and 9 of [4] for the development of a similar theme.

[29] As Etchemendy rightly points out, there is no place in "On the Concept of Logical Consequence" where Tarski attempts to overcome this inadequacy—e.g., where he presents an explicit argument for counting number-theoretic terms as logical. That

These difficulties are increased when we turn from Tarski's ω-inference example to his discussion of Gödel sentences. Recall that, immediately following his discussion of ω-inferences, Tarski writes:

> It follows from the work of K. Gödel that. . . in every theory (apart from theories of a particularly elementary nature), however much we supplement the ordinary rules of inference by new, purely structural rules, it is possible to construct sentences which follow, in the usual sense, from the theorems of this theory, but which nevertheless cannot be proved in this theory on the basis of the accepted rules of inference. ([21], 412)

This passage raises three problems. First, it exacerbates the problem created by Tarski's ω-inference example. Just as ω-inferences aren't captured by standard systems of model theory, neither are inferences from incomplete theories to their Gödel sentences. Hence, Tarski's own model-theoretic account of logical consequence seems to fare no better at capturing "the usual sense" of logical consequence than the derivational accounts which Tarski rejects.[30]

Second, the difficulty with Gödel sentences is harder to patch up than the corresponding difficulty with ω-inferences. In the ω-inference case, one can imagine a plausible argument for treating number-theoretic terms as "logical" (and, so, solving the ω-inference problem in the manner suggested by Etchemendy). In the case of arbitrary Gödel sentences, this solution is essentially unavailable. Because almost any fragment of our language can be used to formulate an incomplete theory (of the sort which generates a canonical Gödel sentence), this solution requires treating almost all terms as logical constants. Since this trivializes both the notion of "logical constant" and that of "logical consequence,"[31] it cannot be an an acceptable solution to Tarski's general problem.[32]

Third, and most important, it's unclear that we even *want* the inference from an incomplete theory to its Gödel sentence to count as an instance of logical consequence. Certainly, we don't want deductively incomplete theories to decide *all* sentences in their languages. We would not, for instance, want PA + "2 is less than Tim's favorite number" to entail *either* "3 is less than Tim's favorite number" or "3 is not less than Tim's favorite number." Nor, I think, do we want every theory with a Gödel sentence to *automatically* entail that sentence. Suppose, for example, that we start with some non-standard model for first-order arithmetic and develop a theory which tries, from the very beginning, to describe this model. (To keep our theory

being said, there are two things which should be noted on Tarski's behalf. First, Tarski was clearly *aware* of the need for a criterion to distinguish between logical and non-logical terms, and he was also aware that he had not yet provided this criterion (see pp. 418–420 of [21]). Second, Tarski does not need a *general* criterion of "logicality" to make sense of his ω-inference example or his discussion of Gödel sentences. In section 5, I will discuss what Tarski *does* need here, and I will show that he has enough resources to avoid Etchemendy's criticisms.

30 In general, Tarski's problems here are caused by the completeness theorem for first-order logic. This theorem ensures that, given a standard, first-order, definition of models, the model-theoretic conception of logical consequence will correspond *exactly* to the derivational conception. As a result, any discrepancies Tarski uncovers between the derivational conception of consequence and the "intuitive" conception of consequence can simply be "read-over" into discrepancies between the model-theoretic conception and the intuitive conception.

31 As discussed in section 2, treating all terms as logical constants reduces logical consequence to material consequence.

32 This point is also made by Etchemendy. See [3], p. 73.

from looking like a mere arithmetic wanna-be, we can formulate it in the language $\{c, f, g, h, i\}$ instead of $\{0, \text{succ}, +, \times, <\}$.) Provided that we start with the right model of arithmetic, our theory will amount to a simple transcription of PA + $\{\neg\text{Con(PA)}\}$. In this case, it should be pretty clear that we *don't* want our theory to entail its own Gödel sentence, since our theory actually proves the *negation* of this sentence—i.e., modulo our transcription we have PA + $\{\neg\text{Con(PA)}\} \vdash \neg\text{Con(PA} + \{\neg\text{Con(PA)}\})$.[33] Hence, it's mysterious why Tarski thinks that these "Gödel inferences" *should* be captured by our account of logical consequence (whether this account is formulated in derivational, model-theoretic, or purely intuitive terms).

These, then, are some of the problems which Tarski's ω-inference example and his discussion of Gödel sentences engender. In each case, we have a certain inference (or class of inferences) which Tarski regards as intuitively logical. In neither case do the inferences in question seem to count as logical under Tarski's own, model-theoretic definition of logical consequence. In both cases, philosophical problems result from attempting to patch up this definition so as to make it fit Tarski's intuitions (problems of arbitrariness in the ω-inference case and problems of trivialization in the Gödel sentences case). Finally, in the case of Gödel sentences (and perhaps even in the case of ω-inferences) it's unclear why Tarski even *wants* to count the relevant inferences as "intuitively logical" in the first place.

5 Tarski and the ω-rule I

To resolve these interpretive problems, it's useful to go back and reexamine the passage where Tarski first introduces his ω-inference example. Tarski writes, "some years ago I gave a quite elementary example of a theory which shows the following peculiarity..."; he then provides a reference to the paper in which this "elementary example" was presented.[34] When we examine this original paper, we notice two facts. First, the paper explicitly assumes that we have access to the full apparatus of second-order logic (in the form of type theory as developed in *Principia Mathematica*). Second, although the paper introduces a number of abbreviations which Tarski employs as aids to exposition, the paper nowhere introduces primitive symbols other than those intrinsic to the formulation of the underlying logic. In particular, the paper nowhere introduces primitive symbols to refer to individual numbers; nor does it introduce a symbol to refer to the whole collection of natural numbers.

These facts have important consequences for our understanding of Tarski's ω-inference example. Because Tarski never introduces primitive symbols to refer to natural numbers, he is forced to define analogs of the natural numbers using only the resources of second-order logic. In particular, he uses finite sets as stand-ins for the natural numbers, and he formulates his ω-inferences in terms of such sets. Given this, a more perspicuous statement of Tarski's ω-inference example might read as follows:

[33] Of course, Tarski might argue that this theory is logically inconsistent. But, given that this theory "tries to" describe a particular structure—and, indeed, given that it actually *succeeds* in describing this structure—it seems implausible to label the theory "inconsistent." To do so is to say that certain structures (like the one in question here) simply can't be described!

[34] The quote is from page 410 of [21]. The paper cited is [25] (see particularly pp. 288–295).

> Some years ago I gave a quite elementary example of a theory which shows the following peculiarity: among its theorems there occur such sentences as:
>
> A_0. The empty set possesses the given property P,
> A_1. Sets containing only one element possess the given property P,
>
> and, in general, all particular sentences of the form
>
> A_n. Sets containing exactly n elements possess the given property P.
>
> On the other hand the universal sentence:
>
> A. Every finite set possesses the given property P,
>
> cannot be proved on the basis of the theory in question by means of the normal rules of inference.

Here, phrases such as "empty set," "set containing exactly n elements," "finite set," etc. are short for the standard second-order formulae corresponding to these phrases.[35]

There are three things to notice about this formulation of Tarski's example. First, the ω-inference in this example is model-theoretically valid: given a standard (second-order) semantics for type theory, the sentence A will be true in any model in which $A_0, \ldots, A_n, \ldots$ are true. This is, I think, why Tarski wasn't too worried about the problem discussed at the beginning of section 4. Although it's quite true that ω-inferences aren't valid in first-order model theory—and that, as a result, first-order, model-theoretic conceptions of logical consequence run into the same problems that derivational conceptions run into—ω-inferences *are* valid in second-order model theory. Therefore, since Tarski formulates both his ω-inference example and his account of logical consequence in second-order terms, there need be no conflict between his model-theoretic account of logical consequence and his insistence that ω-inferences count as (intuitively) logical.[36]

Second, once we recognize the second-order character of Tarski's ω-inference example, we are in a better position to understand his discussion of Gödel's work. One consequence of Gödel's work in [5] is that no derivational system (or, at least, no recursively specifiable derivational system) can capture all of the model-theoretically valid inferences of second-order logic. Therefore, if we regard these inferences as "intuitively logical," then we find a clear sense in which deductive theories don't—and, indeed, *can't*—capture the

[35]See pp. 287-8 of [25] for a formulation of the relevant formulae. On p. 288, Tarski explicitly discusses the fact that these formulae concern finite *sets.*

[36]It is worth noting that Tarski's use of higher-order logic here directly contradicts a claim made (repeatedly) by John Etchemendy. In [3], Etchemendy claims that Tarski's ω-inference example is supposed to show that "standard deductive systems *do not*, and natural extensions of them *cannot*, capture the logical consequence relation *even restricted to first-order languages*" (see [3], p. 72). Similarly, in [4], Etchemendy claims that "both [the ω-inference example and the Gödel's theorem example] involve the consequence relation for first-order languages" (see [4], p. 85). However, even a cursory examination of the original source of Tarski's ω-inference example shows that these claims are false: the paper employs higher-order variables from the outset, and these variables are explicitly used in formulating the ω-incomplete theory to which Tarski later refers.

The fact that Etchemendy was wrong in taking Tarski's ω-inference example to be an illustration of first-order inference is also noticed in [6], [7] and [12].

intuitive notion of logical consequence.[37] On my reading, this is the key fact which underlies Tarski's comments in [21].

Note that this reading of [21] has several advantages. First, it solves the three problems mentioned at the end of the last section: it explains why Tarski uses Gödel's theorems to motivate a distinction between derivational and model-theoretic accounts of consequence, it allows Tarski to formulate this distinction without making *ad hoc* decisions regarding logical constants, and it explains Tarski's insistence that the inferences at issue should be regarded as "intuitively logical."[38] Second, this reading squares well with Tarski's general tendency to treat second-order inferences as part of logic (see, e.g., [23], [22], [19], and the crucial [25]). Finally, and perhaps most importantly, this reading squares with other comments Tarski makes about Godel's theorems—i.e., comments he makes in other papers. In both [25] and [18], for example, Tarski treats higher-order inferences as "logical" and then appeals to Gödel's theorems to prove that deductive systems can't capture the relevant notion of "logic."[39] Hence, the present reading of [21] fits this paper into a broader pattern in Tarski's thought.

This brings me to a final point concerning Tarski's ω-inference example. When we look closely at the original formulation of this example, we see that it supports the fixed-domain conception of models which was presented in section 2 of this paper. To see this, note that [25] is one of the places where Tarski treats the axiom of infinity as a principle of logic (see p. 289 of [25]). Nor is this treatment of Infinity surprising: Tarski needs *something like* the axiom of infinity to ensure that his models contain enough elements to formulate the ω-inferences he is interested in.[40] As a result, Tarski's own understanding of his ω-inference example seems to require (at least) an "infinite domain" conception of models.

This demand for infinite domains, however, tells strongly against variable-domain readings of [21]. After all, part of the point of variable-domain conceptions of models is to allow models to have *arbitrary* cardinality (including arbitrary *finite* cardinality). And this is precisely what Tarski must reject if his analysis of ω-inferences is to be viable. Of course, someone might have a conception of models under which all models have infinite domains, but where these domains are allowed to vary (and even, perhaps, to vary in cardinality).

[37]If desired, this point can be formulated in terms of canonical Gödel sentences. Let T be a recursive, satisfiable, theory in some second-order language. Using the fact that second-order logic has the ability to code ordinary arithmetic (i.e., via the finite-sets trick mentioned in the main text), we can formulate the consistency of T as a sentence of pure (second-order) logic. Then, even though this sentence is a *second-order validity,* the argument of [5] shows that T does not prove this sentence.

All that being said, we should note that Tarski himself never refers to canonical "Gödel sentences." He simply argues that, if we are given a standard deductive theory, we can find *some* sentence which follows intuitively from this theory, but which cannot be derived from it. *This much* follows purely from the facts cited in the main text; there is no need to bring Gödel sentences into the discussion.

[38]Since they are, after all, valid inferences in a perfectly standard system of logic. See section 6 for more on this point.

[39]See pp. 294–5 of [25] and p. 271 of [18]. Note that [25] is the very paper in which Tarski formulates the ω-inference example to which he refers in [21], and that Tarski's discussion of Gödel's theorems in this paper comes only a few pages after his formulation of the ω-inference example.

[40]I.e., to construct all of the finite sets which play the role of numbers in these ω-inferences.

This is, however, a singularly unmotivated conception of models, and I see no reason for attributing it to Tarski. In the end, therefore, I conclude that a proper understanding of Tarski's ω-inference example provides yet further confirmation that Tarski was employing a fixed-domain conception of models when he formulated the analysis of logical consequence in [21].[41]

6 Tarski and the ω-rule II

In this section, I want to address two philosophical questions which the last section may have provoked. First, to what extent does my solution to the ω-inference problem (and the related problem about Gödel sentences) square with the solution given by Etchemendy? On one level, the two solutions are quite similar. They both demand that we treat terms other than the usual first-order connectives as "logical constants"; they both use the fact that logical consequence and material consequence coincide for sentences which contain only logical constants; and they both wind up counting ω-inferences as "logical" *because* these inferences can be formulated using only logical constants (after, of course, the notion "logical constant" has been suitably reinterpreted). As a result, the two solutions can be viewed members of the same species.

That being said, there are two important differences between my solution and Etchemendy's. First, Etchemendy's solution involves attributing to Tarski some rather *ad hoc* modifications of the notion "logical constant" (e.g., by adding in number-theoretic terminology to make ω-inferences count as "logical" and by adding in essentially arbitrary terminology to take care of Gödel's theorems). My solution, on the other hand, requires no *ad hoc* modifications. On my solution, the only logical constants are those from standard formulations of type theory. And, although it is *currently* fashionable to question the status of type theory as logic, at the time Tarski wrote "On the Concept of Logical Consequence," most logicians had no such scruples.[42] Hence, given the standards of his time, there is nothing "*ad hoc*" about Tarski's decision to treat type-theoretic terms as logical; it's simply the "standard" decision for him to make.

Second, Etchemendy's solution allows us to treat number-theoretic terminology as *primitive* when we test ω-inferences for logicality (arguing, of course, that this terminology consists of *logical* primitives rather than non-logical primitives). My solution, on the other hand, requires number-theoretic terminology to be "defined away" into some rather complicated second-order expressions.[43] As a result, my solution preserves

[41]It is interesting to note that [6], [7] and [12] all discuss the use of higher-order logic in [25], but then fail to notice the relevance that this has for the question of fixed vrs. variable domains in [21].

[42]In general, Tarski's contemporaries tended to take the system of *Principia Mathematica* as a standard. Certainly Tarski himself used this system in his own mathematical work, and he seemed to think that he was simply following standard conventions when he did so. See, for instance, [25], [24] and [23]. For my own part, I think Tarski and his contemporaries were right in regarding higher-order logics as "logic." This, however, is an issue for another paper.

[43]As formulated in [25], the premises of ω-inferences don't even contain expressions which refer to specific objects that play the role of individual numbers. The premises simply attribute a property to *all* sets of a specific size, and the conclusion talks about *all* finite sets (i.e., and not about *all numbers*).

It should be noted here that this lack of singular reference could be overcome if Tarski used λ-expressions to define the less of the "surface structure" of Tarski's ω-inference example then Etchemendy's solution does (at least when this example is written out in the manner of [21] rather than that of [25]).

Now, as a matter of pure interpretation, this failure to respect surface structure doesn't bother me very much. Tarski clearly intends his example in [21] to be a mere summary of the more detailed example given in [25]. Hence, it is this latter example which is normative for the purposes of interpretation. And in this latter example, it is clear that Tarski's ω-inferences involve syntactically complicated surrogates for number-theoretic terms. Tarski discusses these surrogates on pages 287–288 of [25], and he explicitly notes that his examples *don't* use primitive number-theoretic terms on page 288 (see fn. 2). From an interpretive point of view, then, the fact that my reading modifies the surface structure of Tarski's example really isn't a problem.

There is, however, a problem in the neighborhood. On my reading of Tarski, ω-inferences formulated in second-order terms come out "logical," while ω-inferences formulated in first-order terms come out "non-logical." Why doesn't this fact, in and of itself, create a problem for Tarski's account of logical consequence? Isn't there something odd about the idea that the *same inference* can come out logical under one formulation and non-logical under another? Shouldn't the logicality of an inference be independent of the way that inference happens to be formulated?

To explain why these questions probably wouldn't bother Tarski, I begin with two preliminary points. First, the underlying issue here has very little to do with the relationship between first and second-order logic: analogous problems can be raised regarding inferences which don't involve this relationship. Consider, for instance, the results of taking "Jack" to mean "All men are mortal," "Jill" to mean "Socrates is a man," and "Hill" to mean "Socrates is mortal." Then there is a clear sense in which the following two arguments capture the same inference:

	1. All men are mortal.		1′. Jack.
	2. Socrates is a man.		2′. Jill.
So,	3. Socrates is mortal.	So,	3′. Hill.
	(G)		(H)

Nevertheless, the first of these arguments meets Tarski's test for logicality, while the second (presumably) does not.[44] This discrepancy—together with the fact that it's clearly the same kind of discrepancy as the one highlighted in the last paragraph—shows that the relationship between first and second-order logic is something of a red herring in the present context.[45]

numbers—e.g., if he defined n as the set of all n-element sets of individuals (cf. Tarski's discussion on 233–4 of [18]). However, this would still involve treating number-theoretic terms as abbreviations for substantially more complicated expressions, so it would still modify the surface structure of Tarski's example. And anyway, it's not how Tarski actually does things.

[44]The only way it could meet the test is if we (ludicrously!) decided to treat "Jack," "Jill," and "Hill" as logical constants.

[45]Two comments about this example are in order. First, the example involves a single inference formulated in both first-order and propositional form. The inference instantiates a valid first-order pattern, but it does not instantiate any valid propositional patterns (this is, in essence, the fact that Tarski's test picks up on). This example highlights a common phenomenon: *many* inferences instantiate valid higher-order patterns without instantiating valid lower-order patterns (e.g., ω-inferences!). In such

Second, *all* of the accounts of logical consequence that Tarski tries out in [21] run into the kinds of difficulties under discussion here. Clearly derivational accounts of logical consequence are sensitive to the ways particular inferences are formulated: few derivational systems would count argument (H) as logical, though most would count argument (G) that way.[46] Similarly, on the assumption that we have enough extra-logical constants to evade the problems discussed in section 1, Tarski's condition (F) will count argument (G) as logical and argument (H) as non-logical (and, for that matter, second-order ω-inferences as logical and first-order ω-inferences as non-logical). Finally, and most obviously, Tarski's own model-theoretic account of consequence entails that the logical status of an inference depends on the way that inference happens to be formulated.

Given, then, that our problems here are not ultimately about the relationship between first and second-order logic, and given that these problems are not peculiar to the model-theoretic account of logical consequence, what *is* at the heart of these problems? On my view, the real issue here is Tarski's understanding of logical consequence as a *formal* relationship. On page 414 of [21], Tarski writes:

> Certain considerations of an intuitive nature will form our starting-point. Consider any class K of sentences and a sentence X which follows from the sentences of this class. . . . Since we are concerned here with the concept of logical, i.e., *formal,* consequence, and thus with a relation which is to be uniquely determined by the form of the sentences between which it holds, this relation cannot be influenced in any way by empirical knowledge, and in particular by knowledge of the objects to which the sentence X of the sentences of the class K refer.

On this understanding, then, logical consequence is characterized by two things: it is a relationship between sentences (or sets of sentences), and it is a relationship which holds (or fails to hold) in virtue of those sentences' syntactic forms.

Keeping this conception of formality in mind, we can see why Tarski might not *want* his account of logical consequence to treat first and second-order formulations of the ω-rule as equivalents. To do so, Tarski's account would have to treat syntactically divergent sentences as logically similar (since it would have to treat the premises of first-order ω-inferences as equivalent to those of second-order ω-inferences). More crucially, it would have to treat syntactically similar sentences as logically divergent (since it would have to distinguish the premises of first-order ω-inferences from the syntactically similar premises which occur in (invalid) inferences involving non-standard models). Each of these moves, however, conflicts with

cases, Tarski's test for logical consequence proves sensitive to the underlying logic in which an inference is formulated; it counts higher-order formulations as logical and lower-order formulations as non-logical).

Second, although the relationship between higher and lower-order logics often gives rise to the kinds of problems at issue here, this relationship is not *essential* to the problems. To see this, we need only rework the example above so as to let "Jack" mean "If Socrates is a man, then Socrates is mortal." In this case, *all* our inferences are propositional, but the problems from the main text still arise.

[46]The fact that derivational conceptions of consequence run into the same difficulties as Tarski's model-theoretic conception may be obscured if we focus solely on the ω-rule case (since standard derivational systems cannot capture *either* first or second-order ω-inferences). However, if I am right that the ω-rule case is simply an instantiation of a broader difficulty and that the example in the main text is another such instantiation, then it should be clear that derivational conceptions pose the same difficulties as Tarski's model-theoretic conception.

Tarski's insistence that logical consequence is formal, that it is, in his words, "determined by the form of the sentences between which it holds."[47]

This, then, provides an explanation for the fact that Tarski's account of logical consequence makes the logical status of first-order ω-inferences differ from that of their second-order cousins: Tarski needs to allow this kind of difference in order to preserve his conception of logic as a *formal* relation.[48] Three further comments are in order here. First, the understanding of formality just sketched is one which runs fairly deep in [21]. We have already seen that Tarski takes this understanding as "intuitive" and as a "starting point" for his investigation of logical consequence. We have also noted, in effect, that *all* of the accounts of logical consequence which Tarski takes seriously in [21] are formal in character.[49] Indeed, at one point Tarski even explores the possibility of deriving an account of logical consequence from the mere notion of formality (see the initial argument for condition (F) on pp. 414–415).

Second, Tarski is clearly aware that his emphasis on formality renders his account of logical consequence sensitive to the ways specific inferences are formulated. Immediately after his introduction of condition (F), Tarski mentions the "necessity of eliminating any defined signs which may possibly occur in the sentences concerned, i.e., of replacing them by primitive signs."[50] Now, if Tarski thought that an account of logical consequence should be indifferent to the ways specific inferences are formulated, then this insistence on eliminating defined terms would be somewhat mysterious: after all, definitional substitution is a fairly innocuous way of "reformulating" a particular inference. On the other hand, if Tarski *expected* his account of logical consequence to be sensitive to the ways specific inferences are formulated, then his insistence on eliminating defined terms would be quite natural. The failure to eliminate defined terms can dramatically simplify the syntactic structure of a sentence, and this, in turn, may seriously affect any account of logical consequence that is sensitive to issues of syntactic form.[51] Hence, Tarski's attitude towards the elimination of defined terms—i.e., his insistence on going through with such elimination—shows that he was well aware that his account of logical consequence tends to make the logicality of an inference depend on the way that inference is formulated.

[47]Similarly, Tarski would probably jib at treating arguments (G) and (H) as equivalents for the purposes of assessing logicality. In order to treat them this way, Tarski would have to overlook the syntactic differences between "All men are mortal" and "Jack," and he would have to overlook the syntactic *similarities* between, e.g., "Jack" and "Jill."

[48]More generally, Tarski's conception of logic as formal requires him to accept a whole menagerie of cases in which the logical status of an inference depends on the way that inference happens to be formulated. See, e.g., the example on page 22. We should note, here, that the issue of formality concerns the sentences *actually used* in formulating an inference, not the overall richness of the background language: argument (H) fairs no better if we assume that our *background* language is second order. In this sense, then, the issues here are quite different from those discussed in connection with condition (F) (see pp. 3–5).

[49]I.e., all of them are formulated so as to make the consequence relation between two sentences (or sets of sentences) depend on those sentences' syntactic forms. This fact explains the fact, noted earlier, that all of Tarski's accounts are susceptible to the objections raised on p. 22.

[50]In context, it's clear that this necessity is supposed to apply *both* to condition (F) itself *and* to the model-theoretic account of logical consequence which immediately follows it (cf. p. 415).

[51]E.g., consider the case where we just *define* "Jack" to mean "All men are mortal," etc.

Nor is Tarski's attitude here particularly surprising in light of his other work. In [18], Tarski argues that we can only understand the concept of truth if we formulate this concept in the right kind of language—i.e., in a language stratified so as to preserve language/metalanguage/meta-metalanguage distinctions.[52] Similarly, in [26], Tarski argues that set theory counts as logical *when it is formulated in type-theoretic terms.* On the other hand, Tarski claims that set theory *does not* count as logical when it is formulated in first-order terms (see [26], pp. 151–153).[53] In both of these cases, then, Tarski claims that certain logical and semantical notions *require* formulation in specific kinds of languages, using specific background logics, in order to make sense. Alternate formulations, although they may *seem* to "get at" the same notions, are non-equivalent for the purposes of serious formal investigation.

At the end of the day, then, I do not think that Tarski would be all that concerned about the problems raised on page 22. Tarski clearly knew that his understanding of logic as formal leads to the result that the logicality of certain inferences depends on the way these inferences are formulated. Equally clearly, this result did not bother him very much. Indeed, the result seems to be of a piece with other things Tarski wrote over the course of his career. From an interpretive standpoint, therefore, the problems on page 22 can—and I think should—be simply dismissed.[54]

7 Conclusion

To conclude, I have used this paper to argue for three claims concerning the interpretation of Tarski's 1936 paper, "On the Concept of Logical Consequence." First, I have argued that Tarski intended to introduce a fixed-domain conception of models in this paper. This reading fits Tarski's text better than others do, and, as I have shown in section 3, it is not susceptible to certain obvious objections. Second, I have argued that the relationship between Tarski's ω-inference example and the status of logical constants is more complicated

[52]See [18], p. 267. Tarski argues that any attempt to get a grip on "truth" in colloquial languages requires us to "reform" these languages first. Since the reforms in question come very close to *replacing* the languages—e.g., by eliminating their "naturalness" and by forcing them to "take on the characteristic features of the formalized languages"—Tarski thinks this project is somewhat pointless. Most importantly, Tarski thinks there is no possibility of understanding "truth" in *unreformed* natural languages or in formal languages that lack the requisite sort of stratification. Any attempt to do so "leads inevitably to confusions and contradictions."

[53]See also Tarski's comments in the discussion transcribed at the end of [15]. It's worth noting that both of these references are a good deal later than [21] itself. Nevertheless, Tarski doesn't seem to think he's changing his mind about anything in these later papers, and the references provide evidence of a general pattern in Tarski's thought.

[54]Of course, there are some larger, non-interpretive, problems in the vicinity. Was Tarski *right* to understand logic as "formal" in the sense described on page 23. Could such an understanding really *justify* an account of logical consequence which treats first and second-order ω-inferences as differently as Tarski's account seems to? For what it's worth, I'll note that I think the answers to these questions are probably "no." However, exploring these issues would take us rather far from Tarski's 1936 paper and would extend *this* paper rather unreasonably. For now, then, I'll simply note that Tarski was in good company in understanding logical consequence as formal: Carnap had a similar understanding of logic, as did Hilbert. Indeed, it was precisely this understanding of logic as formal which lay at the heart of Frege and Hilbert's famous disagreement about the status of model theory (see [1] for an interesting discussion of this disagreement).

than several influential commentators have thought. Tarski does regard number-theoretic terminology as logical, but his reasons for doing so are more complicated—and *far* more plausible—than they are often taken to be. Nor, as I have shown in sections 5 and 6, does Tarski's understanding of ω-inferences run aground on the more obvious objections that might be raised against it. Finally, I have argued that examination of Tarski's ω-inference example provides further support for the claim that Tarski intended to introduce a fixed-domain conception of models. Although adopting a fixed-domain conception of models does not *automatically* solve the problems associated with Tarski's ω-inference example, the interpretation which *does* solve these problems depends heavily on just such a conception.

References

[1] Patricia Blanchette. Frege and Hilbert on consistency. *The Journal of Philosophy*, 93:317–336, 1996.

[2] Rudolf Carnap. *Einführung in die Symbolische Logik*. Springer, Vienna, 1954.

[3] John Echemendy. Tarski on truth and logical consequence. *The Journal of Symbolic Logic*, 53:51–79, 1988.

[4] John Echemendy. *The Concept of Logical Consequence*. Harvard, Cambridge, 1990.

[5] Kurt Gödel. On formally undecidable propositions of *Principia Mathematica* and related systems. In *Collected Works*, pages 145–199. Oxford University Press, New York, 1986.

[6] Mario Gómez-Torrente. Tarski on logical consequence. *Notre Dame Journal of Formal Logic*, 37:125–51, 1996.

[7] Ray Greg. Logical consequence: A defense of Tarski. *The Journal of Philosophical Logic*, 25:617–677, 1996.

[8] David Hilbert and Wilhelm Ackermann. *Principles of Mathematical Logic*. Chelsea Publishing, New York, 1937.

[9] Wilfrid Hodges. Truth in a structure. *Proceedings of the Aristotelian Society*, 86:135–151, 1985–86.

[10] Bertrand Russell. *Introduction to Mathematical Philosophy*. MacMillan, New York, 1919.

[11] Gila Sher. A conception of Tarskian logic. *Pacific Philosophical Quarterly*, 70:341–69, 1989.

[12] Gila Sher. *The Bounds of Logic: A Generalized Viewpoint*. MIT Press, Cambridge, 1991.

[13] Gila Sher. Did Tarski commit 'Tarski's Fallacy'. *The Journal of Symbolic Logic*, 61:653–686, 1996.

[14] Thoralf Skolem. Some remarks on axiomitized set theory. In van Heijenoort [27], pages 290–301.

[15] Thoralf Skolem. Une relativisation des notions mathematiques fondementales. In *Selected Works in Logic*, pages 587–600. Uiversitetsforlaget, Oslo, 1970.

[16] Alfred Tarski. *Einführung in die mathematische Logik und in die Methodologie der Mathematik*. Springer, Vienna, 1937.

[17] Alfred Tarski. *Introduction to Logic and to the Methodology of Deductive Sciences*. Oxford, New York, 1941. Translation of [16].

[18] Alfred Tarski. The concept of truth in formalized languages. In *Logic, Semantics, and Metamathematics* [20], pages 152–278.

[19] Alfred Tarski. Foundations of the calculus of systems. In *Logic, Semantics, and Metamathematics* [20], pages 342–383.

[20] Alfred Tarski. *Logic, Semantics, Metamathematics*. Hackett, Indianapolis, 2nd edition, 1983.

[21] Alfred Tarski. On the concept of logical consequence. In *Logic, Semantics, and Metamathematics* [20], pages 409–20.

[22] Alfred Tarski. On the foundations of Boolean Algebra. In *Logic, Semantics, and Metamathematics* [20], pages 320–341.

[23] Alfred Tarski. On the limitations of the means of expression of deductive theories. In *Logic, Semantics, and Metamathematics* [20], pages 284–392.

[24] Alfred Tarski. Some methodological investigations on the definability of concepts. In *Logic, Semantics, and Metamathematics* [20], pages 296–319.

[25] Alfred Tarski. Some observations on the concepts of ω-consistency and ω-completeness. In *Logic, Semantics, and Metamathematics* [20], pages 278–295.

[26] Alfred Tarski. What are the logical notions? *History and Philosophy of Logic*, 7:443–154, 1986.

[27] Jean van Heijenoort, editor. *From Frege to Gödel*. Harvard, Cambridge, Mass., 1967.