Notre Dame Journal of Formal Logic
Volume 59, Number 2, 2018

Actualism, Serious Actualism, and Quantified Modal Logic

William H. Hanson

Abstract This article studies seriously actualistic quantified modal logics. A key component of the language is an abstraction operator by means of which predicates can be created out of complex formulas. This facilitates proof of a uniform substitution theorem: if a sentence is logically true, then any sentence that results from substituting a (perhaps complex) predicate abstract for each occurrence of a simple predicate abstract is also logically true. This solves a problem identified by Kripke early in the modern semantic study of quantified modal logic. A tableau proof system is presented and proved sound and complete with respect to logical truth. The main focus is on seriously actualistic T (SAT), an extension of T, but the results established hold also for systems based on other propositional modal logics (e.g., K, B, S4, and S5). Following Menzel it is shown that the formal language studied also supports an actualistic account of truth simpliciter.

1 Preliminaries

How should we deal with subject-predicate sentences containing nondenoting singular terms? Suppose "Dweet" is a name that does not denote anything, and consider the following sentences.

(1) Dweet is a sexagenarian.
(2) Dweet enjoys messing about in boats.

It seems perfectly plausible to construe "is a sexagenarian" and "enjoys messing about in boats" as predicates, and thus to take the sentences in question to be of subject-predicate form. I will call such sentences—those that result from applying

Received November 11, 2014; accepted February 5, 2015
First published online November 17, 2017
2010 Mathematics Subject Classification: 03B45
Keywords: quantified modal logic, serious actualism, actualism, predicate abstraction, uniform substitution, tableau proof system
 10.1215/00294527-2017-0022

a predicate to a singular term—*attributive sentences*. I will also use this term for sentences like (3) that apply a relational predicate to two or more singular terms.

(3) Obama is taller than Dweet.

How should truth values be assigned to attributive sentences like (1)–(3) that contain at least one nondenoting singular term?

One answer to this question is that such sentences have no truth values. Another is that they have truth values, some of them being true and others false, just like ordinary attributive sentences. A third answer, the one I advocate and want to explore here, is that all such attributive sentences are false. Indeed they are necessarily false. If there is no such thing as Dweet, then no claim to the effect that Dweet has such and such properties (or stands in such and such relations) could be true, and all sentences that express such claims are necessarily false. This view has been called *serious actualism*.[1]

Intuitively, serious actualism provides a model for understanding ordinary language that is at least as plausible as those provided by the other two alternatives. And when properly developed it has advantages, both technical and philosophical, over them. Or so I claim. Briefly put, serious actualism is superior to the no-truth-value view because it does not require that we abandon two-valued logic. It is superior to the some-are-true-some-are-false alternative because it is less likely to tempt us with metaphysical excess, and because it forces us to draw a sharp line between the logic of fiction and the logic of nonfictional discourse. More fundamentally, serious actualism reflects a basic intuition about objects and predication better than the other two approaches.[2]

I also advocate *actualism*, the view that everything that is, in any coherent sense of "is," is actual. According to actualism there are no merely possible things, things that do not exist but somehow manage to subsist or have some alternative kind of being. Nondenoting names like "Dweet" and "Pegasus" really are nondenoting. They do not denote, designate, or refer in any sense.[3]

In what follows I develop a system of quantified modal logic that is both seriously actualistic and actualistic. A major goal is to extend the thesis of serious actualism to complex predicates without giving up desirable logical properties, the most important of which is *uniform substitution*. Substituting a complex predicate for each occurrence of a simple predicate in a logical truth should yield a logical truth.

Any seriously actualistic attempt to deal with sentences more complex than simple attributive sentences faces two problems. One is easily solved; the solution to other is more difficult. The first problem is that in standard modern logic, connectives and quantifiers are usually treated only as devices for generating more complex sentences out of simpler ones. But in natural language they play a second role as well, allowing us to generate more complex predicates and relational terms out of simpler ones. This point is easily illustrated, using negation, with the following sentence.

(4) Dweet is not a sexagenarian.

Is (4) an attributive sentence, with "is not a sexagenarian" as its predicate, or is it simply the negation of the attributive sentence (1)? If we choose the former answer, we seem to be multiplying primitive predicates needlessly. But if we choose the latter, we are forced to say that (4) is true, since it negates (1), which is false. Yet this seems arbitrary, because "is not a sexagenarian" appears to have as good a claim to being a predicate as "is a sexagenarian." And if (4) is an attributive sentence, then under the seriously actualistic approach to nondesignating singular terms it is false, not true.

This problem is easily solved by adopting a suitable device for distinguishing the scopes of negation and other logical operators. We want to distinguish between

(4a) Dweet is a non-sexagenarian.

and

(4b) It is not the case that Dweet is a sexagenarian.

In the formal language developed in later sections this distinction will be marked using a predicate abstraction operator (λ), that allows us to render (4a) and (4b) as follows:

(4a*) $\langle \lambda x.\neg Sx\rangle(d)$,
(4b*) $\neg\langle \lambda x.Sx\rangle(d)$.

In (4b*) the predicate $\langle \lambda x.Sx\rangle$ (is a sexagenarian) is applied to the singular term d (Dweet), and the resulting sentence is negated. But in (4a*) negation is involved in forming the predicate $\langle \lambda x.\neg Sx\rangle$ (is not a sexagenarian). Predicate abstraction allows us to control the scope of negation and thereby distinguish two senses of (4), using S as the only primitive predicate.

Augmenting modal languages with predicate abstraction is not a new idea, but its usefulness has not been fully exploited. The first-order modal system developed here contains a predicate abstraction operator, and it embodies serious actualism in the manner just described. I will focus on the version of this system based on the propositional modal logic T, which I call SAT (seriously actualistic T), but the results obtained apply to languages based on several other propositional modal logics.[4]

The second and more difficult problem facing the serious actualist in developing a system worthy of being called a logic involves uniform substitution. The result of substituting a complex predicate for each occurrence of a simple predicate in a logically true sentence should also be logically true. I show in Section 5.1 that when complex predicates are defined using predicate abstraction, uniform substitution does indeed preserve logical truth in SAT. But none of the work on serious actualism of which I am aware deals with this matter fully and satisfactorily.[5]

It is hard to overemphasize the importance of uniform substitution. There is nearly universal agreement among logicians that a formal system is not a logic unless it respects logical form in this way. The reason that seriously actualistic modal systems have not been more widely studied is probably that they appear not to support uniform substitution. But this is only because attention has been focused on languages that lack a means of properly handling predication. In such languages uniform substitution does indeed fail to preserve logical truth. This was noticed by Kripke at the very beginning of the modern semantic study of modal logic:

> It is natural to assume that an *atomic* predicate should be *false* in a world H of all those individuals not existing in that world; that is, that the extension of a predicate letter must consist of actually existing individuals. [. . .] We have chosen not to do this because the rule of substitution would no longer hold; theorems would hold for atomic formulae which would not hold when the atomic formulae are replaced by arbitrary formulae. (This answers a question of Putnam and Kalmar.)[6]

Kripke's insight can be illustrated using sentences that express existential generalization, such as

(5) $Fa \rightarrow \exists x Fx$

and

(5*) $\neg Fa \rightarrow \exists x \neg Fx$.

If we adopt the view that Kripke says "is natural to assume," (5) is valid but (5*) is not. If a does not denote at a world and the extension of F is the entire domain of that world, $\neg Fa$ is true but $\exists x \neg Fx$ is false. So if we make the natural assumption that primitive predicates behave in accordance with serious actualism, we must accept the unnatural result that complex predicates do not. Uniform substitution fails.

In Section 5.1, however, I show that when a complex predicate abstract replaces each occurrence of a primitive predicate abstract in a sentence, logical truth is preserved. So if $\phi(y)$ is any formula with free occurrences of y (but no free occurrence of any other variable), every instance of

(6) $\langle \lambda y.\phi(y)\rangle(a) \rightarrow \exists x \langle \lambda y.\phi(y)\rangle(x)$

is valid.

Just as predicate abstraction facilitates a plausible account of serious actualism, so an actuality connective (A) facilitates the expression of actualism itself.[7] Letting $\mathcal{E}(x)$ abbreviate $\exists y(y = x)$, actualism can be formalized as

(7) $\forall x \mathsf{A}\mathcal{E}(x)$.

Although (7) is not valid, it is true. That is, it is true at the actual world element of the intended model of SAT, which is defined in Section 7. Indeed (7) is knowable a priori. For I know, independently of experience, that I and all my surroundings are part of the actual world. Hence I know that everything actually exists.[8] There is a tableau for (7) in Section 8.3.7 that illustrates this point.

It is also worth noting that since SAT contains A it is able to capture the content of sentences that cannot otherwise be expressed in first-order modal languages. An example is

(8) There might have been something that does not actually exist,

which can be rendered

(9) $\Diamond \exists x \neg \mathsf{A}\mathcal{E}(x)$.

Without an actuality connective, there is no way to express (8) in a first-order modal language.[9]

Several other features of SAT should be noted before delving into the details. SAT is a version of free logic and as such has three important features. First, as has already been stated, it makes all attributive sentences with nondenoting singular terms false. Hence, second, unrestricted universal instantiation is not valid. One cannot infer the formal analogue of (2) from that of

(10) Everyone enjoys messing about in boats.

Third, SAT allows models in which quantifiers range over the empty domain. So no existentially quantified sentence is a logical truth, and all inferences from a universal sentence to the corresponding existential sentence are invalid.[10]

To summarize, and to provide some additional information as a guide to the details that follow, SAT has the following noteworthy features.

- The object language contains the usual symbols of a first-order modal language plus a predicate abstraction operator (λ) and an actuality connective (A).

- Names are nonrigid. SAT is thus a contingent identity system. An identity statement containing two names may be true at some worlds and false at others.
- A name need not denote at each world. Indeed a name need not denote at any world.
- If a name denotes at a world, it denotes an object that exists in that world. This makes an actualistic interpretation possible. Indeed the sentence $\forall x \mathsf{A}\mathcal{E}(x)$, which expresses the core tenet of actualism, is true in the intended model.
- An atomic sentence is false at a world unless all the terms it contains denote objects in the domain of that world. Since identity is a primitive predicate, this holds for identity sentences. And since names need not denote, even a self-identity sentence can be false at a world. Thus no self-identity sentence is a logical truth.
- Far from being a flaw, the fact that self-identity sentences are not logically true is an advantage. For it avoids the result that $\exists y \Box (y = a)$ is a truth of logic without artificially restricting necessitation or existential generalization.
- The extension of a predicate abstract at a world, like that of a primitive predicate, is restricted to objects that exist in that world. Thus SAT embodies serious actualism. All sentences of the form $\Box \forall x \Box (\langle \lambda y.\phi \rangle(x) \to \mathcal{E}(x))$ are logical truths.
- There are models in which some or all worlds have empty domains. (In the latter case the domain of the entire model is also empty.)
- The actual world element of a model plays no special role in the definitions of validity and logical consequence. Validity is thus general validity, truth at every world in every model.

In Sections 2–4 I present the syntax and semantics of SAT. Section 5 contains important semantic metatheorems, including uniform substitution and replacement (substitution of logically equivalent subformulas). Section 6 compares SAT with the work of others, and Section 7 presents an actualistic account of truth. Section 8 contains semantic tableau proof rules and some examples of their application. Section 9 is a brief conclusion. In the Appendix I prove the tableau proof rules sound and (weakly) complete with respect to validity as defined for SAT. These proofs extend straightforwardly to several related systems.

2 Syntax

The language under consideration, $\mathfrak{L}$, is a first-order modal language supplemented with a predicate abstraction operator and an actuality connective.

2.1 Symbols The nonlogical symbols of $\mathfrak{L}$ are *individual constants* (names) and *predicate constants* (predicates). The former are the lower case letters a through j; the latter are the upper case letters A through Z. The logical symbols are the *variables*, which are the lower case letters u through z, plus the following symbols

$$\neg \quad \wedge \quad \vee \quad \to \quad \leftrightarrow \quad \Box \quad \Diamond \quad \forall \quad \exists \quad = \quad \mathsf{A} \quad \lambda \quad . \quad \langle \quad \rangle \quad (\quad)$$

The first ten of these are the usual truth-functional and modal connectives, quantifiers, and the identity predicate. The next two are the actuality connective and the abstraction operator. The remaining five are punctuation marks. The term *term* applies to both individual constants and variables.[11]

2.2 Formulas and sentences Formulas and sentences of $\mathfrak{L}$ are as usual. An *atomic formula* is either an n-ary predicate followed by n terms or the identity sign flanked by two terms (the whole being enclosed in parentheses). All occurrences of variables in atomic formulas are *free occurrences*. Thus the following are atomic formulas

$$Bx \quad Gxy \quad (z = c) \quad Xzyz \quad Gyb \quad (d = h) \quad Xaab \quad (b = b),$$

and all the occurrences of variables in them are free. The notions of *formula*, *free occurrence of a variable in a formula*, and *predicate abstract* are defined simultaneously, as follows.

(1) Every atomic formula is a formula, and every occurrence of a variable in an atomic formula is a free occurrence.
(2) If ϕ is a formula, so are $\neg\phi$, $\Box\phi$, $\Diamond\phi$, and $\mathsf{A}\phi$. The free occurrences of variables in these formulas are the same as those in ϕ.
(3) If ϕ and ψ are formulas, so are $(\phi \wedge \psi)$, $(\phi \vee \psi)$, $(\phi \to \psi)$, and $(\phi \leftrightarrow \psi)$. The free occurrences of variables in these formulas are the same as those in ϕ together with those in ψ.
(4) If ϕ is a formula, α is a variable, and α has at least one free occurrence in ϕ, then $\forall\alpha\phi$ and $\exists\alpha\phi$ are formulas. The free occurrences of variables in these formulas are the same as those in ϕ, except for occurrences of α.[12]
(5) If ϕ is a formula, α is a variable, and α has at least one free occurrence in ϕ, then $\langle\lambda\alpha.\phi\rangle$ is a *predicate abstract*. The free occurrences of variables in $\langle\lambda\alpha.\phi\rangle$ are the same as those in ϕ, except for occurrences of α.
(6) If $\langle\lambda\alpha.\phi\rangle$ is a predicate abstract and τ is a term, $\langle\lambda\alpha.\phi\rangle(\tau)$ is a formula. The free occurrences of variables in $\langle\lambda\alpha.\phi\rangle(\tau)$ are the same as those in $\langle\lambda\alpha.\phi\rangle$ together with those in τ.[13]

Any occurrence of a variable in a formula that is not free is *bound*. A formula in which all occurrences of variables are bound is a *sentence*.

It will be convenient, in what follows, to have a way of abbreviating formulas involving iterated predicate abstraction. Consider a formula of the form

$$\big\langle\lambda\alpha_1.\big\langle\lambda\alpha_2.\langle\lambda\alpha_3.\phi\rangle(\tau_3)\big\rangle(\tau_2)\big\rangle(\tau_1),$$

where $\alpha_1, \alpha_2, \alpha_3$ are variables, τ_1, τ_2, τ_3 are terms, and ϕ is a formula. Such formulas will be abbreviated as

$$\langle\lambda\alpha_1, \alpha_2, \alpha_3.\phi\rangle(\tau_1, \tau_2, \tau_3).$$

Similar abbreviations will be used for formulas containing different numbers of iterated predicate abstracts.

3 Semantics

The logic I present in this paper is a seriously actualistic quantified modal logic whose underlying propositional modal logic is the system T. Hence its name, SAT. Its semantic framework will be familiar to those who know Kripke-style modal semantics, and especially familiar to those who know Fitting and Mendelsohn's [7] presentation of it. Indeed only small changes in Fitting and Mendelsohn's definitions of basic semantic notions are sufficient to yield SAT.[14] But small differences in semantic details can make a big difference in the resulting logic, and this is the case with SAT.

Although in what follows I confine attention to SAT, similar proof techniques are readily available, and similar metatheorems hold, when the underlying modal propositional logic is changed to K, B, S4, or S5. Especially noteworthy among these metatheorems are uniform substitution and substitutivity of equivalents (see Section 5) and soundness and completeness of the tableau rules (see Appendix). I call the systems in question SAK, SAB, SAS4, and SAS5. In what follows I will periodically remind the reader that what is being proved applies also to them.

The fundamental semantic notions are those of *model*, *interpretation*, *valuation*, *satisfaction*, *truth*, *validity*, and *consequence*.

3.1 Models A *model* is a quintuple $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$, where $\mathcal{W}$ is a nonempty set, $@ \in \mathcal{W}$, $\mathcal{R}$ is a reflexive binary relation on $\mathcal{W}$, $\mathcal{D}$ is a function from members of $\mathcal{W}$ to (possibly empty) sets (all of which are disjoint from $\mathcal{W}$), and $\mathcal{I}$ assigns extensions to nonlogical symbols of the language at members of $\mathcal{W}$. The domain $\mathcal{D}_{\mathcal{M}} = \bigcup\{\mathcal{D}(w_i) \mid w_i \in \mathcal{W}\}$, the union of the sets of individuals assigned to members of $\mathcal{W}$ by $\mathcal{D}$, is called the *domain of the model* $\mathcal{M}$. (Notice that this definition allows $\mathcal{D}_{\mathcal{M}}$ to be empty.) The possibilist's way of giving intuitive content to these definitions is to say that $\mathcal{W}$ is a set of worlds, @ is the actual world, $\mathcal{R}$ is the relative possibility relation among worlds, $\mathcal{D}$ assigns to each world the (possibly empty) set of individuals that exist in it, $\mathcal{I}$ is an interpretation of the nonlogical symbols of the language, and $\mathcal{D}_{\mathcal{M}}$ is the set of all actual and possible individuals. I argue in Section 7 that models can be understood in a purely actualistic way.

3.2 Interpretations An interpretation $\mathcal{I}$ is a function that assigns individuals to individual constants and sets of ordered n-tuples of individuals to n-ary predicates, both assignments being relative to a member of $\mathcal{W}$. Note that $\mathcal{I}$ is not required to assign anything to an individual constant at a world, but if it does, the individual so assigned must exist at that world. Similarly, $\mathcal{I}$ requires the extension of an n-ary predicate at a world to contain only n-tuples of individuals that exist at that world. This holds for all predicates, including the identity predicate. (So if the individual constant d fails to denote at a world, even $(d = d)$ is false at that world.) It is these two features of $\mathcal{I}$ that make the object language actualistic and seriously actualistic, respectively, and thereby distinguish it from most other quantified modal logics.

More precisely, an *interpretation* $\mathcal{I}$ of a model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$ is the union of a (perhaps partial) function on individual constants and members of $\mathcal{W}$ and a total function on predicate constants and members of $\mathcal{W}$ such that:

(1) For an individual constant ι, and a world $w \in \mathcal{W}$, if $\mathcal{I}(\iota, w)$ is defined, $\mathcal{I}(\iota, w) \in \mathcal{D}(w)$.[15]

(2) For each n-ary predicate θ and each $w \in \mathcal{W}$, $\mathcal{I}(\theta, w)$ is a set of ordered n-tuples of elements of $\mathcal{D}(w)$. Specifically, if θ is $=$, $\mathcal{I}(\theta, w)$ is the identity relation on $\mathcal{D}(w)$.[16]

These two clauses embody actualism and serious actualism, respectively.[17]

3.3 Valuations of variables: Designations of terms Three more semantic functions must be defined before I can define satisfaction.

First, a *valuation* $\mathcal{V}$ relative to a model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$ is a (total) function from the set of variables of the language into $\mathcal{D}_{\mathcal{M}}$. (Thus if $\mathcal{D}_{\mathcal{M}}$ is empty, $\mathcal{V}$ is the null set.)

Next, where $\mathcal{V}$ and $\mathcal{U}$ are valuations relative to a model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$, $w \in \mathcal{W}$, α is a variable, and $\mathcal{D}(w)$ is not empty, $\mathcal{U}$ is an *α-variant-at-w* of $\mathcal{V}$ if

$\mathcal{V}$ and $\mathcal{U}$ differ at most in what they assign to α, and $\mathcal{U}(\alpha) \in \mathcal{D}(w)$. (I stipulate that if $\mathcal{D}(w)$ is empty, no valuation (not even $\mathcal{V}$ itself) is an α-variant-at-w of $\mathcal{V}$. In this case universally quantified formulas are satisfied by all valuations at w and existentially quantified formulas are satisfied by none; see Section 3.4.)

Finally, I define the *designation of a term* at a world of a model relative to a valuation. Recall that a term may be either a variable or an individual constant. Since it is models (via their constituent interpretations) that assign designations to individual constants, but valuations that assign designations to variables, it will be convenient to have a single notation that expresses the designation of a term. For this purpose I adopt the notation $(\mathcal{V} \star \mathcal{I})(\tau, w)$ of Fitting and Mendelsohn (where $\mathcal{V}$ is a valuation relative to a model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$, $w \in \mathcal{W}$, and τ is a term).

(1) If τ is a variable, $(\mathcal{V} \star \mathcal{I})(\tau, w) = \mathcal{V}(\tau)$.
(2) If τ is an individual constant, and $\mathcal{I}(\tau, w)$ is defined, $(\mathcal{V} \star \mathcal{I})(\tau, w) = \mathcal{I}(\tau, w)$. Otherwise $(\mathcal{V} \star \mathcal{I})(\tau, w)$ is undefined.

Note that under these definitions variables designate rigidly, but individual constants are allowed to designate nonrigidly.

3.4 Satisfaction and satisfiability The definition of *satisfaction* is now straightforward. First, abbreviate valuation $\mathcal{V}$ *satisfies* formula ϕ at world w of model $\mathcal{M}$ as $\mathcal{M}_{\mathcal{V}}(\phi, w) = 1$, and valuation $\mathcal{V}$ *does not satisfy* formula ϕ at world w of model $\mathcal{M}$ as $\mathcal{M}_{\mathcal{V}}(\phi, w) = 0$. Satisfaction of a formula ϕ relative to a model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$, a valuation $\mathcal{V}$, and a world $w \in W$ (where θ is an n-ary predicate, α is a variable, τ and $\tau_1, \ldots, \tau_n$ are terms, and ϕ, χ, and ψ are formulas) is defined as follows.

(1) If ϕ is the atomic formula $\theta\tau_1 \ldots \tau_n$, $\mathcal{M}_{\mathcal{V}}(\phi, w) = 1$ if $\langle (\mathcal{V} \star \mathcal{I})(\tau_1, w), \ldots, (\mathcal{V} \star \mathcal{I})(\tau_n, w) \rangle \in \mathcal{I}(\theta, w)$; otherwise $\mathcal{M}_{\mathcal{V}}(\phi, w) = 0$.[18]
(2) If ϕ is $\neg\psi$, $\mathcal{M}_{\mathcal{V}}(\phi, w) = 1$ if $\mathcal{M}_{\mathcal{V}}(\psi, w) = 0$; otherwise $\mathcal{M}_{\mathcal{V}}(\phi, w) = 0$.
(3) If ϕ is $(\psi \wedge \chi)$, $\mathcal{M}_{\mathcal{V}}(\phi, w) = 1$ if $\mathcal{M}_{\mathcal{V}}(\psi, w) = 1$ and $\mathcal{M}_{\mathcal{V}}(\chi, w) = 1$; otherwise $\mathcal{M}_{\mathcal{V}}(\phi, w) = 0$.
(4)–(6) The clauses for $\vee$, $\rightarrow$, and $\leftrightarrow$ are similar (with obvious modifications) to clause 3.
(7) If ϕ is $\Box\psi$, $\mathcal{M}_{\mathcal{V}}(\phi, w) = 1$ if $\mathcal{M}_{\mathcal{V}}(\psi, w^*) = 1$ for all $w^* \in \mathcal{W}$ such that $w\mathcal{R}w^*$; otherwise $\mathcal{M}_{\mathcal{V}}(\phi, w) = 0$.
(8) If ϕ is $\Diamond\psi$, $\mathcal{M}_{\mathcal{V}}(\phi, w) = 1$ if $\mathcal{M}_{\mathcal{V}}(\psi, w^*) = 1$ for at least one $w^* \in \mathcal{W}$ such that $w\mathcal{R}w^*$; otherwise $\mathcal{M}_{\mathcal{V}}(\phi, w) = 0$.
(9) If ϕ is $\mathsf{A}\psi$, $\mathcal{M}_{\mathcal{V}}(\phi, w) = 1$ if $\mathcal{M}_{\mathcal{V}}(\psi, @) = 1$; otherwise $\mathcal{M}_{\mathcal{V}}(\phi, w) = 0$.
(10) If ϕ is $\forall\alpha\psi$, $\mathcal{M}_{\mathcal{V}}(\phi, w) = 1$ if $\mathcal{M}_{\mathcal{U}}(\psi, w) = 1$ for all valuations $\mathcal{U}$ that are α-variants-at-w of $\mathcal{V}$; otherwise $\mathcal{M}_{\mathcal{V}}(\phi, w) = 0$.
(11) If ϕ is $\exists\alpha\psi$, $\mathcal{M}_{\mathcal{V}}(\phi, w) = 1$ if $\mathcal{M}_{\mathcal{U}}(\psi, w) = 1$ for at least one valuation $\mathcal{U}$ that is an α-variant-at-w of $\mathcal{V}$; otherwise $\mathcal{M}_{\mathcal{V}}(\phi, w) = 0$.
(12a) If ϕ is $\langle\lambda\alpha.\psi\rangle(\tau)$, $(\mathcal{V} \star \mathcal{I})(\tau, w)$ is defined, and $(\mathcal{V} \star \mathcal{I})(\tau, w) \in \mathcal{D}(w)$, $\mathcal{M}_{\mathcal{V}}(\phi, w) = 1$ if $\mathcal{M}_{\mathcal{U}}(\psi, w) = 1$, where $\mathcal{U}$ is the α-variant-at-w of $\mathcal{V}$ such that $\mathcal{U}(\alpha) = (\mathcal{V} \star \mathcal{I})(\tau, w)$; otherwise $\mathcal{M}_{\mathcal{V}}(\phi, w) = 0$.
(12b) If ϕ is $\langle\lambda\alpha.\psi\rangle(\tau)$ and either $(\mathcal{V} \star \mathcal{I})(\tau, w)$ is undefined or $(\mathcal{V} \star \mathcal{I})(\tau, w) \notin \mathcal{D}(w)$, $\mathcal{M}_{\mathcal{V}}(\phi, w) = 0$.

A set of formulas Σ is *satisfiable* if and only if there is a model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$, a valuation $\mathcal{V}$ relative to $\mathcal{M}$, and a world $w \in \mathcal{W}$, such that for each formula $\sigma_i \in \Sigma$, $\mathcal{M}_{\mathcal{V}}(\sigma_i, w) = 1$. A formula ϕ is *satisfiable* if and only if $\{\phi\}$ is satisfiable.

3.5 Truth at a world in a model; Truth in a model As in the standard Tarskian approach to semantics, a sentence of SAT is satisfied by a valuation at a world in a model just in case it is satisfied by all such valuations. Truth and falsity for sentences can thus be defined in the usual way. A sentence ϕ is *true at world* w of model $\mathcal{M}$ if and only if all valuations $\mathcal{V}$ relative to $\mathcal{M}$ satisfy ϕ at w. Similarly, ϕ is *false at world* w of $\mathcal{M}$ if and only if no valuation $\mathcal{V}$ relative to $\mathcal{M}$ satisfies ϕ at w. A sentence ϕ is *true in model* $\mathcal{M}$ if and only if ϕ is true at world @ of $\mathcal{M}$.

Notice that truth and falsity are defined only for sentences. This is deliberate. Extending these definitions so that they apply to all formulas would not be a harmless change. It would be tantamount to introducing *possibilist* universal quantifiers (quantifiers that range over the domain of a model) although allowing them to stand only at the beginnings of sentences. Such quantifiers flout the fundamental principle of actualism. And defining truth in this way can lead to serious confusion, as was shown by Kripke [19] in his refutation of an alleged proof of the Converse Barcan Formula given by Prior.[19]

3.6 Validity, consequence, and equivalence Definitions of validity, logical consequence, and logical equivalence, again only for sentences, can now be given in the standard way. (These definitions apply to SAK, SAB, SAS4, and SAS5, as well as to SAT.) A sentence ϕ is *valid* (abbreviated $\vDash \phi$) if and only if it is true at every world of every model. A sentence ϕ is a *logical consequence* of a set of sentences Γ (abbreviated $\Gamma \vDash \phi$) if and only if for every world in every model, if each member of Γ is true at that world so is ϕ. Two sentences ϕ and ϕ' are *logically equivalent* if and only if $\vDash (\phi \leftrightarrow \phi')$. It is easy to verify that the rule of necessitation (if $\vDash \phi$, then $\vDash \Box\phi$) holds.

The notion of validity just defined (truth at each world of each model) is called *general validity*. It should not be confused with the weaker notion of *real-world validity*, which requires only truth at the actual-world element, @, of each model. Elsewhere I have argued that general validity better captures our intuitive notion of logical truth.[20]

4 Notable Features of the Semantics

4.1 Attributive sentences We are now in a position to see how SAT handles the examples with which we began in Section 1. Consider a model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$ and a particular world $w \in \mathcal{W}$ such that $\mathcal{D}(w)$ is not empty. For predicate constants S and N, and individual constant d, suppose that $\mathcal{I}(S, w)$ and $\mathcal{I}(N, w)$ are nonempty subsets of $\mathcal{D}(w)$, $\mathcal{I}(N, w)$ is the complement of $\mathcal{I}(S, w)$ with respect to $\mathcal{D}(w)$, and $\mathcal{I}(d, w) \notin \mathcal{D}(w)$. Then $\langle \lambda x.\neg Sx \rangle(d)$ and $\langle \lambda x.Nx \rangle(d)$ are both false at w while $\neg\langle \lambda x.Sx \rangle(d)$ and $\neg\langle \lambda x.\neg Nx \rangle(d)$ are both true. If we think of S, N, and d as "is a sexagenarian," "is a non-sexagenarian," and "Dweet," respectively, and w as the world of interest, we have formal representations of sentences (4a) and (4b) of the Introduction. Since Dweet does not exist at w, all attributive sentences involving "Dweet" are false, and the denials of all such sentences are true.

Without predicate abstraction we would need both S and N as primitive predicates to symbolize sentences (4a) and (4b). For Nd is false at w while $\neg Sd$ is true. Predicate abstraction lets the serious actualist say what he wants to say, whether a predicate or its complement is taken as primitive.[21]

4.2 Identity Since individual constants may denote different objects at different worlds (and need not denote anything at some or even any worlds), SAT allows identity sentences to be true at some worlds and false at others.[22] It is thus a contingent identity logic. Adding this feature to serious actualism entails that self-identity sentences are false at worlds in which the individual constant involved does not denote. And as the following paragraphs show, this in turn allows us to avoid having as theorems of logic statements to the effect that ι necessarily exists, for every individual constant ι.[23]

To facilitate exposition, let $\mathcal{E}(\tau)$ abbreviate $\exists\alpha(\alpha = \tau)$, where τ is a term and α is a variable distinct from τ. Clearly $\mathcal{E}(\tau)$ is satisfied by a valuation at a world just in case the object denoted by τ is a member of the domain of that world. So $\mathcal{E}$ can be thought of as an existence predicate.

Every sentence of the form

(Id) $\forall\alpha(\alpha = \alpha)$

is valid, but claims of self-identity that make use of individual constants are not. Thus neither $d = d$ nor $\langle\lambda x.(x = x)\rangle(d))$ is valid. Indeed both are false at world w of the model $\mathcal{M}$ of Section 4.1. And while all sentences of the forms

(□Id1) $\forall\alpha(\alpha = \alpha \to \Box(\mathcal{E}(\alpha) \to (\alpha = \alpha)))$,
(□Id2) $\forall\alpha_1\forall\alpha_2(\alpha_1 = \alpha_2 \to \Box(\mathcal{E}(\alpha_1) \to (\alpha_1 = \alpha_2)))$,

and

(ND) $\forall\alpha_1\forall\alpha_2(\alpha_1 \neq \alpha_2 \to \Box(\alpha_1 \neq \alpha_2))$

are valid, no sentence of the form $\forall\alpha(\alpha = \alpha \to \Box(\alpha = \alpha))$, $\forall\alpha_1\forall\alpha_2(\alpha_1 = \alpha_2 \to \Box(\alpha_1 = \alpha_2))$, or $\forall\alpha\Box(\alpha = \alpha)$ is.[24]

It is also noteworthy that existence claims and claims of self-identity are logically equivalent. That is, all instances of

$$\mathcal{E}(\iota) \leftrightarrow (\iota = \iota)$$

are valid. And since the rule of necessitation preserves validity, if $d = d$ were valid, then $\mathcal{E}(d)$ and $\Box\mathcal{E}(d)$—abbreviations of $\exists x(x = d)$ and $\Box\exists x(x = d)$—would also be valid. Of course the same would hold for every individual constant. But surely we do not want to have all, or any, sentences of the form

(NE) $\Box\mathcal{E}(\iota)$

as theorems of logic. The consistent serious actualist avoids this problem by unflinchingly applying her basic principle to self-identities. No sentence of the form

(SI) $\iota = \iota$

is a logical truth.[25]

Failure to take this principled step is not only inconsistent with serious actualism, it leads to other difficulties. In Menzel's [24] painstaking reconstruction of Prior's seriously actualistic modal logic, all sentences of the form (SI) are valid, but no sentence of the form (NE) is. Menzel achieves this result only by denying that the rule of necessitation applies to (SI) or to valid sentences that depend on it.[26] But this seems ad hoc.

So in SAT statements of self-identity and tautologies are not logically equivalent. While in classical logic $d = d$ and $Sd \vee \neg Sd$ are both logical truths, in SAT only the latter is. Similarly $\forall x(Sx \vee \neg Sx) \rightarrow (Sd \vee \neg Sd)$ is valid, but $\forall x(x = x) \rightarrow (d = d)$ is not.[27] The proper comparison is between $\langle \lambda y.(y = y)\rangle(d)$ and $\langle \lambda y.(Sy \vee \neg Sy)\rangle(d)$, which are logically equivalent though not valid.

With respect to substitutivity of identicals, all sentences of the forms

(Sub1) $(\iota_1 = \iota_2) \rightarrow (\langle \lambda\alpha.\phi(\alpha)\rangle(\iota_1) \rightarrow \langle \lambda\alpha.\phi(\alpha)\rangle(\iota_2))$

and

(Sub2) $\forall\alpha_1\forall\alpha_2(\alpha_1 = \alpha_2 \rightarrow (\langle \lambda\alpha_3.\phi(\alpha_3)\rangle(\alpha_1) \rightarrow \langle \lambda\alpha_3.\phi(\alpha_3)\rangle(\alpha_2)))$

are valid, but not all sentences of the form $((\iota_1 = \iota_2) \wedge \phi(\iota_2)) \rightarrow \phi(\iota_1)$ are. There are thus instances of the corresponding argument form in which the conclusion is not a consequence of the premises.[28]

4.3 Quantifiers Quantification is world-relative in SAT, and primitive predicates and predicate abstracts are true at a world only of objects that exist at that world. Because of these two features some of the classical principles involving quantifiers hold only in restricted forms.

Existential generalization in the form

(EG) $\langle \lambda\alpha_2.\psi(\alpha_2)\rangle(\iota) \rightarrow \exists\alpha_1\langle \lambda\alpha_2.\psi(\alpha_2)\rangle(\alpha_1)$

is valid. Thus the sentences $\langle \lambda x.Nx\rangle(d) \rightarrow \exists y\langle \lambda yNx\rangle(y)$ and $\langle \lambda x.\neg Sx\rangle(d) \rightarrow \exists y\langle \lambda x.\neg Sx\rangle(y)$ are valid, but not every sentence of the form $\psi(\iota) \rightarrow \exists\alpha\psi(\alpha)$ is. For example, $\neg Sd \rightarrow \exists x\neg Sx$ is not. Similarly, not all instances of $(\psi(\iota) \wedge \mathcal{E}(\iota)) \rightarrow \exists\alpha\psi(\alpha)$ are valid. Although this schema is valid in standard nonmodal free logics, the presence of $\mathcal{E}(\iota)$ in the antecedent is not sufficient to guarantee validity when modal connectives are available. For example, the sentence $(\Box Sd \wedge \mathcal{E}(d)) \rightarrow \exists x\Box Sx$ is invalid. Individual constants may denote different objects at different worlds. So Sd can be true at every world accessible from a world w even though no object in the domain of w exhibits S at all such worlds.

When applied to predicate abstracts quantifiers do not exhibit their usual dual properties. Thus in spite of the validity of (EG), not all sentences of the form

$$\forall\alpha_1\big\langle\lambda\alpha_2.\phi(\alpha_2)\big\rangle(\alpha_1) \rightarrow \big\langle\lambda\alpha_2.\phi(\alpha_2)\big\rangle(\iota)$$

are valid. Indeed

$$\forall x\langle \lambda y.Sy\rangle(x) \rightarrow \langle \lambda y.Sy\rangle(d)$$

is not. Even if the antecedent is true, d must designate in order for the consequent to be true.

Universal instantiation in the form

(UI) $\forall\alpha_1\langle \lambda\alpha_2.\phi(\alpha_2)\rangle(\alpha_1) \rightarrow (\mathcal{E}(\iota) \rightarrow \langle \lambda\alpha_2.\phi(\alpha_2)\rangle(\iota))$,

however, is valid. And just as in the case of (EG) the use of predicate abstraction in (UI) is essential, since not every sentence of the form

(*faux* UI) $\forall\alpha_1\phi(\alpha_1) \rightarrow (\mathcal{E}(\iota) \rightarrow \phi(\iota))$

is valid. For example, the sentence

$$\forall x\Box Sx \rightarrow \big(\mathcal{E}(d) \rightarrow \Box Sd\big)$$

is not. There is a tableau for this sentence, and a countermodel, in Section 8.3.

4.4 The Barcan and Converse Barcan Formulas SAT provides sensible treatments of the Barcan and Converse Barcan Formulas. Consider first the following form of the Converse Barcan Formula:

(CBF) $\exists\alpha_1\Diamond\langle\lambda\alpha_2.\phi(\alpha_2)\rangle(\alpha_1)\to\Diamond\exists\alpha_1\langle\lambda\alpha_2.\phi(\alpha_2)\rangle(\alpha_1)$.

(CBF) is valid, as seems entirely plausible. Evaluated at any world w, (CBF) says (roughly) that if there is an object in w that has a property in some possible world w', then there is a possible world in which some object has this property. Given the antecedent, and given that according to serious actualism an object can have a property at a world only if it exists in that world, w' as described is sufficient to make the consequent true. If an object from w has a property in w', that object must exist in w'.[29]

Standard counterexamples to this form of the Converse Barcan Formula assume that true predications can be made of objects at worlds in which those objects do not exist.[30] Serious actualism rejects this assumption. In so doing it provides a plausible explanation of why these counterexamples so often seem wrong to the uninitiated. They are wrong.

Of course the Converse Barcan Formula can also be expressed using necessity and universal quantification. (CBF) is logically equivalent to

$$\Box\forall\alpha_1\neg\big\langle\lambda\alpha_2.\phi(\alpha_2)\big\rangle(\alpha_1)\to\forall\alpha_1\Box\neg\big\langle\lambda\alpha_2.\phi(\alpha_2)\big\rangle(\alpha_1),$$

which is also valid. But in this form the fact that negations precede predicate abstracts is crucially important. Not all sentences of the form

(*faux* CBF) $\Box\forall\alpha_1\langle\lambda\alpha_2.\phi(\alpha_2)\rangle(\alpha_1)\to\forall\alpha_1\Box\langle\lambda\alpha_2.\phi(\alpha_2)\rangle(\alpha_1)$

are valid. Standard counterexamples apply.[31] There are tableaus for instances of (CBF) and (*faux* CBF) in Section 8.3.

It is worth noting that Stalnaker [32] gives a version of the Converse Barcan Formula that is entailed by, but does not entail, (*faux* CBF). Translated into my notation it is

(Stal CBF) $\Box\forall\alpha_1\langle\lambda\alpha_2.\phi(\alpha_2)\rangle(\alpha_1)\to\forall\alpha_1\langle\lambda\alpha_2.\Box\phi(\alpha_2)\rangle(\alpha_1)$.

(Stal CBF) is not valid in SAT or in Stalnaker's system. (For more about the latter, see Section 6.2.)

Turning now to the Barcan Formula, it is not difficult to see that both of the following formulations have invalid instances:

$$\Diamond\exists\alpha_1\big\langle\lambda\alpha_2.\phi(\alpha_2)\big\rangle(\alpha_1)\to\exists\alpha_1\Diamond\big\langle\lambda\alpha_2.\phi(\alpha_2)\big\rangle(\alpha_1),$$
$$\forall\alpha_1\Box\big\langle\lambda\alpha_2.\phi(\alpha_2)\big\rangle(\alpha_1)\to\Box\forall\alpha_1\big\langle\lambda\alpha_2.\phi(\alpha_2)\big\rangle(\alpha_1).$$

Standard counterexamples again apply.[32]

4.5 Extensionality of predicate abstraction Predicate abstraction in SAT is in a certain sense extensional, as the following valid sentences show:

$$\forall z\big(\big\langle\lambda x.\langle\lambda y.Fy\rangle(x)\big\rangle(z)\leftrightarrow\langle\lambda y.Fy\rangle(z)\big),$$
$$\big\langle\lambda x.\langle\lambda y.Fy\rangle(x)\big\rangle(d)\leftrightarrow\langle\lambda y.Fy\rangle(d),$$
$$\big\langle\lambda z,w.\langle\lambda x,y.Rxy\rangle(z,w)\big\rangle(d,o)\leftrightarrow\langle\lambda x,y.Rxy\rangle(d,o),$$
$$\forall u\forall v\big(\big\langle\lambda z,w.\langle\lambda x,y.Rxy\rangle(z,w)\big\rangle(u,v)\leftrightarrow\langle\lambda x,y.Rxy\rangle(u,v)\big).$$

If predicate abstracts are taken to represent properties, then having a property and having the property of having that property are not distinguished in the semantics, similarly for relations.

5 Some Important Semantic Metatheorems

Since antiquity logicians have realized the importance of logical form. If a sentence is logically true, then every sentence of the same form is also logically true. In predicate logic this principle manifests itself in the principle of uniform substitution for predicates. Substituting the same complex predicate for each occurrence of a simple predicate in a logically true sentence should yield another logical truth. Yet seriously actualistic logics have commonly been thought to violate this fundamental principle. I show in Section 5.1 that when predication is understood properly, using predicate abstraction, uniform substitution of predicates does indeed preserve logical truth. This holds for SAT and its kindred systems, and it is one of their main advantages over other attempts to formalize seriously actualistic modal logics.

Another important feature of SAT is that it supports a theorem (sometimes called a *replacement theorem*) that sanctions substituting logically equivalent subformulas for each other within a given formula. If, in a formula, an occurrence of a subformula is replaced by a logically equivalent formula, the result is logically equivalent to the original. Of course such a theorem depends on an appropriate definition of logical equivalence for *formulas*. (The definition of logical equivalence given in Section 3.6 applies only to sentences.) In Section 5.2 I define logical equivalence for formulas and prove a replacement theorem.

Section 5.3 contains two theorems about validity and logical consequence of sentences. Section 5.4 contains a deduction theorem.[33]

5.1 Uniform substitution for predicate abstracts If a sentence ϕ is logically true, then any sentence that results from substituting a new predicate (simple or complex) for each occurrence of a simple predicate in ϕ should be logically true as well. Yet it appears that uniform substitution does not preserve logical truth in seriously actualistic systems. For example, where F and G are primitive predicates, although

$$\exists x \Diamond Fx \rightarrow \Diamond \exists x Fx$$

is valid in SAT,

$$\exists x \Diamond \neg Gx \rightarrow \Diamond \exists x \neg Gx$$

is not.[34] Yet the second sentence has the same general form as the first, since it results from substituting $\neg Gx$ for Fx. I believe failure of uniform substitution is the reason seriously actualistic systems of quantified modal logic have not been widely studied.

Indeed Kripke cited this failure as a major problem very early in the development of model-theoretic modal semantics. In Section 1 I quoted the salient passage from [19], one of his most important and most frequently cited early papers. Yet if properly understood, uniform substitution does preserve validity in seriously actualistic systems.[35]

Proper understanding means taking predicate abstraction seriously as the preferred means of expressing predication, and thus understanding uniform substitution as substitution of a predicate abstract for each occurrence of a primitive predicate abstract throughout a sentence.[36] For example, if we substitute $\langle \lambda y.\neg Gy \rangle$ for $\langle \lambda y.Fy \rangle$ throughout the valid sentence

(11) $\exists x \Box \langle \lambda y.Fy \rangle(x) \rightarrow \Box \exists x \langle \lambda y.Fy \rangle(x)$,

we obtain

(12) $\exists x \Box \langle \lambda y.\neg Gy \rangle(x) \rightarrow \Box \exists x \langle \lambda y.\neg Gy \rangle(x)$,

which is also valid. Theorem 1 shows that the foregoing is not an isolated example.

Theorem 1 *Uniform Substitution for Predicate Abstracts (USPA)*

Consider a sentence Φ, an atomic formula $\theta\alpha_1 \ldots \alpha_n$, where $\alpha_1, \ldots, \alpha_n$ are distinct variables, and a formula ψ. Suppose that each of the variables $\alpha_1, \ldots, \alpha_n$ has at least one free occurrence in ψ, and no other variable occurs free in ψ. Suppose also that θ occurs in Φ only as part of the predicate abstract $\langle \lambda\alpha_1, \ldots, \alpha_n.\theta\alpha_1 \ldots \alpha_n \rangle$, and that Ψ is the result of simultaneously substituting $\langle \lambda\alpha_1, \ldots, \alpha_n.\psi \rangle$ for each occurrence of $\langle \lambda\alpha_1, \ldots, \alpha_n.\theta\alpha_1 \ldots \alpha_n \rangle$ in Φ. Under these conditions, if Φ is valid, then so is Ψ.

Proof Consider the case where $n = 1$. To facilitate the proof, I define the extension of a unary predicate and the extension of a predicate abstract at a world in a model. If θ is a unary predicate, the *extension of predicate θ at world w in model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$* (abbreviated $\text{Ext}_{\mathcal{M}_w}(\theta)$) is $\mathcal{I}(\theta, w)$. By the definition of an interpretation (see Section 3.2) $\text{Ext}_{\mathcal{M}_w}(\theta) \subseteq \mathcal{D}(w)$. If ψ is a formula, α is a variable, α has at least one free occurrence in ψ, and no variable other than α occurs free in ψ, the *extension of predicate abstract $\langle \lambda\alpha.\psi \rangle$ at world w in model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$* (abbreviated $\text{Ext}_{\mathcal{M}_w}(\langle \lambda\alpha.\psi \rangle)$) is $\{\mathcal{V}(\alpha) \mid \mathcal{M}_{\mathcal{V}}(\langle \lambda\alpha.\psi \rangle(\alpha), w) = 1\}$, where $\mathcal{V}$ ranges over all valuations relative to $\mathcal{M}$. By clause 12 in the definition of satisfaction (see Section 3.4) it follows that $\text{Ext}_{\mathcal{M}_w}(\langle \lambda\alpha.\psi \rangle) \subseteq \mathcal{D}(w)$.

Since θ is a unary predicate, $\theta\alpha$ is a formula. So $\text{Ext}_{\mathcal{M}_w}(\langle \lambda\alpha.\theta \rangle)$ is $\{\mathcal{V}(\alpha) \mid \mathcal{M}_{\mathcal{V}}(\langle \lambda\alpha.\theta\alpha \rangle(\alpha), w) = 1\}$, where $\mathcal{V}$ ranges over all valuations relative to $\mathcal{M}$. It follows from the definition of satisfaction (see Section 3.4) and the definition of $\text{Ext}_{\mathcal{M}_w}(\theta)$ (above) that $\text{Ext}_{\mathcal{M}_w}(\langle \lambda\alpha.\theta\alpha \rangle) = \text{Ext}_{\mathcal{M}_w}(\theta) = \mathcal{I}(\theta, w)$.

Now suppose that Ψ is invalid. Hence there is a model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$, a world $w \in \mathcal{W}$, and a valuation $\mathcal{V}$, such that $\mathcal{M}_{\mathcal{V}}(\Psi, w) = 0$. Consider the model $\mathcal{M}^* = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I}^* \rangle$, which is exactly like $\mathcal{M}$ except that, for all $w^* \in W$, $\mathcal{I}^*(\theta, w^*) = \text{Ext}_{\mathcal{M}_{w^*}}(\langle \lambda\alpha.\psi \rangle)$. Since $\text{Ext}_{\mathcal{M}_{w^*}}(\langle \lambda\alpha.\psi \rangle) \subseteq \mathcal{D}(w^*)$, $\mathcal{I}^*$ is well defined. A straightforward induction on the complexity of Φ shows that $\mathcal{M}^*_{\mathcal{V}}(\Phi, w) = 0$. Hence Φ is invalid. The proof easily generalizes to cases where $n > 1$.[37] ■

In addition to sanctioning the move from the validity of (11) to that of (12), USPA also allows us to infer the validity of (14) from that of (13):

(13) $\exists x \forall y \langle \lambda u, v.Ruv \rangle(x, y) \rightarrow \forall y \exists x \langle \lambda u, v.Ruv \rangle(x, y)$,

(14) $\exists x \forall y \langle \lambda u, v.\neg \Diamond \exists w(Fuw \wedge Gwv) \rangle(x, y) \rightarrow \forall y \exists x \langle \lambda u, v.\neg \Diamond \exists w(Fuw \wedge Gwv) \rangle(x, y)$.

The reader will be able to generate other examples of inferences sanctioned by USPA.

5.2 Replacement (substitution of equivalents) Well-behaved logics typically support a definition of logical equivalence for formulas and a replacement theorem. Such theorems generally guarantee that given a formula ϕ, replacing any subformula ψ of ϕ by a formula ψ' equivalent to ψ, yields a formula ϕ' equivalent to the original ϕ.

In nonmodal first-order logic it is sufficient for this purpose to declare formulas χ and χ' logically equivalent if $\forall\alpha_1 \ldots \forall\alpha_n(\chi \leftrightarrow \chi')$ is valid, where $\alpha_1, \ldots, \alpha_n$ include all the variables that have free occurrences in χ or χ'. But such a definition will not support substitution of equivalents for modal logics that allow different worlds in a model to have different domains of individuals. For these logics we need a stronger notion of logical equivalence.

For a formula ϕ, let the *$\forall\Box$-closure of ϕ* be any sentence formed by prefixing ϕ with a sequence of interspersed universal quantifiers and necessity connectives. This sequence may be of any finite length and the universal quantifiers and necessity connectives may appear in any order. None of the quantifiers may be vacuous, however, since vacuous quantifiers are banned by the definition of a formula. A $\forall\Box$-closure of a sentence will consist of that sentence prefixed with zero or more necessity connectives.

For a formula ϕ I write $\Vdash \phi$ to indicate that every $\forall\Box$-closure of ϕ is valid, and (extending the corresponding definition for nonmodal first-order logic given two paragraphs back) I say that two formulas ψ and ψ' are *logically equivalent* if and only if $\Vdash (\psi \leftrightarrow \psi')$.[38] Armed with this definition the statement and proof of the replacement theorem are straightforward.[39]

Theorem 2 *Replacement (Substitution of Equivalent Subformulas)*

Let ψ, ψ', ϕ, and ϕ' be formulas. If $\Vdash (\psi \leftrightarrow \psi')$, and ϕ' is like ϕ except that ϕ' contains an occurrence of ψ' at a place where ϕ contains an occurrence of ψ, then $\Vdash (\phi \leftrightarrow \phi')$.

Proof This is straightforward, adapting the proof of replacement in Mates [22, pp. 135–36]. First prove the following lemma. If $\Vdash (\psi \leftrightarrow \psi')$, then $\Vdash (\neg\psi \leftrightarrow \neg\psi')$, $\Vdash ((\psi \wedge \chi) \leftrightarrow (\psi' \wedge \chi))$, $\Vdash (\Box\psi \leftrightarrow \Box\psi')$, $\Vdash (\mathsf{A}\psi \leftrightarrow \mathsf{A}\psi')$, $\Vdash (\forall\alpha\psi \leftrightarrow \forall\alpha\psi')$, and $\Vdash (\langle\lambda\alpha.\psi\rangle(\tau) \leftrightarrow \langle\lambda\alpha.\psi'\rangle(\tau))$, where ψ, ψ', and χ are formulas, α is a variable, and τ is a term.[40] Theorem 2 then follows by induction on the complexity of ϕ. ■

5.3 Sentences: Closure, validity, consequence, equivalence Two further results concerning sentences are easily established.

Theorem 3 *Closure and Validity*

For any sentence ϕ, $\Vdash \phi$ if and only if $\vDash \phi$.

Proof The "if" part is immediate by the rule of necessitation. For the "only if" part, suppose $\nvDash \phi$. Then ϕ is false at some world w of some model $\mathcal{M}$. Construct a model $\mathcal{M}'$ that is exactly like $\mathcal{M}$ except for containing an additional world w' such that $w'\mathcal{R}w$. Then $\Box\phi$ is false at w', and hence $\nvDash \Box\phi$. So it is false that $\Vdash \phi$. ■

Theorem 4 *Logical Consequence and Logical Equivalence*

For any sentences ϕ and ϕ', the following three conditions are equivalent:

(1) $\Vdash (\phi \leftrightarrow \phi')$,
(2) $\vDash (\phi \leftrightarrow \phi')$,
(3) $\phi \vDash \phi'$ and $\phi' \vDash \phi$.

Proof The equivalence of (1) and (2) is immediate by Theorem 3. The equivalence of (2) and (3) holds in view of the standard definition of consequence in Section 3.6. ■

It is noteworthy that, as in standard first-order nonmodal logic, two sentences are logically equivalent if and only if each is a consequence of the other as (2) and (3) of Theorem 4 show.

5.4 Deduction theorem A deduction theorem exactly analogous to that of standard first-order nonmodal logic also holds.

Theorem 5 *Deduction Theorem*

For any sentences ϕ and ϕ', and any set of sentences Γ, $\Gamma \cup \{\phi\} \vDash \phi'$ if and only if $\Gamma \vDash \phi \to \phi'$.

Proof This is obvious in view of the standard definition of consequence in Section 3.6. ■

6 Comparisons with Other Systems

Modal logics that are seriously actualistic with respect to *atomic* sentences have been studied by several authors, but because the languages considered lack a predicate abstraction operator, none of them support uniform substitution.[41] The work of Plantinga, Jager, and Stephanou is nonetheless noteworthy. I discuss it in Section 6.1.

Stalnaker [31] gives a subtle and insightful discussion of logical form in a first-order nonmodal language supplemented with a predicate abstraction operator. He also briefly discusses a seriously actualistic modal extension of this language, emphasizing the importance of the scope of the abstraction operator in determining the logical form of a sentence. And he suggests (see [31, pp. 335–36]), but does not prove, that in this modal extension uniform substitution preserves logical truth. He further explores this same modal extension in [32].[42] I discuss his work in Section 6.2. The work of Garson and Chihara is also discussed briefly in Section 6.2 to help locate SAT within the literature.

6.1 Actualism and serious actualism: Plantinga, Jager, Stephanou Actualism and serious actualism have long been advocated by Plantinga. He defines actualism as the claim that there neither are nor could have been things that do not exist.[43] If we adopt Exists as a predicate for existence, this can be expressed as

(Act1) $\neg(\exists x \neg \text{Exists}(x) \vee \Diamond \exists x \neg \text{Exists}(x))$,

which is equivalent to

(Act2) $\Box \forall x\, \text{Exists}(x)$.

Plantinga's idea seems to be that actualists and possibilists will agree that Exists holds of all and only actual objects, and that the dispute between them is over the scope of the quantifiers. Since actualists can make no sense of objects that are possible but not actual, they take quantifiers to range over just actual objects. This would seem to make (Act1) and (Act2) true. But for possibilists quantifiers range over both actual and merely possible objects, apparently making (Act1) and (Act2) false.

Unfortunately, if we try to express the foregoing reasoning using standard modal logics, a problem arises. Since Exists is defined as holding of just actual objects, it can presumably be expressed as

$$\mathsf{A}\mathcal{E}(x)$$

(where A is an actuality connective and $\mathcal{E}(x)$ is defined as $\exists y(y = x)$, as in Section 4.2). (Act2) should then be expressible as

(Act2*) $\Box\forall x \mathsf{A}\mathcal{E}(x)$.

But (Act2*) does not capture the interpretation of (Act2) that I attribute to Plantinga in the previous paragraph. For in standard varying domain modal semantics the range of a quantifier at a world is the domain of that world, and some of these domains are typically construed as containing possible but nonactual objects.[44] So in these logics, SAT included, (Act2*) is false at the actual world element of the intended model. In view of this, I suggest that actualism is properly expressed by

(Act3) $\forall x \mathsf{A}\mathcal{E}(x)$,

which is true at the actual world element of the intended model $\mathcal{M}_{\mathcal{I}}^{\mathcal{A}}$, defined in Section 7.

Because an actualist cannot consistently allow nonactual objects to be constituents of the intended model of his allegedly actualistic modal language, there may seem to be a problem with this view. But an actualist can, following Menzel [23], construe possible objects and the possible worlds they inhabit as set-theoretic constructs that use only actual objects to represent ways things might have been, not as (or having as constituents) anything that is merely possible. So although actualists deny the existence of merely possible objects, they need not deny themselves the use of varying domain semantics.

In spite of the foregoing problem with Plantinga's account of actualism, there is much in his work that supports and is supported by SAT. His discussion in [26] of what he calls predicative and impredicative singular propositions is a case in point. The former affirm, and the latter deny, a property of an object.[45] Plantinga's rebuttal of what he calls the *classical argument* for possible but nonexistent objects makes use of this distinction.[46] Specifically, (using his numbering) he distinguishes the impredicative

(13*) Socrates does not have the property of existing

from the predicative

(13**) Socrates has the property of nonexistence.[47]

Plantinga says, rightly, that (13*) is true in just those worlds in which

(23*) Socrates exists

is false, and that (13**) is false in every world.

Predicate abstraction is ideally suited to mark these distinctions. In SAT Plantinga's (13*) and (13**) are naturally expressed as

(*) $\neg\langle\lambda y.\mathcal{E}(y)\rangle(s)$

and

(**) $\langle\lambda y.\neg\mathcal{E}(y)\rangle(s)$,

respectively, and their truth values are determined exactly as he says. The sentence (*) is true at a world in a model just in case s fails to denote at that world, and the denial of (**) is true at every world in every model. So in SAT

$(\neg^{**})$ $\neg\langle\lambda y.\neg\mathcal{E}(y)\rangle(s)$

and indeed

$(\neg\Diamond^{**})$ $\neg\Diamond\langle\lambda y.\neg\mathcal{E}(y)\rangle(s)$

are valid. Not only is it false that Socrates is nonexistent, it is impossible that he be so. And the use of predicate abstraction makes it clear that this latter claim, embodied in $(\neg\Diamond^{**})$, is a far cry from

$$\neg\Diamond\neg\big\langle\lambda y.\mathcal{E}(y)\big\rangle(s),$$

and from several orthographically similar sentences, all of which are false in the intended model $\mathcal{M}^{\mathcal{A}}_{\mathcal{I}}$.[48,49]

Plantinga also advocates serious actualism, which he expresses as the necessitation of "…no object could have a property without existing."[50] This cannot be expressed in a first-order language without quantifying over properties, but it can be approximated schematically as

(SA) $\Box\forall\alpha_1\Box(\langle\lambda\alpha_2.\phi(\alpha_2)\rangle(\alpha_1) \rightarrow \mathcal{E}(\alpha_1))$,

all instances of which are valid in SAT.

Plantinga nowhere formalizes his preferred modal logic; for this he refers his readers to the work of Jager. In [16], Jager presents his system A, in which an atomic formula is satisfied at a world by an n-tuple of objects only if each of those objects exists at that world. System A is thus seriously actualistic with respect to atomic formulas. But it lacks predicate abstraction and thus is not fully seriously actualistic in the way SAT is. A contains the usual connectives and quantifiers, but it lacks individual constants and an actuality connective. For each model the variables of A take their values from the domain of that model, and they denote the same object at each world therein. But the range of quantifiers at a world is limited to the domain of that world. Thus in its treatment of variables and quantifiers A is exactly like SAT. System A differs from SAT in that negation and necessity are given *de re* interpretations, which I will represent with $\neg_r$ and $\Box_r$, and which are easily defined in SAT. For example, $\neg_r Fx$ and $\Box_r Fx$ can be expressed as

$$\big\langle\lambda y.\big(\mathcal{E}(y) \wedge \neg Fy\big)\big\rangle(x)$$

and

$$\big\langle\lambda y.(\mathcal{E}(y) \wedge Fy \wedge \Box\big(\mathcal{E}(y) \rightarrow Fy\big)\big\rangle(x),$$

respectively.[51]

Jager [17] extends system A to the system he calls D, to facilitate distinguishing *de re* and *de dicto* modalities and negations. He does this by supplementing A in a way that makes it possible to exempt particular occurrences of variables within the scope of $\neg_r$ or $\Box_r$ from the requirement, described in the previous paragraph, that they denote an object at the world of evaluation in order for the formula in which they appear to be satisfied at that world. Specifically, he adds an operator he calls a *dictifier* (∇) and, for each variable α, infinitely many *position variants*, α^1, α^2, $\alpha^3, \ldots,$ of α. (In the semantics of D each position variant of a variable is assigned the same value as the variable itself.) Thus if $\nabla x^1 Rx^1 y^1$ appears within the scope of $\neg_r$ or $\Box_r$, x^1 (but not y^1) is exempt from the requirement in question.

Using these resources, and adding individual constants such as s and p to the language of D to facilitate representing English sentences (as Jager himself does), the sentences

Socrates has the property of necessarily-teaching-Plato

and

Plato has the property of necessarily-being-taught-by-Socrates

can be expressed as

$$\Box_r \nabla pTsp$$

and

$$\Box_r \nabla sTsp,$$

respectively. In SAT these sentences can be expressed, at least as perspicuously, as

$$\langle \lambda x.\Box(\mathcal{E}(x) \to Txp)\rangle(s)$$

and

$$\langle \lambda x.\Box(\mathcal{E}(x) \to Tsx)\rangle(p),$$

respectively.

More generally, it is easy to show that every Jager model and variable assignment can be exactly replicated by a SAT model and valuation.[52] Given this fact, and using the SAT constructions for expressing Jager's *de re* negation and *de re* necessity given three paragraphs back, it is easily proved that a formula of system D is satisfied by a variable assignment at a world of a Jager model if and only if it is satisfied at the same world of the corresponding SAT model by the corresponding valuation.

Finally, it should be noted that Stephanou [38] has given a long and subtle defense of serious actualism. Although, like Plantinga, he understands serious actualism as quantifying over properties and relations as well as individuals, he recognizes that it is partially expressed by first-order formulas of the form

(SA1) $\forall \alpha_1 \Box(\phi(\alpha_1) \to \exists \alpha_2(\alpha_1 = \alpha_2))$,

where $\phi(\alpha_1)$ is an atomic formula in which α_1, and perhaps other variables, have free occurrences. He calls the schema (SA1) *predicate actualism*, and he is keenly aware that it is plausible only when $\phi(\alpha_1)$ is construed as an atomic formula. For he notes that from (SA1) and

$$\forall \alpha_1 \Box(\neg\phi(\alpha_1) \to \exists \alpha_2(\alpha_1 = \alpha_2))$$

it follows that

$$\forall \alpha_1 \Box \exists \alpha_2(\alpha_1 = \alpha_2),$$

a result he wants to avoid.

In [38] Stephanou does not present a system of formal semantics, although he considered such systems in [36] and [37]. And although (SA1) is a valid schema in the systems he presents in the latter two papers, uniform substitution of complex formulas for atomic formulas does not preserve validity in any of them. For none of these systems contains predicate abstraction or any means of obtaining its effect. So while Stephanou's semantic systems treat primitive predicates in accordance with serious actualism, they lack a mechanism for passing that treatment on to more complex formulas.

6.2 Names and abstraction: Garson, Chihara, Stalnaker Of all the systems discussed by Garson [8] in his encyclopedic survey of modal logics SAT is most similar to Q3.[53] Both are free logics in which quantifiers are world-relative, variables are rigid, and individual constants are nonrigid. They differ in that SAT contains predicate abstraction and embodies serious actualism, features that Q3 lacks.

Because it contains predicate abstraction SAT supports a straightforward free-logic version of universal instantiation: (UI) given in Section 4.3. The version of universal instantiation Garson gives for Q3 (due to Hintikka) is considerably more complex. And although the individual constants of SAT and Q3 are nonrigid, in both systems they can be made to approximate rigid designators. For the individual constant d all that is required is the stipulation $\exists x\Box(d = x)$, or $\exists x(\Box(d = x) \wedge \Box\Box(d = x))$, or $\exists x(\Box(d = x) \wedge \Box\Box(d = x) \wedge \Box\Box\Box(d = x))$, and so on.[54] We can thus stipulate that d denotes one and the same object in all possible worlds (including the actual world), or that it does this in all possible worlds and in all possibly possible worlds, and so on. (In SAS4 and SAS5 this entire set of sentences can be replaced by the single sentence $\exists x\Box(d = x)$.)

Chihara's systems M_1 and M^* are similar to SAT in that they are seriously actualistic free logics with world-relative quantifiers. They differ from it in that their individual constants are model-wide rigid designators. They also differ in lacking predicate abstraction and identity. SAT is closer to M_1, since M^* is simply M_1 plus Chihara's (seemingly false) Principle of Compossibility.[55]

More importantly for my purposes Stalnaker [32] develops a quantified modal logic with predicate abstraction that is similar to, but interestingly different from, SAT.[56] A striking feature of this approach is that quantifiers apply directly to predicate abstracts (and to unary predicates) rather than to formulas. Stalnaker says this:

> ... gives a clearer representation of the logic of quantification because it separates two conceptually distinct operations that are performed by variable-binding operators. First is the implicit formation of complex predicates from complex sentences by introducing blanks—free variables—in the sentences. The second is generalization: the formation of general claims from predicates—the claims that everything in the domain, or at least one thing in the domain, satisfies the predicate that is implicitly represented by the open sentence. In our language, the abstraction operator makes the first of these operations explicit, turning an open sentence into an expression that has the syntactic role as well as the semantic function of a predicate. Then the quantifier has only the job of expressing generality.[57]

Stalnaker's predicate abstraction operator (represented by a cap over a variable) plays the same role as the lambda operator plays in SAT. Where S is a unary predicate, $\hat{x}Sx$ and $\hat{x}Sxd$ correspond to $\langle\lambda x.Sx\rangle$ and $\langle\lambda x.Sx\rangle(d)$. The sentences $\forall\hat{x}Sx$ and $\forall S$ correspond to $\forall y\langle\lambda x.Sx\rangle(y)$ and $\forall ySy$. I will focus on his T-based system, which I will call Stal-T, to facilitate comparison with SAT.[58]

In addition to predicate abstraction and quantifiers Stal-T, like SAT, contains the usual truth-functional and modal connectives. Indeed SAT and Stal-T are nearly identical logics. Yet even though Stalnaker is clearly interested in the proper representation of logical form, he does not discuss uniform substitution for predicate abstracts (cf. Section 5.1 above). Neither does he discuss substitution of equivalent subformulas for one another (cf. Section 5.2 above). His main focus is on showing how to combine extensional first-order free logic and propositional modal logic without having to retract anything from either and adding only "...two principles that concern the interaction of modality with predication, and modality with identity."[59] He gives slightly different axiomatizations of this logic in the two versions of the paper, and in both versions he claims soundness and completeness. But neither soundness nor completeness is proved in either version.

SAT and Stal-T are alike in that the extensions of predicates and predicate abstracts at a world are restricted to objects in the domain of that world. Both are thus seriously actualistic with respect to all predicates, unlike the systems of Jager and Stephanou (discussed in Section 6.1) that are seriously actualistic only with respect to primitive predicates. They are also alike in that the domain of a world in a model may be empty, an individual constant need not denote at a world and may denote different objects at different worlds, and the denotation (if any) of an individual constant at a world must be in the domain of that world. They differ mainly in that Stal-T, but not SAT, counts vacuous predicate abstracts as well formed, and applies the terms "valid" and "invalid" to formulas containing free occurrences of variables. If the syntax and semantics of Stal-T are modified to be like those of SAT in these two ways, and if we restrict attention to formulas of SAT not containing the actuality connective (which Stal-T lacks), corresponding sentences of the two languages will take the same truth value at each world of each model.[60]

The sentences of Stal-T are unambiguous, but inserting $\langle$ and $\rangle$ as additional brackets to mark off predicate abstracts will improve readability and facilitate comparison with SAT. So, for example, where S is a primitive unary predicate I will write $\exists\langle S\rangle$ for $\exists S$, $\exists\langle\hat{x}Sx\rangle$ for $\exists\hat{x}Sx$, $\neg\forall\langle\hat{x}\neg Sx\rangle$ for $\neg\forall\hat{x}\neg Sx$, and $\forall\langle\hat{x}\Box\langle\hat{y}Sy\rangle x\rangle$ for $\forall\hat{x}\Box\hat{y}Syx$.[61] The transformations that preserve equivalence are as follows.

Starting with a sentence of SAT replace subformulas of the form $\forall\alpha\phi$ with $\forall\langle\hat{\alpha}\phi\rangle$, $\exists\alpha\phi$ with $\exists\langle\hat{\alpha}\phi\rangle$, and $\langle\lambda\alpha.\phi\rangle(\beta)$ with $\langle\hat{\alpha}\phi\rangle\beta$. To go in the other direction first insert $\langle$ and $\rangle$ around predicate abstracts as in the previous paragraph. Next, working from the inside out, replace subformulas of the form $\forall\langle\psi\rangle$ and $\exists\langle\psi\rangle$ with $\forall\alpha\langle\psi\rangle\alpha$ and $\exists\alpha\langle\psi\rangle\alpha$, respectively, where the variable α does not appear in ψ. Then replace subformulas of the form $\langle\hat{\alpha}\phi\rangle\beta$ with $\langle\lambda\alpha.\phi\rangle(\beta)$. This will yield a sentence of SAT. Under these transformations the resulting sentence will be equivalent to the original, given the modifications of Stal-T syntax and semantics noted two paragraphs back.

Stalnaker points out that both the Barcan Formula and its converse have invalid instances when expressed in his language as

$$\forall\hat{x}\Box\phi \to \Box\forall\hat{x}\phi$$

and

$$\Box\forall\hat{x}\phi \to \forall\hat{x}\Box\phi.$$

But he does not mention that all instances of the following form of the Converse Barcan Formula

$$\exists\hat{x}\Diamond\phi \to \Diamond\exists\hat{x}\phi$$

are valid in Stal-T.

For example, taking ϕ to be Sx (where S is a primitive unary predicate) this becomes

$$\exists\hat{x}\Diamond Sx \to \Diamond\exists\hat{x}Sx,$$

and adding brackets yields

$$\exists\langle\hat{x}\Diamond Sx\rangle \to \Diamond\exists\langle\hat{x}Sx\rangle.$$

The transformation process given above then yields

$$\exists y \langle \lambda x.\Diamond Sx\rangle(y) \rightarrow \Diamond\exists y\langle \lambda x.Sx\rangle(y),$$

which is equivalent to

$$\exists y\Diamond\langle \lambda x.Sx\rangle(y) \rightarrow \Diamond\exists y\langle \lambda x.Sx\rangle(y),$$

an instance of the valid schema (CBF) given in Section 4.4 above.

Indeed, where $\phi(\alpha_2)$ is a formula containing free occurrences of α_2 (but no free occurrence of any other variable) all sentences of the form

(CBF*) $\exists\hat{\alpha}_1\Diamond\hat{\alpha}_2\phi(\alpha_2)\alpha_1 \rightarrow \Diamond\exists\hat{\alpha}_1\hat{\alpha}_2\phi(\alpha_2)\alpha_1$

are valid in Stal-T. But not all sentences of the form

(*faux* CBF*) $\Box\forall\hat{\alpha}_1\hat{\alpha}_2\phi(\alpha_2)\alpha_1 \rightarrow \forall\hat{\alpha}_1\Box\hat{\alpha}_2\phi(\alpha_2)\alpha_1$

are. And (CBF*) and (*faux* CBF*) are equivalent to the translations into Stal-T of (CBF) and (*faux* CBF) of Section 4.4. So Stal-T distinguishes between the two versions of the Converse Barcan Formula just as SAT does.

7 Truth *Simpliciter* and Actualism

Thus far I have been concerned almost exclusively with matters of logic: logical truth, logical consequence, logical equivalence. These notions were defined, using models, in a way that embodies serious actualism. Indeed we have seen that serious actualism is reflected in a logically true sentence of the language. (Recall (SA) from Section 6.1.) But as yet I have said almost nothing about truth *simpliciter*.

Given model-theoretic semantics it is natural to designate a particular model as the intended model and then identify truth *simpliciter* with truth at the actual world element of this model. A modal realist would specify the intended model, $\mathcal{M}_\mathcal{I}$ ($\mathcal{I}$ for intended), as $\langle \mathcal{W}_\mathcal{I}, @_\mathcal{I}, \mathcal{R}_\mathcal{I}, \mathcal{D}_\mathcal{I}, \mathcal{I}_\mathcal{I}\rangle$, where $\mathcal{W}_\mathcal{I}$ really is the set of all possible worlds, $@_\mathcal{I}$ is the actual world, $\mathcal{R}_\mathcal{I}$ is the relation of relative possibility among the worlds, $\mathcal{D}_\mathcal{I}(@_\mathcal{I})$ is the set of all actual individuals, $\mathcal{D}_{\mathcal{M}_\mathcal{I}}$ is the set of all actual and possible individuals, and so on. Truth *simpliciter* would then be truth at $@_\mathcal{I}$ of $\mathcal{M}_\mathcal{I}$. But such an approach is anathema to actualists, who eschew all talk of possible worlds, possible objects, and alleged relations among such alleged entities. So construed $\mathcal{M}_\mathcal{I}$ does not exist.

Fortunately Menzel [23] has shown in considerable detail how the foregoing approach can be shorn of metaphysical excess and made acceptable to actualists. Actualists need only acknowledge that they understand and accept as primitive basic modal terms (like those italicized in this paragraph). Menzel's innovation is based on construing selected pure sets as surrogates for actual and possible objects and using them to create an actualistically acceptable model that *would be* isomorphic to $\mathcal{M}_\mathcal{I}$ if $\mathcal{M}_\mathcal{I}$ existed.[62] Under this approach $@_\mathcal{I}$ is replaced by a set-theoretic construct that accurately models the world, how things are. The other members of $\mathcal{W}_\mathcal{I}$ are replaced by similar set-theoretic constructs, each of which *might have been* a model of the world. $\mathcal{D}_\mathcal{I}(@_\mathcal{I})$ becomes the set of surrogates of all actual things, and for each $w_i \in \mathcal{W}_\mathcal{I}$, $\mathcal{D}_\mathcal{I}(w_i)$ becomes the set of things that *would be* surrogates of the actual things if w_i *were* a model of how things are. In this way an actualist who accepts modal terms as primitive can construct a model that *would be* isomorphic to $\mathcal{M}_\mathcal{I}$ if the latter existed.[63]

I will call this actualistic intended model $\mathcal{M}^{\mathcal{A}}_{\mathcal{I}} = \langle \mathcal{W}^{\mathcal{A}}_{\mathcal{I}}, @^{\mathcal{A}}_{\mathcal{I}}, \mathcal{R}^{\mathcal{A}}_{\mathcal{I}}, \mathcal{D}^{\mathcal{A}}_{\mathcal{I}}, \mathcal{I}^{\mathcal{A}}_{\mathcal{I}} \rangle$ ($\mathcal{A}$ for actualistic, $\mathcal{I}$ for intended). $\mathcal{W}^{\mathcal{A}}_{\mathcal{I}}$ is thus the set of all possible world surrogates, $@^{\mathcal{A}}_{\mathcal{I}}$ the actual world surrogate, $\mathcal{D}^{\mathcal{A}}_{\mathcal{I}}(@^{\mathcal{A}}_{\mathcal{I}})$ the set of all actual individual surrogates, and so on. Truth *simpliciter* is now just truth at $@^{\mathcal{A}}_{\mathcal{I}}$. Necessary truth is truth at all $w^{\mathcal{A}}_{\mathcal{I}} \in \mathcal{W}^{\mathcal{A}}_{\mathcal{I}}$ such that $@^{\mathcal{A}}_{\mathcal{I}} \mathcal{R}^{\mathcal{A}}_{\mathcal{I}} w^{\mathcal{A}}_{\mathcal{I}}$.

Given this terminology the sentence (Act3) of Section 6.1, which expresses actualism, is true but not necessary:

(Act3) $\forall x \mathsf{A}\mathcal{E}(x)$

(Act3) is true because it says that everything is actual, an obvious—indeed an a priori—truth. But on the plausible assumption that there might have been things that do not actually exist, it is not true at every $w^{\mathcal{A}}_{\mathcal{I}} \in \mathcal{W}^{\mathcal{A}}_{\mathcal{I}}$ such that $@^{\mathcal{A}}_{\mathcal{I}} \mathcal{R}^{\mathcal{A}}_{\mathcal{I}} w^{\mathcal{A}}_{\mathcal{I}}$, and hence not necessary. So (Act3) is an example of a contingent a priori truth.[64]

It should also be noted that because I define validity as general validity (truth at every world in every model) (Act3) is not valid. Were validity defined as real-world validity (truth at the actual-world element of every model) things would be different. (Act3) would be classified as valid, but the rule of necessitation would no longer preserve validity. In [11] and [12] I have argued at length that general validity has a better claim to the title of logical truth than does real-world validity.

8 Tableaus

In this section I develop a system of tableau proofs for SAT and its kindred systems.[65] In the Appendix I prove that these systems are sound and (weakly) complete with respect to the corresponding notions of logical consequence. These tableau systems follow closely those of Fitting and Mendelsohn [7].

8.1 The Basics

8.1.1 $\mathfrak{L}^\star$, the language of tableaus The language $\mathfrak{L}^\star$ is a slight extension of the language $\mathfrak{L}$ presented in Section 2. It differs from $\mathfrak{L}$ only in containing additional individual symbols called *grounded terms*. Grounded terms will be explained shortly, but first I need to describe the prefixes that appear in tableau nodes.

Each tableau node consists a formula of $\mathfrak{L}^\star$ preceded by a prefix, this prefix being a sequence of (numerals for) positive integers and the symbol @ separated by periods. The first symbol in a sequence is always 1 or @. Thus the following are sequences: 1, @.1.1.3, 1.4.5.2, 1.1.2. If σ is a prefix and n is a positive integer, then $\sigma.n$ is σ followed by a period followed by n. Thus if σ is 1.3.4 or @.1.1, and n is 2, then $\sigma.n$ is 1.3.4.2 or @.1.1.2, respectively. Intuitively, tableau prefixes play the role of the worlds of a model, and if nodes containing the prefixes σ and $\sigma.n$ appear on a tableau branch it means that $\sigma \mathcal{R} \sigma.n$.

The *grounded terms* of $\mathfrak{L}^\star$ are individual symbols that have the prefixes of tableau nodes as subscripts. Grounded terms are of two different kinds, *grounded names* and *parameters*. Grounded names are the same symbols as the individual constants of $\mathfrak{L}$ (the lower case letters a through j), each with a tableau prefix as subscript. Parameters are the lower case letters k through t (not used in $\mathfrak{L}$) with this same kind of subscript. Examples of grounded names are $a_{1.1.2}$, $b_{1.2.@.2.1}$, c_1, and $d_{1.2.@.2.1.@}$. Similarly, $k_{1.1.2}$, $p_{1.2.@.2.1}$, q_1, and $r_{1.2.@.2.1.@}$ are parameters.[66]

Parameters are used by Fitting and Mendelsohn [7] in tableau proofs. They are introduced by the tableau rule for negated universally quantified sentences (and by some other rules) to ensure that the individual term with which a sentence is instantiated is new to the branch in which the sentence appears. Although not essential (grounded names that are new to a branch could be used instead) they are convenient. They make it easy to tell at a glance which terms appearing on a branch are part of the sentences of $\mathfrak{L}$ with which the tableau began and which were introduced by applying tableau rules.

$\mathfrak{L}^\star$ is like $\mathfrak{L}$ except that the new grounded terms function like rigidly designating names. Syntactically, they behave exactly like the individual constants of $\mathfrak{L}$. That is, they may appear in atomic formulas and as the terms that follow predicate abstracts. (Thus if $\langle\lambda\alpha.\phi\rangle$ is a predicate abstract and τ_σ is a grounded term, $\langle\lambda\alpha.\phi\rangle(\tau_\sigma)$ is a formula of $\mathfrak{L}^\star$. The free variable occurrences in $\langle\lambda\alpha.\phi\rangle(\tau_\sigma)$ are the same as those in $\langle\lambda\alpha.\phi\rangle$.) But grounded terms never play the variable-binding role in quantification or predicate abstraction.

Semantically, the grounded terms of $\mathfrak{L}^\star$ behave like the variables of $\mathfrak{L}$ in that they are assigned their *designata* by a valuation $\mathcal{V}$ rather than by an interpretation $\mathcal{I}$ of a model $\mathcal{M} = \langle\mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I}\rangle$. So given a grounded term τ_σ, a model $\mathcal{M} = \langle\mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I}\rangle$, and a valuation $\mathcal{V}$ relative to $\mathcal{M}$, $\mathcal{V}(\tau_\sigma)$ is an element of $\mathcal{D}_\mathcal{M}$. Thus unlike an individual constant, whose denotation (if any) at a world is always drawn from the domain of that world, a grounded term is assigned a single, model-wide denotation from the domain of the model.[67]

With these stipulations in place, the definition of *satisfaction* of a formula of $\mathfrak{L}^\star$ by a valuation $\mathcal{V}$ at a world w of a model $\mathcal{M}$ is exactly like that given for $\mathfrak{L}$ in Section 3.4. Similarly for the definitions of the *satisfiability* of a formula of $\mathfrak{L}^\star$ and *satisfiability* of a set of formulas of $\mathfrak{L}^\star$. But Sections 3.5 and 3.6 define terms (truth at a world in a model, truth in a model, validity, consequence, and equivalence) that apply only to sentences of $\mathfrak{L}$. These terms are also undefined for formulas of $\mathfrak{L}^\star$ that are not sentences of $\mathfrak{L}$ for the reason given in Section 3.5. Formulas of $\mathfrak{L}^\star$ containing grounded terms are like formulas of $\mathfrak{L}$ containing free occurrences of variables. Calling such a formula true would be tantamount to prefixing the formula with possibilist universal quantifiers and calling the resulting sentence true (cf. Section 3.5).

8.1.2 The general structure of tableaus The purpose of tableaus is to generate proofs of logical consequence and validity for arguments and sentences, respectively, of $\mathfrak{L}$. Hence the formulas that appear as parts of the initial nodes of a tableau are always sentences of $\mathfrak{L}$.

A tableau test for an argument of $\mathfrak{L}$ with premises $\gamma_1, \gamma_2, \ldots, \gamma_n$ and conclusion ϕ begins with the following $n + 1$ nodes:

$$\begin{array}{l} 1\ \gamma_1 \\ 1\ \gamma_2 \\ \vdots \\ 1\ \gamma_n \\ 1\ \neg\phi \end{array}$$

A tableau test for a single sentence ϕ of $\mathfrak{L}$ begins with the single node

$$1\ \neg\phi$$

The result of applying a tableau rule to a node or nodes of a tableau will be the addition of one or more nodes, each a prefixed formula of $\mathfrak{L}^\star$, to the tableau. None of these formulas will contain free occurrences of variables.

Some tableau rules produce branching. The construction of a branch terminates when no more rules can be applied to any node on the branch, or when some prefixed formula and the negation of this formula with the same prefix ($\sigma\ \phi$ and $\sigma\ \neg\phi$) appear as nodes on the branch. In the latter case the branch is said to be *closed*. A branch that is not closed is *open*. A tableau is *closed* if all of its branches are closed.

A closed tableau beginning with $1\ \gamma_1$, $1\ \gamma_2, \ldots, 1\ \gamma_n$, and $1\ \neg\phi$ (where the γ_i ($1 \le i \le n$) and $\neg\phi$ are, as specified above, all sentences of $\mathfrak{L}$) is a *derivation of ϕ from* $\{\gamma_1, \gamma_2, \ldots, \gamma_n\}$. If Γ is a set of sentences of $\mathfrak{L}$, each γ_i ($1 \le i \le n$) $\in \Gamma$, and there is a derivation of ϕ from $\{\gamma_1, \gamma_2, \ldots, \gamma_n\}$, ϕ is said to be *derivable from* Γ (abbreviated $\Gamma \vdash \phi$). Similarly, a closed tableau beginning with $1\ \neg\phi$ (where, again, ϕ is a sentence of $\mathfrak{L}$) is a *proof* of ϕ. If there is such a proof, ϕ is said to be a *theorem* (abbreviated $\vdash \phi$).

In the Appendix I prove that a tableau is closed if and only if the sentences of $\mathfrak{L}$ with which it begins are inconsistent (i.e., there is no interpretation under which all of these sentences are true). Thus validity coincides with theoremhood and, for finite sets of premises, logical consequence coincides with derivability. The foregoing holds for each of the five systems under consideration. These proofs are based on the fact that an open tableau branch mimics a model *cum* valuation of $\mathfrak{L}^\star$. The intuitive idea is that node prefixes designate worlds (@ designates the actual world), $\sigma\ \phi$ asserts that formula ϕ is satisfied by the valuation at world σ, and the appearance of $\sigma.n$ as a node prefix indicates that $\sigma \mathcal{R} \sigma.n$.

8.2 Tableau rules

8.2.1 Truth-functional connectives Rules for conjunction, negated conjunction, and double negation are standard and can be represented as follows:

$$\frac{\sigma\ \phi \wedge \psi}{\begin{array}{l}\sigma\ \phi\\ \sigma\ \psi\end{array}} \qquad \frac{\sigma\ \neg(\phi \wedge \psi)}{\sigma\ \neg\phi \mid \sigma\ \neg\psi} \qquad \frac{\sigma\ \neg\neg\phi}{\sigma\ \phi}$$

The conjunction and double negation rules add the indicated nodes to each branch in which the top node appears. The negated conjunction rule produces branching in the usual way. Rules for the other truth-functional connectives are also standard and can be derived from those given here.

8.2.2 Modal and actuality connectives The SAT rules for necessity and negated necessity are the standard ones for T. The negated necessity rule is

$$\frac{\sigma\ \neg\Box\phi}{\sigma.n\ \neg\phi} \quad \sigma.n \text{ is new to the branch}$$

For each branch containing $\sigma\ \neg\Box\phi$ this rule is applied just once, and integer n is chosen so that prefix $\sigma.n$ is new to the branch. The necessity rule for SAT takes two forms.

$$(1)\quad \frac{\sigma\ \Box\phi}{\sigma\ \phi}$$

(2) $\dfrac{\sigma \ \Box\phi}{\sigma.n \ \phi}$ if the prefix $\sigma.n$ appears somewhere on the branch

Form 2 generally adds multiple nodes to a branch. And as new prefixes come to appear on a branch form 2 is reapplied, using these prefixes, to nodes to which it has already been applied using other prefixes.

Tableau rules for the other logics under consideration are also standard. SAK uses only form 2 of the necessity rule. For SAB, SAS4, and SAS5 three additional forms are needed.

(3) $\dfrac{\sigma.n \ \Box\phi}{\sigma \ \phi}$

(4) $\dfrac{\sigma \ \Box\phi}{\sigma.n \ \Box\phi}$ if the prefix $\sigma.n$ appears somewhere on the branch

(5) $\dfrac{\sigma.n \ \Box\phi}{\sigma \ \Box\phi}$

SAB retains forms 1 and 2 and adds form 3. SAS4 uses forms 1 and 4. SAS5 uses forms 1, 4, and 5. Since $\Diamond$ and $\Box$ are interdefinable, rules for $\Diamond$ and $\neg\Diamond$ can be derived from the ones given here.[68]

Rules for the actuality connective reflect the idea that a sentence is actually true at a world just in case it is true at the actual world.

$$\frac{\sigma \ \mathsf{A}\phi}{@ \ \phi} \qquad\qquad \frac{\sigma \ \neg\mathsf{A}\phi}{@ \ \neg\phi}$$

8.2.3 Quantifiers Where α is a variable, let $\psi(\alpha)$ be a formula of $\mathfrak{L}^\star$ containing at least one free occurrence of α (but no free occurrence of any other variable). Let τ_σ and π_σ be a grounded term and a parameter, respectively, each subscripted with the tableau prefix σ.[69] The formula $\psi(\tau_\sigma)$ stands for the result of replacing all free occurrences of α in $\psi(\alpha)$ with τ_σ, and similarly for $\psi(\pi_\sigma)$ and π_σ. The universal quantifier and negated universal quantifier rules are as follows, where τ_σ is any grounded term that already appears on the branch, and π_σ is a parameter that is new to the branch.

$\dfrac{\sigma \ \forall\alpha\psi(\alpha)}{\sigma \ \psi(\tau_\sigma)}$ τ_σ appears somewhere on the branch $\qquad$ $\dfrac{\sigma \ \neg\forall\alpha\psi(\alpha)}{\sigma \ \neg\psi(\pi_\sigma)}$ π_σ is new to the branch

For each branch containing the node $\sigma \ \forall\alpha\psi(\alpha)$, and each grounded term τ_σ that appears as part of a node on that branch, the universal quantifier rule adds $\sigma \ \psi(\tau_\sigma)$ to the branch. (If no grounded term appears with the tableau prefix σ as subscript in any node on a branch containing $\sigma \ \forall\alpha\psi(\alpha)$, the universal quantifier rule is not applied to $\sigma \ \forall\alpha\psi(\alpha)$ on that branch. This reflects the fact that the domain of a world may be empty.) For each branch containing the node $\sigma \ \neg\forall\alpha\psi(\alpha)$, the negated universal quantifier rule is applied just once, and the parameter π_σ must be chosen so that it is new to the branch.

The universal quantifier rule (and the dual negated existential quantifier rule) reflect the fact that SAT is an actualistic free logic. Actualism requires quantification to be world-relative, and a free logic permits instantiation at a world only with terms that denote objects in the domain of that world. The negated universal quantifier rule (and the dual existential rule) also reflect actualism in requiring that the term being instantiated denote an object in the domain of the world in question.

8.2.4 Predicate abstraction According to serious actualism, an attributive sentence is true at a world only if the object denoted by the term to which the abstract applies exists at that world. To fully reflect this requirement, the (unnegated) predicate abstraction rule takes three forms. Form 1 "grounds" or "rigidifies" an individual constant governed by a predicate abstract, thus transforming it into a grounded name. Forms 2 and 3 apply to terms that are already grounded. But since the subscript of a grounded term may not be the same as the prefix of the node in which the sentence containing it appears, two rules are needed.

Suppose σ and σ_1 are distinct prefixes, ι is an individual constant, and τ_σ and τ_{σ_1} are grounded terms. Forms 1 and 2 instantiate the predicate abstract with ι_σ and τ_σ, respectively. Form 3 instantiates it with τ_{σ_1}, and it introduces a second node. This second node reflects the fact that, since $\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_{\sigma_1})$ is true at world σ, τ_{σ_1} denotes an object that exists in σ as well as in σ_1.[70]

(1) $\dfrac{\sigma\ \langle\lambda\alpha.\psi(\alpha)\rangle(\iota)}{\sigma\ \psi(\iota_\sigma)}$ ι is an individual constant of $\mathfrak{L}$ (ι_σ is thus a grounded name of $\mathfrak{L}^\star$)

(2) $\dfrac{\sigma\ \langle\lambda\alpha.\psi(\alpha)\rangle(\tau_\sigma)}{\sigma\ \psi(\tau_\sigma)}$ τ_σ is a grounded term

(3) $\dfrac{\sigma\ \langle\lambda\alpha.\psi(\alpha)\rangle(\tau_{\sigma_1})}{\begin{array}{l}\sigma\ \psi(\tau_{\sigma_1})\\ \sigma\ (\pi_\sigma = \tau_{\sigma_1})\end{array}}$ τ_{σ_1} is a grounded term, and $\sigma \neq \sigma_1$; π_σ is a parameter that is new to the branch

It should be noted that form 1 of the predicate abstraction rule and form 1 of the atomic formula rule are the only rules that introduce grounded names into tableaus.

The rule for negated predicate abstraction also has three forms. Each form applies to a tableau node that, except for the negation sign, is the same as the node in the correspondingly numbered form of the positive predicate abstraction rule.[71]

(1) $\dfrac{\sigma\ \neg\langle\lambda\alpha.\psi(\alpha)\rangle(\iota)}{\sigma\ \neg\psi(\iota_\sigma)}$ ι is an individual constant, and the grounded name ι_σ appears somewhere on the branch

(2) $\dfrac{\sigma\ \neg\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_\sigma)}{\sigma\ \neg\psi(\tau_\sigma)}$ τ_σ is a grounded term

(3) $\dfrac{\begin{array}{l}\sigma\ \neg\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_{\sigma_1})\\ \sigma_2\ (\upsilon_\sigma = \tau_{\sigma_1})\end{array}}{\sigma\ \neg\psi(\tau_{\sigma_1})}$ τ_{σ_1} is a grounded term, and $\sigma \neq \sigma_1$; σ_2 is any prefix; υ_σ is any term grounded with σ

In form 1 the presence of the grounded name ι_σ somewhere on the branch assures us that the individual constant ι designates, at world σ, a member of the domain of σ. (The individual constant ι need not appear with subscript σ in the premise of the rule.) If ι_σ did not appear on the branch, we could not be sure that ι designates anything in the domain of σ. But then we could not be sure that $\neg\psi(\iota)$ is true in σ. For suppose θ is a primitive predicate, ι does not designate at σ, and $\psi(\iota)$ is $\neg\theta(\iota)$. Then $\theta(\iota)$ is false at σ and so is $\neg\psi(\iota)$.

The rationale for form 2 should be obvious. Form 3 is similar to from 1 in that the predicate abstract applies to a term that is not grounded in σ. Thus assurance is needed that τ_{σ_1} really does denote an object in the domain of world σ. The appearance of σ_2 $(\upsilon_\sigma = \tau_{\sigma_1})$ somewhere on the branch provides this. If τ_{σ_1} did not denote an object in the domain of world σ, $\neg\psi(\tau_{\sigma_1})$ might be false at σ, as in the case of form 1.

8.2.5 Identity The self-identity rule reflects the fact that in SAT a sentence of the general form $\beta = \beta$ is true at a world only if β designates an object that exists at that world. Since the appearance of a grounded term τ_σ in any node on a branch indicates that the object it denotes exists at world σ, that object must be self-identical at σ. Thus where τ_σ is any grounded term, the self-identity rule takes the following form:[72]

$$\frac{}{\sigma\ (\tau_\sigma = \tau_\sigma)} \qquad \tau_\sigma \text{ appears somewhere on the branch}$$

As usual, the substitutivity of identity rule sanctions substituting coreferential rigid designators for each other, but in SAT it requires some new notation and takes a rather complicated form. Suppose σ_1, σ_2, σ_3, and σ_4 are prefixes of branch nodes (not necessarily distinct), τ_{σ_3} and υ_{σ_4} are grounded terms, and ϕ and ψ are formulas of $\mathfrak{L}^\star$ that contain no free occurrences of any variable. Suppose also that ϕ and ψ are alike except that ψ contains occurrences of υ_{σ_4} at one or more places where ϕ contains occurrences of τ_{σ_3}. If $\tau_{\sigma_3} = \upsilon_{\sigma_4}$ holds at any world, $\mathcal{V}$ assigns the same member of $\mathcal{D}_\mathcal{M}$ to τ_{σ_3} and υ_{σ_4}. So if $\tau_{\sigma_3} = \upsilon_{\sigma_4}$ holds at σ_1 and ϕ holds at σ_2, ψ also holds at σ_2. (And the object denoted by τ_{σ_3} and υ_{σ_4} exists in the domains of worlds σ_1, σ_3, and σ_4, and possibly in the domain of world σ_2.) The substitutivity rule is thus

$$\frac{\begin{array}{l}\sigma_1\ \tau_{\sigma_3} = \upsilon_{\sigma_4}\\ \sigma_2\ \phi\end{array}}{\sigma_2\ \psi}$$

This rule is the only one that has two premises. Thus it can be applied only when both of these premises appear on the same branch.

8.2.6 Atomic formulas of $\mathfrak{L}^\star$ An atomic formula is satisfied by an n-tuple of objects at a world only if all those objects exist at that world. So each individual term that appears in such a formula, whether grounded or ungrounded, must designate an object that exists at that world. This fact must be reflected in the tableau rules, and so, unlike other tableau systems, SAT has a rule for atomic formulas. The rule has three forms, and some additional notation will facilitate understanding. Where θ is an n-ary predicate and τ and υ are terms, $\theta(\ldots\tau\ldots)$ will stand for an atomic formula in which τ has one or more occurrences, and $\theta(\ldots\upsilon\ldots)$ for the result of replacing each occurrence of τ in $\theta(\ldots\tau\ldots)$ with υ.[73]

The first two forms of the rule are:

$$(1)\quad \frac{\sigma\ \theta(\ldots\iota\ldots)}{\sigma\ \theta(\ldots\iota_\sigma\ldots)} \quad \begin{array}{l}\iota \text{ is an individual constant of } \mathfrak{L}\\ (\iota_\sigma \text{ is thus a grounded name of } \mathfrak{L}^\star)\end{array}$$

$$(2)\quad \frac{\sigma\ \theta(\ldots\tau_{\sigma_1}\ldots)}{\sigma\ (\pi_\sigma = \tau_{\sigma_1})} \quad \begin{array}{l}\tau_{\sigma_1} \text{ is a grounded term, and } \sigma \neq \sigma_1\\ \pi_\sigma \text{ is a parameter that is new to the branch}\end{array}$$

These two forms of the rule reflect the fact that an atomic sentence is true at a world only if each of the terms it contains denotes an object that exists at that world. As noted in Section 8.2.4, form 1 of the predicate abstraction rule and form 1 of the atomic formula rule are the only rules that introduce grounded names into tableaus.[74]

The third form of the atomic formula rule is unlike any other tableau rule in that it allows a subscript to be removed from a grounded name.[75] It is also unlike the first two forms in that it may be applied to some but not all occurrences of an individual symbol in an atomic formula. Let ι be an individual constant, and let ι_σ be that same individual constant subscripted with σ. ι_σ is thus a grounded name. If ι_σ appears in an atomic formula as part of a tableau node with prefix σ, one or more occurrences of ι_σ may be replaced with ι. To call attention to the fact that form 3 of the rule allows replacement of some but not all occurrences of a term in an atomic formula, I will use a new notation. Let $\theta(\iota_\sigma)$ and $\theta(\iota)$ stand for atomic formulas that are alike except that $\theta(\iota)$ contains occurrences of ι at one or more places where $\theta(\iota_\sigma)$ contains occurrences of ι_σ.

The third form of the atomic formula rule is:

$$(3)\quad \frac{\sigma\ \theta(\iota_\sigma)}{\sigma\ \theta(\iota)} \quad \begin{array}{l}\iota_\sigma \text{ is a grounded name of } \mathfrak{L}^\star \\ (\iota \text{ is thus an individual constant of } \mathfrak{L})\end{array}$$

Like the first two forms, form 3 also reflects the fact that an atomic sentence is true at a world only if each of the terms it contains denotes an object that exists at that world.[76]

8.3 Examples of tableaus In the following examples tableau rules are referred to by abbreviated names. For example, $\neg\rightarrow$ denotes the negated conditional rule, $\neg\lambda 2$ denotes the second form of the negated abstraction rule, and Subs= denotes the substitution of identity rule.

8.3.1 The Converse Barcan Formula (CBF) The Converse Barcan Formula is SAT valid, as the following tableau for a simple instance of (CBF) shows.

1	$\neg\big(\exists x\Diamond\langle\lambda y.Fy\rangle(x) \rightarrow \Diamond\exists x\langle\lambda y.Fy\rangle(x)\big)$	1.
1	$\exists x\Diamond\langle\lambda y.Fy\rangle(x)$	2. (From 1 by $\neg\rightarrow$.)
1	$\neg\Diamond\exists x\langle\lambda y.Fy\rangle(x)$	3. (From 1 by $\neg\rightarrow$.)
1	$\Diamond\langle\lambda y.Fy\rangle(p_1)$	4. (From 2 by $\exists$.)
1.1	$\langle\lambda y.Fy\rangle(p_1)$	5. (From 4 by $\Diamond$.)
1.1	Fp_1	6. (From 5 by $\lambda 3$.)
1.1	$(q_{1.1} = p_1)$	7. (From 5 by $\lambda 3$.)
1.1	$\neg\exists x\langle\lambda y.Fy\rangle(x)$	8. (From 3 by $\neg\Diamond 2$.)
1.1	$\neg\langle\lambda y.Fy\rangle(q_{1.1})$	9. (From 8 by $\neg\exists$, using $q_{1.1}$.)
1.1	$\neg Fq_{1.1}$	10. (From 9 by $\neg\lambda 2$.)
1.1	$\neg Fp_1$	11. (From 7 and 10 by Subs =.)
	$\times$	

The use of variables other than x and y would not affect this proof. And since none of the rules for atomic formulas have been applied, replacing Fy with any complex formula $\phi(y)$ would not avoid closure.

*8.3.2 An imposter (*faux *CBF)* In Section 4.4 I dubbed formulas of the form

$$\Box\forall\alpha_1\big\langle\lambda\alpha_2.\phi(\alpha_2)\big\rangle(\alpha_1) \rightarrow \forall\alpha_1\Box\big\langle\lambda\alpha_2.\phi(\alpha_2)\big\rangle(\alpha_1)$$

(*faux* CBF). Not all formulas of this form are SAT valid, as the following tableau shows.

1	$\neg(\Box\forall x\langle\lambda y.Fy\rangle(x) \to \forall x\Box\langle\lambda y.Fy\rangle(x))$	1.
1	$\Box\forall x\langle\lambda y.Fy\rangle(x)$	2. (From 1 by $\neg\to$.)
1	$\neg\forall x\Box\langle\lambda y.Fy\rangle(x)$	3. (From 1 by $\neg\to$.)
1	$\forall x\langle\lambda y.Fy\rangle(x)$	4. (From 2 by $\Box$1.)
1	$\neg\Box\langle\lambda y.Fy\rangle(p_1)$	5. (From 3 by $\neg\forall$.)
1	$\langle\lambda y.Fy\rangle(p_1)$	6. (From 4 by $\forall$, using p_1.)
1	Fp_1	7. (From 6 by λ1.)
1.1	$\neg\langle\lambda y.Fy\rangle(p_1)$	8. (From 5 by $\neg\Box$.)
1.1	$\forall x\langle\lambda y.Fy\rangle(x)$	9. (From 2 by $\Box$2.)
	$\circ$	

8.3.3 Predicate abstraction applied to atomic formulas is inessential A main theme of this article is the importance of predicate abstraction for a coherent account of logical form in quantified modal logic. The tableaus in this section and the next illustrate this by showing that while

$$\langle\lambda x.Fx\rangle(a) \leftrightarrow Fa$$

is SAT valid,

$$\langle\lambda x.\neg Fx\rangle(a) \leftrightarrow \neg Fa$$

is not. Here is the first tableau.

1 $\neg(\langle\lambda x.Fx\rangle(a) \leftrightarrow Fa)$ 1.

1 $\langle\lambda x.Fx\rangle(a)$	2.	1 $\neg\langle\lambda x.Fx\rangle(a)$	4.
1 $\neg Fa$	3.	1 Fa	5.
1 Fa_1	6.	1 Fa_1	8.
1 Fa	7.	1 $\neg Fa_1$	9.
$\times$		$\times$	

In the left branch λ1 applied to line 2 yields line 6, and Atomic 3 applied to line 6 yields line 7. In the right branch Atomic 1 applied to line 5 yields line 8, which is the first node in the right branch that contains a_1. The appearance of a_1 makes it possible to apply $\neg\lambda$1 to line 4, which yields line 9 and closure. If a_1 did not appear on the right branch, $\neg\lambda$1 could not be applied to line 4. And if the only appearances of a on the right branch were as parts of nonatomic sentences, Atomic 1 could not be applied and a_1 would not come to appear on the branch. So the right branch closes only because a appears on it as part of an *atomic* sentence. The next section contains a tableau (for a sentence with a similar form) that fails to close for exactly this reason.

8.3.4 Predicate abstraction applied to nonatomic formulas is *essential* If Fx and Fa in $\langle\lambda x.Fx\rangle(a) \leftrightarrow Fa$ are replaced by nonatomic formulas $\phi(x)$ and $\phi(a)$, the tableau for the resulting formula may not close. For example, consider the tableau for $\langle\lambda x.\neg Fx\rangle(a) \leftrightarrow \neg Fa$.

$$1\ \neg\big(\langle\lambda x.\neg Fx\rangle(a) \leftrightarrow \neg Fa\big)\quad 1.$$

$1\ \langle\lambda y.\neg Fy\rangle(a)$ 2.	$1\ \neg\langle\lambda x.\neg Fx\rangle(a)$ 4.
$1\ \neg\neg Fa$ 3.	$1\ \neg Fa$ 5.
$1\ Fa$ 6.	$\circ$
$1\ Fa_1$ 7.	
$1\ \neg Fa_1$ 8.	
$\times$	

The left branch is essentially the same as that of the tableau in the previous section except for the use of double negation. But on the right branch no rule can be applied to line 4 or line 5. Since the sentence on line 5 is not atomic the grounded name a_1 cannot be introduced, and in its absence $\neg\lambda 1$ cannot be applied to line 4. The right branch is thus open.

8.3.5 Other imposters Consider the formulas $\exists x\Diamond Fx \rightarrow \Diamond\exists xFx$ and $\exists x\Diamond\neg Fx \rightarrow \Diamond\exists x\neg Fx$. In most systems of quantified modal logic both are considered instances of the Converse Barcan Formula, and either both are valid or both invalid. But in SAT neither is an instance of (CBF) because neither contains predicate abstracts. Yet in view of Sections 8.3.1, 8.3.3, and 8.3.4, it should not be surprising that the tableau for $\exists x\Diamond Fx \rightarrow \Diamond\exists xFx$ closes but the one for $\exists x\Diamond\neg Fx \rightarrow \Diamond\exists x\neg Fx$ does not. (I leave it to the reader to verify that the former tableau is nearly identical to the one given in Section 8.3.1 but requires only nine steps.) The tableau for $\exists x\Diamond\neg Fx \rightarrow \Diamond\exists x\neg Fx$ follows.

1	$\neg(\exists x\Diamond\neg Fx \rightarrow \Diamond\exists x\neg Fx)$	1.
1	$\exists x\Diamond\neg Fx$	2. (From 1 by $\neg\rightarrow$.)
1	$\neg\Diamond\exists x\neg Fx$	3. (From 1 by $\neg\rightarrow$.)
1	$\Diamond\neg Fp_1$	4. (From 2 by $\exists$.)
1.1	$\neg Fp_1$	5. (From 4 by $\Diamond$.)
1.1	$\neg\exists x\neg Fx$	6. (From 3 by $\neg\Diamond 2$.)
1	$\neg\exists x\neg Fx$	7. (From 3 by $\neg\Diamond 1$.)
1	$\neg\neg Fp_1$	8. (From 7 by $\neg\exists$.)
1	Fp_1	9. (From 8 by$\neg\neg$.)
	$\circ$	

It is instructive to compare line 5 of the tableau in Section 8.3.1 with line 5 of the above tableau. Each is inferred from line 4 of its respective tableau by the $\Diamond$ rule. In the tableau of Section 8.3.1 the $\lambda 3$ rule is applied to line 5, which leads ultimately to closure. But here no rule can be applied to line 5, and the tableau does not close.[77]

*8.3.6 A tableau for an instance of (*faux *UI)* In the discussion of the logical properties of quantifiers in Section 4.3 I mentioned sentences of a form I called (*faux* UI), and I gave an instance of it that I claimed is not SAT valid. Here is the sentence:

$$\forall x\Box Sx \rightarrow \big(\mathcal{E}(d) \rightarrow \Box Sd\big).$$

Its invalidity is shown by the following tableau.

1	$\neg\big(\forall x \Box Sx \rightarrow \big(\exists y(y = d) \rightarrow \Box Sd\big)\big)$	1.
1	$\forall x \Box Sx$	2. (From 1 by $\neg \rightarrow$.)
1	$\neg\big(\exists y(y = d) \rightarrow \Box Sd\big)$	3. (From 1 by $\neg \rightarrow$.)
1	$\exists y(y = d)$	4. (From 3 by $\neg \rightarrow$.)
1	$\neg\Box Sd$	5. (From 3 by $\neg \rightarrow$.)
1	$(p_1 = d)$	6. (From 4 by $\exists$.)
1	$(p_1 = d_1)$	7. (From 6 by Atomic 1.)
1	$\Box Sp_1$	8. (From 2 by $\forall$.)
1	$\Box Sd_1$	9. (From 2 by $\forall$.)
1	Sp_1	10. (From 8 by $\Box$1.)
1	Sd_1	11. (From 9 by $\Box$1.)
1.1	$\neg Sd$	12. (From 5 by $\neg\Box$.)
1.1	Sp_1	13. (From 8 by $\Box$2.)
1.1	Sd_1	14. (From 9 by $\Box$2.)
1.1	$(q_{1.1} = d_1)$	15. (From 13 by Atomic 2.)
1.1	$(r_{1.1} = d_1)$	16. (From 14 by Atomic 2.)
	$\circ$	

This tableau is not complete. The rules for self-identity, substitution of identity, and universal quantifier can be applied in several additional ways. But the tableau will not close. Indeed its single branch suggests a simple countermodel to *faux* (UI). In this model, $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$, $\mathcal{W}$ contains two worlds, w_1 and $w_{1.1}$, such that $w_1 \mathcal{R} w_1$, $w_{1.1} \mathcal{R} w_{1.1}$, and $w_1 \mathcal{R} w_{1.1}$; @ may be either w_1 or $w_{1.1}$. The domains of these worlds will be the same one-membered set. Anticipating the strategy employed in the completeness lemma (see Section A.2), let $\mathcal{D}(w_1) = \mathcal{D}(w_{1.1}) = \{\{p_1, d_1, q_{1.1}, r_{1.1}\}\}$. (Where σ_1, σ_2, and σ_3 are tableau prefixes and τ_{σ_2} and υ_{σ_3} are grounded terms, $\{\{p_1, d_1, q_{1.1}, r_{1.1}\}\}$ is the set of equivalence classes of grounded terms determined by formulas of the form $(\tau_{\sigma_2} = \upsilon_{\sigma_3})$ such that, for some σ_1, σ_1 $(\tau_{\sigma_2} = \upsilon_{\sigma_3})$ is a node on the branch.) The extension of the predicate S at each world is the domain of that world, which in this case is also the domain of the model. That is, $\mathcal{I}(S, w_1) = \mathcal{I}(S, w_{1.1}) = \{\{p_1, d_1, q_{1.1}, r_{1.1}\}\}$. The individual constant d denotes $\{p_1, d_1, q_{1.1}, r_{1.1}\}$ at w_1, but does not denote at $w_{1.1}$. That is, $\mathcal{I}(d, w_1) = \{p_1, d_1, q_{1.1}, r_{1.1}\}$, but $\mathcal{I}(d, w_{1.1})$ is undefined. For any valuation $\mathcal{V}$ relative to this model $\mathcal{M}$ (since the domain of the model has only one member there is only one such valuation) $\mathcal{M}_{\mathcal{V}}(\forall x \Box Sx \rightarrow (\mathcal{E}(d) \rightarrow \Box Sd)) = 0$.[78]

8.3.7 A tableau for (Act3) I argued in Section 6.1 that actualism is properly expressed in SAT as

(Act3) $\forall x \mathsf{A} \mathcal{E}(x)$.

Although (Act3) is true in the intended model it is not SAT valid, as the following tableau shows.

1	$\neg\forall x \mathsf{A} \exists y(y = x)$	1.
1	$\neg\mathsf{A}\exists y(y = p_1)$	2. (From 1 by $\neg\forall$.)
@	$\neg\exists y(y = p_1)$	3. (From 2 by $\neg\mathsf{A}$.)
	$\circ$	

The negated existential rule cannot be applied to line 3 because no parameter with subscript @ appears on the branch. But this tableau can be used to provide insight into the distinction between logical truth and a priori truth. For suppose we want to prove that (Act3) is actually true. A plausible way of proceeding would be to use the tableau method but start by supposing that (Act3) is false in the actual world. So we would get:

$$
\begin{array}{lll}
@ & \neg\forall x \mathsf{A}\exists y(y = x)) & 1. \\
@ & \neg\mathsf{A}\exists y(y = p_{@})) & 2. \text{ (From 1 by } \neg\forall.) \\
@ & \neg\exists y(y = p_{@}) & 3. \text{ (From 2 by } \neg\mathsf{A}.) \\
@ & \neg(p_{@} = p_{@}) & 4. \text{ (From 3 by } \neg\exists.) \\
@ & (p_{@} = p_{@}) & 5. \text{ (By Self} =.) \\
 & \times &
\end{array}
$$

The assumption that (Act3) is actually false leads to a contradiction. So if one knows, independently of experience (as we all do), that the world one inhabits is the actual world, then the reasoning embodied in the preceding tableau is an a priori proof of (Act3).[79] (Act3) is true a priori even though it is not necessary and not logically true.

9 Conclusion

I have shown that an actualistic and seriously actualistic quantified modal logic with desirable formal and philosophical properties is possible. These properties include, most notably, uniform substitution of complex predicate abstracts for simple ones in logical truths, replacement of logically equivalent subformulas within sentences *salva veritate*, and a sound and complete proof system. Indeed by defining SAT and proving that it has these properties I have shown that such a logic is actual. I have also compared SAT with several other treatments of the issues with which it deals. I believe it deserves serious consideration in ongoing logical and philosophical discussions of modality.

Appendix: Soundness and Completeness

It is not difficult to verify that the soundness and completeness results given here for SAT hold also for SAK, SAB, SAS4, and SAS5.

A.1 Soundness The intuitive idea behind the soundness proof is simple. To demonstrate soundness of the tableau proof system it is sufficient to show that if a set of prefixed formulas of $\mathfrak{L}^\star$ appearing on a tableau branch is satisfiable in a certain sense, then the set obtained by adding prefixed formulas that result from applying tableau rules to members of that set is satisfiable in this same sense. Satisfiability of a set of formulas of $\mathfrak{L}$ was defined in Section 3.4. The expanded notion of satisfiability of a set of prefixed formulas of $\mathfrak{L}^\star$, used in the soundness proof, is defined as follows.[80]

A set S of prefixed formulas of $\mathfrak{L}^\star$ is *satisfiable in a model* $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I}\rangle$ *relative to a valuation* $\mathcal{V}$ if there is a function Θ that assigns to each prefix σ occurring in S a world $\Theta(\sigma) \in \mathcal{W}$ such that:

(1) For every prefix σ in S $\Theta(\sigma)\mathcal{R}\Theta(\sigma)$, and if $\sigma.n$ also appears as a prefix in S, $\Theta(\sigma)\mathcal{R}\Theta(\sigma.n)$.

(2) If @ occurs as a prefix in S, $\Theta(@) = @$.

(3) If the parameter π_σ occurs in S, then $\mathcal{V}(\pi_\sigma) \in \mathcal{D}(\Theta(\sigma))$.

(4) Let ι be an individual constant of $\mathfrak{L}$, and let ι_σ be the grounded name of $\mathfrak{L}^\star$ that results from subscripting ι with the tableau prefix σ. If ι and ι_σ both appear in S, then $\mathcal{I}(\iota, \Theta(\sigma))$ is defined, $\mathcal{I}(\iota, \Theta(\sigma)) \in \mathcal{D}(\Theta(\sigma))$, and $\mathcal{V}(\iota_\sigma) = \mathcal{I}(\iota, \Theta(\sigma))$.[81]

(5) For any prefix σ and formula ϕ, if $\sigma\ \phi$ is in S (i.e., if the pair consisting of ϕ prefixed with σ is a member of S), then $\mathcal{M}_\mathcal{V}(\phi, \Theta(\sigma)) = 1$.

A *branch of a tableau is satisfiable* if the set of prefixed formulas appearing on that branch is satisfiable in some model relative to some valuation. A *tableau is satisfiable* if one or more of its branches is.

The connection between a set of prefixed formulas appearing on an open tableau branch and a model-valuation pair $(\mathcal{M}, \mathcal{V})$ that satisfies them, relative to a function Θ, is straightforward. Given a branch that is satisfiable in this sense, Θ maps tableau prefixes to worlds of $\mathcal{M}$ in a way that induces on the prefixes the roles of relation $\mathcal{R}$ and world @ (clauses 1 and 2). Clauses 3 and 4 require that a term grounded with a given prefix denotes an object in the domain of the world associated with that prefix.[82] Clause 5 requires that each prefixed formula on the branch be satisfied at the world corresponding to its prefix.

The soundness lemma is the heart of the soundness proof. Once it is established, the soundness theorem itself follows easily. Proof of the soundness lemma (and of the completeness lemma in the next section) will be facilitated by the following substitution lemma for $\mathfrak{L}^\star$.

Substitution lemma (for $\mathfrak{L}^\star$):

Let $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$ be a model of $\mathfrak{L}^\star$, and let $\mathcal{V}$ be a valuation relative to $\mathcal{M}$. Where $\psi(\alpha)$ is a formula of $\mathfrak{L}^\star$ containing at least one free occurrence of the variable α (but no free occurrence of any other variable), and $w \in \mathcal{W}$:

Part 1. Let τ_σ be a grounded term of $\mathfrak{L}^\star$, $\mathcal{V}(\tau_\sigma) \in \mathcal{D}(w)$, and $\psi(\tau_\sigma)$ the result of replacing all free occurrences of α in $\psi(\alpha)$ with τ_σ. Suppose $\mathcal{U}$ is the α-variant-at-w of $\mathcal{V}$ such that $\mathcal{U}(\alpha) = \mathcal{V}(\tau_\sigma)$. Then $\mathcal{M}_\mathcal{U}(\psi(\alpha), w) = \mathcal{M}_\mathcal{V}(\psi(\tau_\sigma), w)$.

Part 2. Let ι be an individual constant of $\mathfrak{L}$, ι_σ be a grounded name of $\mathfrak{L}^\star$ (that is, ι_σ is the grounded name of $\mathfrak{L}^\star$ that results from subscripting ι with the tableau-node prefix σ), $\mathcal{V}(\iota_\sigma) \in \mathcal{D}(w)$, and $\psi(\iota_\sigma)$ the result of replacing all free occurrences of α in $\psi(\alpha)$ with ι_σ. Suppose $\mathcal{U}$ is the α-variant-at-w of $\mathcal{V}$ such that $\mathcal{U}(\alpha) = \mathcal{V}(\iota_\sigma)$. Then $\mathcal{M}_\mathcal{U}(\psi(\alpha), w) = \mathcal{M}_\mathcal{V}(\psi(\iota_\sigma), w)$.

Proof of the Part 1 is straightforward by induction on the complexity of $\psi(\alpha)$. Since a grounded name is a grounded term, Part 2 is a special case of Part 1. It is convenient to have the Lemma stated in this form in some of the following proofs. ■

Soundness lemma: Tableau rules preserve satisfiability.

If a tableau rule is applied to a satisfiable tableau, the result is another satisfiable tableau.

Proof It is sufficient to restrict attention to application of a tableau rule to a node or nodes that appear in a satisfiable branch of a tableau, and to show that at least one satisfiable branch results from application of the rule. So suppose $\mathcal{B}$ is a branch of a tableau $\mathcal{T}$, S is the set of prefixed formulas of $\mathfrak{L}^\star$ that appear on $\mathcal{B}$, and there is a model $\mathcal{M}$, a valuation $\mathcal{V}$, and a function Θ that together satisfy 1–5. $\mathcal{B}$ is thus a satisfiable branch of $\mathcal{T}$. We want to show that whenever a rule is applied to a member of $\mathcal{B}$, the result is at least one branch $\mathcal{B}'$ that extends $\mathcal{B}$ and is such that, for the set

S' of prefixed formulas of $\mathfrak{L}^\star$ that appear on $\mathcal{B}'$, there is a model $\mathcal{M}'$, a valuation $\mathcal{V}'$, and a function Θ' that together satisfy 1–5.

Proofs of the cases for truth-functional and modal connectives are simple and straightforward, as in Fitting and Mendelsohn [7, pp. 58–9]. The cases for the actuality connective are trivial.

For the universal quantifier case, assume that the prefixed formula $\sigma \;\; \forall\alpha\psi(\alpha)$ is in S, that clauses 1–5 hold, and hence that $\mathcal{M}_{\mathcal{V}}(\forall\alpha\psi(\alpha), \Theta(\sigma)) = 1$. By clause 10 of Section 3.4, $\mathcal{M}_{\mathcal{U}}(\psi(\alpha), \Theta(\sigma)) = 1$, for every valuation $\mathcal{U}$ that is an α-variant-at-$\Theta(\sigma)$ of $\mathcal{V}$. By 3 and 4, $\mathcal{V}(\tau_\sigma) \in \mathcal{D}(\Theta(\sigma))$, for each grounded term τ_σ that appears on $\mathcal{B}$ (i.e., for each parameter or grounded name with subscript σ that appears on $\mathcal{B}$). So part 1 of the substitution lemma yields $\mathcal{M}_{\mathcal{V}}(\psi(\tau_\sigma), \Theta(\sigma)) = 1$, for each grounded term τ_σ that appears on $\mathcal{B}$. So the branch $\mathcal{B}'$ that results from one or more applications of the universal quantifier rule to $\sigma \;\; \forall\alpha\psi(\alpha)$ is satisfiable in $\mathcal{M}$ relative to $\mathcal{V}$ and Θ.

For the negated universal quantifier case, assume that the prefixed formula $\sigma \;\; \neg\forall\alpha\psi(\alpha)$ is in S, that clauses 1–5 hold, and hence that $\mathcal{M}_{\mathcal{V}}(\neg\forall\alpha\psi(\alpha), \Theta(\sigma)) = 1$. So $\mathcal{M}_{\mathcal{V}}(\forall\alpha\psi(\alpha), \Theta(\sigma)) = 0$. By clause 10 of Section 3.4, $\mathcal{M}_{\mathcal{U}}(\psi(\alpha), \Theta(\sigma)) = 0$, for some valuation $\mathcal{U}$ that is an α-variant-at-$\Theta(\sigma)$ of $\mathcal{V}$. Since each node on a tableau branch is just finitely many steps from the origin, for any given node only finitely many distinct parameters appear in the nodes up to and including it. So a parameter π_σ that does not appear on $\mathcal{B}$ will always be available for application of the negated universal quantifier rule. Let $\mathcal{U}'$ be the π_σ-variant-at-$\Theta(\sigma)$ of $\mathcal{U}$ such that $\mathcal{U}'(\pi_\sigma) = \mathcal{U}(\alpha)$. So by part 1 of the substitution lemma, $\mathcal{M}_{\mathcal{U}'}(\psi(\pi_\sigma), \Theta(\sigma)) = 0$. Hence $\mathcal{M}_{\mathcal{U}'}(\neg\psi(\pi_\sigma), \Theta(\sigma)) = 1$. Then $\mathcal{U}'$ satisfies clauses 1, 2, and 4, since $\mathcal{V}$ does. And $\mathcal{U}'$ satisfies clause 3 because it makes all the same assignments to parameters other than π_σ as $\mathcal{V}$ does, and because $\mathcal{U}'(\pi_\sigma) \in \mathcal{D}(\Theta(\sigma))$. So the branch $\mathcal{B}'$ that results from application of the negated universal quantifier rule to $\sigma \;\; \neg\forall\alpha\psi(\alpha)$, and thereby adds the node $\sigma \;\; \neg\psi(\pi_\sigma)$ to $\mathcal{B}$, is satisfiable in $\mathcal{M}$ relative to $\mathcal{U}'$ and Θ.

For the predicate abstraction cases, recall that a predicate abstract $\langle\lambda\alpha.\psi(\alpha)\rangle$ is satisfied by a term at a world if and only if the term denotes an object in the domain of that world and the object satisfies $\psi(\alpha)$. Form 1 of the rule for unnegated predicate abstraction applies to prefixed formulas of the form $\sigma \;\; \langle\lambda\alpha.\psi(\alpha)\rangle(\iota)$, where ι is an individual constant. In this case the predicate abstract is instantiated with the grounded name ι_σ, which may or may not already appear somewhere on $\mathcal{B}$. If ι_σ does already appear on $\mathcal{B}$, it was introduced either by form 1 of the atomic formula rule or by a previous application of the very rule under consideration. So suppose $\sigma \;\; \langle\lambda\alpha.\psi(\alpha)\rangle(\iota)$ is in S, that clauses 1–5 hold, and hence that $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\iota), \Theta(\sigma)) = 1$.

There are two subcases as follows.

(a) If ι_σ already appears on $\mathcal{B}$, $\mathcal{I}(\iota, \Theta(\sigma))$ is defined, $\mathcal{I}(\iota, \Theta(\sigma)) \in \mathcal{D}(\Theta(\sigma))$, and $\mathcal{V}(\iota_\sigma) = \mathcal{I}(\iota, \Theta(\sigma))$ (by clause 4). Since $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\iota), \Theta(\sigma)) = 1$, there is an α-variant-at-$\Theta(\sigma)$ of $\mathcal{V}$, $\mathcal{U}$, such that $\mathcal{U}(\alpha) = (\mathcal{V} \star \mathcal{I})(\iota, \Theta(\sigma))$ and $\mathcal{M}_{\mathcal{U}}(\psi(\alpha), \Theta(\sigma)) = 1$ (clause 12a of Section 3.4). But since ι is an individual constant, $(\mathcal{V} \star \mathcal{I})(\iota, \Theta(\sigma)) = \mathcal{I}(\iota, \Theta(\sigma))$, and so by the identities already established $\mathcal{U}(\alpha) = \mathcal{V}(\iota_\sigma)$. So by part 2 of the substitution lemma, $\mathcal{M}_{\mathcal{V}}(\psi(\iota_\sigma), \Theta(\sigma)) = 1$.

(b) If ι_σ does not appear on $\mathcal{B}$, then, since $\mathcal{M}_\mathcal{V}(\langle\lambda\alpha.\psi(\alpha)\rangle(\iota), \Theta(\sigma)) = 1$, $\mathcal{I}(\iota, \Theta(\sigma))$ is nonetheless defined and $\mathcal{I}(\iota, \Theta(\sigma)) \in \mathcal{D}(\Theta(\sigma))$ (clause 12a of Section 3.4). Furthermore $\mathcal{M}_\mathcal{U}(\psi(\alpha), \Theta(\sigma)) = 1$, where $\mathcal{U}$ is the α-variant-at-$\Theta(\sigma)$ of $\mathcal{V}$ such that $\mathcal{U}(\alpha) = (\mathcal{V} \star \mathcal{I})(\iota, \Theta(\sigma)) = \mathcal{I}(\iota, \Theta(\sigma))$ (by clause 12a of Section 3.4 and the definition of $(\mathcal{V} \star \mathcal{I})$ in Section 3.3). But there is no guarantee that $\mathcal{V}(\iota_\sigma) = \mathcal{I}(\iota, \Theta(\sigma))$. So let $\mathcal{V}'$ be the ι_σ-variant-at-$\Theta(\sigma)$ of $\mathcal{V}$ such that $\mathcal{V}'(\iota_\sigma) = \mathcal{I}(\iota, \Theta(\sigma))$. Clearly $\mathcal{M}$, $\mathcal{V}'$, and Θ satisfy clauses 1–3 of the definition given at the beginning of the present section, and when the node $\sigma\ \psi(\iota_\sigma)$ is added to $\mathcal{B}$ by application of form 1 of the unnegated predicate abstraction rule, clause 4 is also satisfied. Let $\mathcal{U}'$ be the α-variant-at-$\Theta(\sigma)$ of $\mathcal{V}'$ such that $\mathcal{U}'(\alpha) = \mathcal{V}'(\iota_\sigma)$. So by part 1 of the substitution lemma (with $\mathcal{V}'$ as $\mathcal{V}$, ι_σ as τ_σ, $\Theta(\sigma)$ as w, and $\mathcal{U}'$ as $\mathcal{U}$), $\mathcal{M}_{\mathcal{U}'}(\psi(\alpha), \Theta(\sigma)) = \mathcal{M}_{\mathcal{V}'}(\psi(\iota_\sigma), \Theta(\sigma))$. Since $\mathcal{M}_\mathcal{U}(\psi(\alpha), \Theta(\sigma)) = 1$ and $\mathcal{U}'$ differs from $\mathcal{U}$ only in what it assigns to ι_σ, which does not appear in $\psi(\alpha)$, $\mathcal{M}_{\mathcal{U}'}(\psi(\alpha), \Theta(\sigma)) = 1$. Hence $\mathcal{M}_{\mathcal{V}'}(\psi(\iota_\sigma), \Theta(\sigma)) = 1$.

So the branch $\mathcal{B}'$ that results from application of form 1 of the rule for unnegated predicate abstraction to $\sigma\ \langle\lambda\alpha.\psi(\alpha)\rangle(\iota)$, and thereby adds $\sigma\ \psi(\iota_\sigma)$ to $\mathcal{B}$, is satisfiable in $\mathcal{M}$ relative to $\mathcal{V}$ and Θ, or satisfiable in $\mathcal{M}$ relative to $\mathcal{V}'$ and Θ, depending on whether ι_σ already appears on $\mathcal{B}$.

Having given detailed proofs for the quantifier rules and one of the predicate abstraction rules, I will just sketch the proofs for the remaining cases.

The proof for form 2 of the predicate abstraction rule is similar to that for form 1, but simpler, since the grounded term τ_σ that is instantiated appears in the premise. The proof for form 3 combines features of the proofs for form 2 and the negated quantifier rule.

The proofs for the first two forms of the negated predicate abstraction rule are straightforward in view of the explanation given in Section 8.2.4, where they are introduced, and the proofs for predicate abstraction. Form 3 does not require a separate proof since it can be treated as a derived rule.

The proofs for the two identity rules are straightforward.

Forms 1 and 2 of the atomic formula rule are similar, respectively, to forms 1 and 3 of the predicate abstraction rule, and their proofs are similar. The proof for form 3 of the atomic formula rule is trivial. ■

Theorem 6 *Tableau Soundness for $\mathfrak{L}$*

If ϕ is a sentence of $\mathfrak{L}$ and Γ is a set of sentences of $\mathfrak{L}$, then if ϕ is derivable from Γ, ϕ is a consequence of Γ. In abbreviated form, if $\Gamma \vdash \phi$, then $\Gamma \vDash \phi$.

Proof Suppose that ϕ is not a consequence of Γ. Then $\Gamma \cup \{\neg\phi\}$ is satisfiable. Hence there is a model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I}\rangle$, and a world $w \in \mathcal{W}$, such that for any valuation $\mathcal{V}$ relative to $\mathcal{M}$ and each sentence $\psi \in \Gamma \cup \{\neg\phi\}$, $\mathcal{M}_\mathcal{V}(\psi, w) = 1$. So for any finite subset Γ' of Γ and each $\psi \in \Gamma' \cup \{\neg\phi\}$, $\mathcal{M}_\mathcal{V}(\psi, w) = 1$. For any such Γ' consider the set of ordered pairs $S = \{\langle 1, \psi\rangle \mid \psi \in \Gamma'\} \cup \{\langle 1, \neg\phi\rangle\}$. S is a set of prefixed formulas of $\mathfrak{L}$ and hence of $\mathfrak{L}^\star$. If we define the function Θ so that $\Theta(1) = w$, clauses 1–5 at the beginning of this section are satisfied. Thus S is a set of prefixed formulas of $\mathfrak{L}^\star$ that is satisfiable in $\mathcal{M}$ relative to $\mathcal{V}$. Since S is finite, its members constitute the initial sentences of a tableau. Call this tableau $S^\mathcal{T}$. By definition, $S^\mathcal{T}$ is a satisfiable tableau.

By the soundness lemma, any application of a tableau rule to $S^{\mathcal{T}}$ results in another satisfiable tableau. This means that in every such extension of $S^{\mathcal{T}}$ there is a branch $\mathcal{B}$, a model $\mathcal{M}$, and a valuation $\mathcal{V}$ such that, for every node $\sigma\ \phi$ appearing on $\mathcal{B}$, $\mathcal{M}_{\mathcal{V}}(\phi, \Theta(\sigma)) = 1$. But by the definition of satisfaction, no formula and its negation can both be assigned the value 1. Hence $\mathcal{B}$ is not closed, and thus ϕ is not derivable from Γ. ■

A.2 Completeness Although the details of the completeness proof are complex, the basic idea behind it is simple. If a tableau has an open branch in which all the rules that can be applied have been applied, that branch contains information sufficient to determine an interpretation and a valuation under which all the formulas of $\mathfrak{L}^{\star}$ on the branch are satisfied. The worlds of this interpretation are the prefixes of nodes appearing on the branch, and the individuals are the equivalence classes of grounded terms determined by the identity statements appearing on the branch. The atomic formulas and grounded names appearing on the branch determine the assignment of extensions to predicates and individuals to individual constants, respectively. The valuation assigns to a grounded term the equivalence class of which that term is a member, and it makes an arbitrary assignment of these equivalence classes to variables.

As is common in tableau proof systems, some tableaus will contain infinitely long branches. This is because some rules may have to be applied repeatedly to the same node (or pair of nodes) of a tableau. The rules for necessity, universal quantifier, and substitutivity of identity fall into this category. I will call a branch *complete* if all applicable rules have been applied to all its nodes.

Consider a complete and open branch $\mathcal{B}$ of a tableau $\mathcal{T}$, and define a *categorical model* $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$ and a *categorical valuation* $\mathcal{V}$ (relative to $\mathcal{M}$) *induced by* $\mathcal{B}$ as follows:

$\mathcal{W}$: $\mathcal{W}$ is the set of prefixes that appear on $\mathcal{B}$

$@$: $@$ is $@$, if $@$ appears as the prefix of a node on $\mathcal{B}$; otherwise $@$ is 1

$\mathcal{R}$: $\mathcal{R} = \{\langle \sigma, \sigma \rangle \mid \sigma \text{ appears on } \mathcal{B}\} \cup \{\langle \sigma, \sigma.n \rangle \mid \sigma \text{ and } \sigma.n \text{ both appear on } \mathcal{B}\}$.

In order to define $\mathcal{D}$, the function that assigns sets of individuals to the worlds in $\mathcal{W}$, some further notation is needed. As before, let $\sigma_1, \sigma_2, \sigma_3, \ldots$ be prefixes of nodes that appear on $\mathcal{B}$, and let $\tau_{\sigma_1}, \tau_{\sigma_2}, \tau_{\sigma_3}, \ldots$ and $\upsilon_{\sigma_1}, \upsilon_{\sigma_2}, \upsilon_{\sigma_3}, \ldots$ be grounded terms. Also, let $\mathcal{G}$ be the set of grounded terms that appear on $\mathcal{B}$, and let $\mathcal{E}$ be the set of formulas of $\mathfrak{L}^{\star}$ of the form $(\tau_{\sigma_2} = \upsilon_{\sigma_3})$ such that, for some σ_1, $\sigma_1\ (\tau_{\sigma_2} = \upsilon_{\sigma_3})$ is a node on $\mathcal{B}$. Since $\mathcal{B}$ is a complete branch, and in view of the tableau rules for identity in Section 8.2.5, $\mathcal{E}$ defines an equivalence relation on $\mathcal{G}$. Let $\mathcal{Y}$ be the set of equivalence classes (of members of $\mathcal{G}$) determined by $\mathcal{E}$. If $\tau \in \mathcal{G}$, let $\overline{\tau}$ denote the member of $\mathcal{Y}$ of which τ is a member. (Thus for any tableau prefixes σ_2 and σ_3, and grounded terms τ_{σ_2} and υ_{σ_3}, $\overline{\tau_{\sigma_2}} = \overline{\upsilon_{\sigma_3}}$ if and only if there is a prefix σ_1 such that $\sigma_1\ (\tau_{\sigma_2} = \upsilon_{\sigma_3})$ is a node on $\mathcal{B}$. For example, if $1.2\ (a_{1.2.@} = p_{1.1.2})$ is a node on $\mathcal{B}$, then $\overline{a_{1.2.@}} = \overline{p_{1.1.2}}$. That is, the grounded name $a_{1.2.@}$ and the parameter $p_{1.1.2}$ denote the same individual, namely, the equivalence class of grounded terms of which both are members.)

$\mathcal{D}$: $\mathcal{D}(\sigma_1) = \{y \mid y \in \mathcal{Y}$ and at least one member of y is a grounded term with subscript $\sigma_1\}$.

So the domain of world σ_1 consists of all and only those members of $\mathcal{Y}$ that contain at least one grounded term with subscript σ_1. Since the domain of a model is the union of the domains of its worlds (see Section 3.1), $\mathcal{D}_{\mathcal{M}} = \mathcal{Y}$. That is, the domain of the model is the set of equivalence classes of grounded terms defined on $\mathcal{G}$ by $\mathcal{E}$. Thus for every prefix $\sigma_1 \in \mathcal{W}$, $\mathcal{D}(\sigma_1) \subseteq \mathcal{D}_{\mathcal{M}}$.

The interpretation $\mathcal{I}$ of model $\mathcal{M}$ makes assignments to the individual constants and predicates of $\mathfrak{L}$.[83] Let ι be an individual constant of $\mathfrak{L}$, let σ_1 be a prefix of a node that appears on $\mathcal{B}$, and let ι_{σ_1} be the grounded name (and hence grounded term) that results from subscripting ι with σ_1.

$\mathcal{I}$: If ι_{σ_1} appears as part of a formula on $\mathcal{B}$, $\mathcal{I}(\iota, \sigma_1) = \overline{\iota_{\sigma_1}}$. Otherwise $\mathcal{I}(\iota, \sigma_1)$ is undefined.

Since $\overline{\iota_{\sigma_1}} \in \mathcal{D}(\sigma_1)$, $\mathcal{I}$ satisfies the constraint imposed by clause 1 of Section 3.2, namely, that if an individual constant of $\mathfrak{L}$ designates at a world, its designation is a member of the domain of that world.

Next, let θ be an n-ary predicate (perhaps $=$), σ_1 a prefix that appears on $\mathcal{B}$ (perhaps @), τ_{σ_1} a grounded term, $\overline{\tau_{\sigma_1}}$ the equivalence class of grounded terms such that $\tau_{\sigma_1} \in \overline{\tau_{\sigma_1}}$, $\langle \ldots, \overline{\tau_{\sigma_1}}, \ldots \rangle$ an n-tuple of equivalence classes of grounded terms of which $\overline{\tau_{\sigma_1}}$ is the m^{th} member, and $\theta(\ldots \tau_{\sigma_1} \ldots)$ an atomic formula of $\mathfrak{L}^\star$ in which all the individual symbols are grounded terms with subscript σ_1 and of which τ_{σ_1} is the m^{th}.

$\mathcal{I}$: $\mathcal{I}(\theta, \sigma_1) = \{\langle \ldots, \overline{\tau_{\sigma_1}}, \ldots \rangle \mid \sigma_1\ \theta(\ldots \tau_{\sigma_1} \ldots)$ is a node on $\mathcal{B}\}$.

If θ is the identity predicate, this becomes

$\mathcal{I}$: $\mathcal{I}(=, \sigma_1) = \{\langle \overline{\tau_{\sigma_1}}, \overline{\upsilon_{\sigma_1}} \rangle \mid \sigma_1\ (\tau_{\sigma_1} = \upsilon_{\sigma_1})$ is a node on $\mathcal{B}\}$.

In view of the definitions of set $\mathcal{Y}$ and function $\mathcal{D}$, for each $\sigma_1 \in \mathcal{W}$, $\mathcal{I}(\theta, \sigma_1)$ assigns an n-ary relation on $\mathcal{D}(\sigma_1)$ to the predicate θ. So $\mathcal{I}$ satisfies the first part of clause 2 of Section 3.2, namely, the requirement that the extension of a predicate at a world contain only objects from the domain of that world. Furthermore, if θ is $=$, $\mathcal{I}(=, \sigma_1)$ is the set of ordered pairs $\langle \overline{\tau_{\sigma_1}}, \overline{\tau_{\sigma_1}} \rangle$, where $\overline{\tau_{\sigma_1}}$ is an equivalence class of grounded terms, and $\overline{\tau_{\sigma_1}} \in \mathcal{D}(\sigma_1)$. That is, for each $\sigma_1 \in \mathcal{W}$, $\mathcal{I}(=, \sigma_1)$ is the identity relation on $\mathcal{D}(\sigma_1)$. So $\mathcal{I}$ also satisfies the second part of clause 2 of Section 3.2. This completes the definition of a categorical model $\mathcal{M}$ induced by a complete and open branch $\mathcal{B}$.

Given such a model $\mathcal{M}$, the categorical valuation $\mathcal{V}$ relative to $\mathcal{M}$ is easily defined. Recall (see Section 8.1.1) that grounded terms of $\mathfrak{L}^\star$ (i.e., parameters and grounded names) are treated semantically like variables and thus have their values assigned by $\mathcal{V}$. Where $\mathcal{G}$, as before, is the set of grounded terms that appear on $\mathcal{B}$, and $\mathcal{A}$ is the set of variables of $\mathfrak{L}$, $\mathcal{V}$ is a function from $\mathcal{G} \cup \mathcal{A}$ into $\mathcal{D}_{\mathcal{M}}$ such that

$\mathcal{V}$: For every grounded term $\tau_{\sigma_1} \in \mathcal{G}$, $\mathcal{V}(\tau_{\sigma_1}) = \overline{\tau_{\sigma_1}}$,

$\mathcal{V}$: For every variable $\alpha \in \mathcal{A}$, $\mathcal{V}(\alpha)$ is some arbitrarily selected member of $\mathcal{D}_{\mathcal{M}}$.

This completes the definition of a categorical valuation $\mathcal{V}$ induced by a complete and open branch $\mathcal{B}$ (relative to a categorical model $\mathcal{M}$ induced by this same branch $\mathcal{B}$).[84]

Completeness lemma: A categorical model and valuation satisfy all the formulas on the branch that induces them.

Let $\mathcal{B}$ be a complete and open branch of a tableau $\mathcal{T}$, and let $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$ be a categorical model and $\mathcal{V}$ a categorical valuation (relative to $\mathcal{M}$) induced by $\mathcal{B}$. For any formula ϕ of $\mathfrak{L}^\star$, and any prefix σ:

Part 1. If $\sigma\ \phi$ is a node on $\mathcal{B}$, then $\mathcal{M}_{\mathcal{V}}(\phi, \sigma) = 1$;

Part 2. If $\sigma\ \neg\phi$ is a node on $\mathcal{B}$, then $\mathcal{M}_{\mathcal{V}}(\phi, \sigma) = 0$.

Proof is by induction on the number of connectives and quantifiers in ϕ.

Basis: Part 1.

Case a. Suppose that ϕ is an atomic formula made up of the n-ary predicate θ and n grounded terms, τ_σ, all of which have σ as their grounding subscript. Represent ϕ as $\theta(\ldots\tau_\sigma\ldots)$. By hypothesis $\mathcal{B}$ is complete. So, by the definition of $\mathcal{I}$, the extension of θ at σ contains $\langle\ldots,\overline{\tau_\sigma},\ldots\rangle$ if and only if $\sigma\ \theta(\ldots\tau_\sigma\ldots)$ is a node on $\mathcal{B}$. And by the definition of $\mathcal{V}$, $\mathcal{V}(\tau_\sigma) = \overline{\tau_\sigma}$. So if $\sigma\ \phi$ is a node on $\mathcal{B}$, $\mathcal{M}_{\mathcal{V}}(\phi, \sigma) = 1$.

Case b. Suppose that $\sigma\ \phi$ is a node on $\mathcal{B}$, where ϕ is an atomic formula made up of the n-ary predicate θ and n terms, and where these terms may include individual constants of $\mathfrak{L}$ as well as grounded terms (parameters and grounded names). Individual constants of $\mathfrak{L}$ are ungrounded (and hence have no subscripts), and parameters and grounded names may be subscripted with tableau prefixes other than σ, the prefix of the node $\sigma\ \phi$ with which we are concerned. By hypothesis $\mathcal{B}$ is complete, so all possible applications of all rules will have been made to all nodes on $\mathcal{B}$. So if the formula ϕ of node $\sigma\ \phi$ contains individual constants, repeated application of form 1 of the atomic formula rule will have added a node $\sigma\ \phi^*$, where ϕ^* is like ϕ except that each occurrence of an individual constant ι has been replaced by the corresponding grounded name ι_σ. So ι_σ will appear as part of a formula on $\mathcal{B}$, and thus by the definition of $\mathcal{I}$, $\mathcal{I}(\iota, \sigma) = \overline{\iota_\sigma}$.

Similarly, if $\phi*$ contains (one or more occurrences of) a grounded term τ with a subscript other than σ, application of form 2 of the atomic formula rule will have added a node of the form $\sigma\ (\pi_\sigma = \tau)$, where π_σ is a parameter that has not previously appeared on $\mathcal{B}$. So by the definitions of $\mathcal{Y}$ and $\mathcal{V}$, $\overline{\pi_\sigma} = \overline{\tau}$, and thus $\mathcal{V}(\tau) = \overline{\pi_\sigma}$. And application of the substitutivity of identity rule will have added a node $\sigma\ \phi*^{\pi_\sigma}$, where $\phi*^{\pi_\sigma}$ is like $\phi*$ except that τ has been replaced by π_σ throughout $\phi*$. Since $\mathcal{B}$ is complete, this will have happened for each term appearing in $\sigma\ \phi*$ that is grounded with a subscript other than σ, and occurrences of distinct terms of this kind will have been replaced by distinct parameters, each of which was new to the branch when it appeared. Call the resulting formula $\phi**$. So the node $\sigma\ \phi**$ will appear on $\mathcal{B}$. Each term that appears in $\phi**$ is subscripted with σ, and so as in Case (a) above $\mathcal{M}_{\mathcal{V}}(\phi**, \sigma) = 1$. And since $\mathcal{V}(\pi_\sigma) = \mathcal{V}(\tau)$, for each parameter π_σ that was substituted for a grounded term τ in the creation of $\phi**$ from $\phi*$, $\phi**$ will contain, at each of its individual-term positions, a term that denotes the same member of $\mathcal{D}(\sigma)$ as is denoted by the term in the corresponding position in $\phi*$. So $\mathcal{M}_{\mathcal{V}}(\phi*, \sigma) = \mathcal{M}_{\mathcal{V}}(\phi**, \sigma)$, and hence $\mathcal{M}_{\mathcal{V}}(\phi*, \sigma) = 1$.

Finally, recall that ϕ^* was defined as differing from ϕ only in containing the grounded name ι_σ at places where ϕ contains the corresponding individual constant ι. By the definition of $\mathcal{I}$, if ι is an individual constant of $\mathfrak{L}$ and the grounded name ι_σ appears as part of a formula on $\mathcal{B}$, then $\mathcal{I}(\iota, \sigma) = \overline{\iota_\sigma}$. Thus since $\mathcal{M}_{\mathcal{V}}(\phi^*, \sigma) = 1$, it follows that $\mathcal{M}_{\mathcal{V}}(\phi, \sigma) = 1$. So if $\sigma\ \phi$ is a node on $\mathcal{B}$, $\mathcal{M}_{\mathcal{V}}(\phi, \sigma) = 1$.

Basis: Part 2.

Case a. As in Case (a) of Part 1, let ϕ be an atomic formula made up of the n-ary predicate θ and n grounded terms τ_σ, all of which have σ as their grounding subscript. Represent ϕ as $\theta(\ldots\tau_\sigma\ldots)$. Suppose $\sigma\ \neg\theta(\ldots\tau_\sigma\ldots)$ is a node on $\mathcal{B}$. By hypothesis $\mathcal{B}$ is open, so $\sigma\ \theta(\ldots\tau_\sigma\ldots)$ is a not a node on $\mathcal{B}$. Suppose for *reductio* that $\mathcal{M}_{\mathcal{V}}(\theta(\ldots\tau_\sigma\ldots),\sigma) = 1$. Since $\theta(\ldots\tau_\sigma\ldots)$ is an atomic formula of $\mathfrak{L}^\star$, by clause 2 of Section 3.2 each term τ_σ in ϕ denotes a member of $\mathcal{D}(\sigma)$. And by the definition of $\mathcal{I}$ of the categorical model given earlier in this section, a node of the form $\sigma\ \theta(\ldots\tau_\sigma\ldots)$ appears on $\mathcal{B}$. $\mathcal{B}$ is thus closed, which contradicts the assumption of the completeness lemma. So by *reductio* $\mathcal{M}_{\mathcal{V}}(\theta(\ldots\tau_\sigma\ldots),\sigma) = 0$.

Case b. Let ϕ be as in Case (b) of Part 1. That is, ϕ is an atomic formula made up of the n-ary predicate θ and n terms, where these terms may include individual constants of $\mathfrak{L}$ as well as parameters and grounded names that are subscripted with σ or with tableau prefixes other than σ. Suppose $\sigma\ \neg\phi$ is a node on $\mathcal{B}$. By hypothesis $\mathcal{B}$ is open, so $\sigma\ \phi$ is a not a node on $\mathcal{B}$. Again suppose for *reductio* that $\mathcal{M}_{\mathcal{V}}(\phi,\sigma) = 1$. Since ϕ is an atomic formula of $\mathfrak{L}^\star$, by Sections 3.2 and 3.4 each term in ϕ denotes a member of $\mathcal{D}(\sigma)$, even though those terms may be unsubscripted or subscripted with prefixes other than σ.[85] And by the definition of $\mathcal{I}$ of the categorical model $\mathcal{M}$ given earlier in this section, a node of the form $\sigma\ \theta(\ldots\tau_\sigma\ldots)$ appears on $\mathcal{B}$, where each term τ_σ in $\theta(\ldots\tau_\sigma\ldots)$ is grounded with σ and denotes the same member of $\mathcal{D}(\sigma)$ as the term at the corresponding position in ϕ (even though the latter terms may be unsubscripted or subscripted with a prefix other than σ). It can then be shown that $\sigma\ \phi$ is a node on $\mathcal{B}$, which contradicts the assumption that $\mathcal{B}$ is open. So by *reductio* $\mathcal{M}_{\mathcal{V}}(\phi,\sigma) = 0$.

The proof that $\sigma\ \phi$ is a node on $\mathcal{B}$ is straightforward but complex. Rather than giving a fully general proof, I give an example that should be sufficient to convince the reader of its truth. Let θ be the five-place predicate F, and let ϕ be the atomic formula $Fa_1\, b_{1.2}\, p_{1.1}\, b\, c$. Also, let σ be the tableau prefix 1.1. Then by our assumptions $1.1\ \neg Fa_1\, b_{1.2}\, p_{1.1}\, b\, c$ is a node on $\mathcal{B}$, and (for *reductio*) $\mathcal{M}_{\mathcal{V}}(Fa_1\, b_{1.2}\, p_{1.1}\, b\, c,\ 1.1) = 1$. Since the individual constants b and c appear in an atomic formula that has the value 1 at world 1.1, by the semantics of $\mathfrak{L}$ (see Sections 3.2 and 3.4 above) $\mathcal{I}(b, 1.1)$ and $\mathcal{I}(c, 1.1)$ are defined and each is a member of $\mathcal{D}(1.1)$. And since the members of $\mathcal{D}(1.1)$ are equivalence classes of grounded terms that appear on $\mathcal{B}$, $b_{1.1}$ and $c_{1.1}$ must both appear somewhere on $\mathcal{B}$, by the definition of interpretation $\mathcal{I}$ of the categorical model $\mathcal{M}$ given earlier in this section. But then by this same definition, $\mathcal{I}(b, 1.1) = \overline{b_{1.1}}$ and $\mathcal{I}(c, 1.1) = \overline{c_{1.1}}$. Similarly, by the definition of the valuation $\mathcal{V}$ associated with the categorical model $\mathcal{M}$ defined earlier in this section, $\mathcal{V}(a_1) = \overline{a_1}$, $\mathcal{V}(b_{1.2}) = \overline{b_{1.2}}$, and $\mathcal{V}(p_{1.1}) = \overline{p_{1.1}}$. So by the semantics of $\mathfrak{L}$ (see Sections 3.2, 3.4 above)

(1) $\langle\overline{\mathrm{a}_1}, \overline{\mathrm{b}_{1.2}}, \overline{\mathrm{p}_{1.1}}, \overline{\mathrm{b}_{1.1}}, \overline{\mathrm{c}_{1.1}}\rangle \in \mathcal{I}(F,\ 1.1)$.

But then there must be grounded terms subscripted with 1.1 that are members of the equivalence classes $\overline{a_1}$ and $\overline{b_{1.2}}$. For the semantics of $\mathfrak{L}$ (see Sections 3.2 and 3.4 above) states that a predicate is satisfied by an n-tuple of objects at a world, only if each of those objects is a member of the domain of that world. Suppose these grounded terms are the parameter $q_{1.1}$ and the grounded name $d_{1.1}$, respectively. Then $q_{1.1}$ and $d_{1.1}$ will appear on $\mathcal{B}$ in nodes of the form $\sigma_1\ a_1 = q_{1.1}$ and $\sigma_2\ b_{1.2} = d_{1.1}$, where σ_1 and σ_2 may be any prefixes. Thus

(2) $\overline{a_1} = \overline{q_{1.1}}$ and $\overline{b_{1.2}} = \overline{d_{1.1}}$.

Substituting $\overline{q_{1.1}}$ for $\overline{a_1}$ and $\overline{d_{1.1}}$ for $\overline{b_{1.2}}$ in 1 yields

(3) $\langle\overline{\mathrm{q}_{1.1}}, \overline{\mathrm{d}_{1.1}}, \overline{\mathrm{p}_{1.1}}, \overline{\mathrm{b}_{1.1}}, \overline{\mathrm{c}_{1.1}}\rangle \in \mathcal{I}(F,\ 1.1)$.

But then by the definition of $\mathcal{I}$ of the categorical model $\mathcal{M}$ given earlier in this section

(4) 1.1 $Fq_{1.1}\, d_{1.1}\, p_{1.1}\, b_{1.1}\, c_{1.1}$ is a node on $\mathcal{B}$.

From 4 and the fact that $\sigma_1\ a_1 = q_{1.1}$ and $\sigma_2\ b_{1.2} = d_{1.1}$ appear on $\mathcal{B}$, the tableau rule for substitutivity of identity yields

(5) 1.1 $Fa_1\, b_{1.2}\, p_{1.1}\, b_{1.1}\, c_{1.1}$ is a node on $\mathcal{B}$.

Finally, by two applications of form 3 of the tableau rule for atomic formulas,

(6) 1.1 $Fa_1\, b_{1.2}\, p_{1.1}\, b\, c$ is a node on $\mathcal{B}$.

Hence $\mathcal{B}$ is closed, which contradicts the assumption that it is open. So by *reductio* $\mathcal{M}_{\mathcal{V}}(Fa_1\, b_{1.2}\, p_{1.1}\, b\, c,\ 1.1) = 0$. This completes the proof for Case (b) when ϕ is $Fa_1\, b_{1.2}\, p_{1.1}\, b\, c$, σ is 1.1, and $\sigma\ \neg\phi$ is a node on $\mathcal{B}$.

It should be clear that an analogous proof can be given for any atomic formula made up of an n-ary predicate and n terms, where these terms may include individual constants as well as parameters and grounded names that are subscripted with σ or with tableau prefixes other than σ. Proof of the basis of the induction is thus complete.

Inductive step Assume the result holds for formulas with fewer than n logical operators, and show that it holds for formulas with n operators. Cases for the connectives and the universal quantifier are straightforward and are left to the reader. (The universal quantifier case is facilitated by use of part 1 of the substitution lemma.)

Predicate abstraction: ϕ is $\langle\lambda\alpha.\psi(\alpha)\rangle$. So we must show the following.

Where ι is an individual constant:
(1) if $\sigma\ \langle\lambda\alpha.\psi(\alpha)\rangle(\iota)$ is a node on $\mathcal{B}$, then $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\iota), \sigma) = 1$;
(2) if $\sigma\ \neg\langle\lambda\alpha.\psi(\alpha)\rangle(\iota)$ is a node on $\mathcal{B}$, then $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\iota), \sigma) = 0$.
Where τ_σ is a grounded term:
(3) if $\sigma\ \langle\lambda\alpha.\psi(\alpha)\rangle(\tau_\sigma)$ is a node on $\mathcal{B}$, then $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_\sigma), \sigma) = 1$;
(4) if $\sigma\ \neg\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_\sigma)$ is a node on $\mathcal{B}$, then $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_\sigma), \sigma) = 0$.
Where τ_{σ_1} is a grounded term, and $\sigma_1 \neq \sigma$:
(5) if $\sigma\ \langle\lambda\alpha.\psi(\alpha)\rangle(\tau_{\sigma_1})$ is a node on $\mathcal{B}$, then $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_{\sigma_1}), \sigma) = 1$;
(6) if $\sigma\ \neg\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_{\sigma_1})$ is a node on $\mathcal{B}$, then $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_{\sigma_1}), \sigma) = 0$.

The case where τ_σ is a grounded term (involving 3 and 4 above) is the simplest. If $\sigma\ \langle\lambda\alpha.\psi(\alpha)\rangle(\tau_\sigma)$ is a node on $\mathcal{B}$, then, since $\mathcal{B}$ is a complete branch, a node of the form $\sigma\ \psi(\tau_\sigma)$ also appears on $\mathcal{B}$ (from application of form 2 of the predicate abstraction rule). So by the inductive hypothesis $\mathcal{M}_{\mathcal{V}}(\psi(\tau_\sigma), \sigma) = 1$, and by part 1 of the substitution lemma $\mathcal{M}_{\mathcal{U}}(\psi(\alpha), \sigma) = 1$, where $\mathcal{U}$ is the α-variant-at-σ of $\mathcal{V}$ such that $\mathcal{U}(\alpha) = \mathcal{V}(\tau_\sigma)$. But then by the semantics of predicate abstraction (clause 12a of Section 3.4) $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_\sigma), \sigma) = 1$. So 3 is established.

On the other hand if $\sigma\ \neg\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_\sigma)$ is a node on $\mathcal{B}$, then, since $\mathcal{B}$ is complete, a node of the form $\sigma\ \neg\psi(\tau_\sigma)$ also appears on $\mathcal{B}$ (from application of form 2 of the negated predicate abstraction rule). So by the inductive hypothesis $\mathcal{M}_{\mathcal{V}}(\psi(\tau_\sigma), \sigma) = 0$, and by part 1 of the substitution lemma $\mathcal{M}_{\mathcal{U}}(\psi(\alpha), \sigma) = 0$, where $\mathcal{U}$ is the α-variant-at-σ of $\mathcal{V}$ such that $\mathcal{U}(\alpha) = \mathcal{V}(\tau_\sigma)$. But then by the semantics of predicate abstraction (clause 12a of Section 3.4) $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_\sigma), \sigma) = 0$. So 4 is established.

Consider next the case where ι is an individual constant. The proof of 1 is exactly like that of 3 except that it appeals to form 1 of the predicate abstraction rule and part 2 of the substitution lemma. The proof of 2 divides into two cases depending on whether or not ι_σ appears on $\mathcal{B}$. If it does, the proof is like that of 4, except that it uses form 1 of the negated predicate abstraction rule and part 2 of the substitution lemma. If ι_σ does not appear on $\mathcal{B}$, then in particular σ $(\iota_\sigma = \iota_\sigma)$ is not a node on $\mathcal{B}$, and so by definition of the canonical model $\mathcal{M}$, $\mathcal{I}(\iota, \sigma)$ is undefined for all $\sigma \in \mathcal{W}$. So by clause 12b of Section 3.4 $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\iota), \sigma) = 0$.

To prove 5, consider the case where σ $\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_{\sigma_1})$ is a node on $\mathcal{B}$, τ_{σ_1} is a grounded term, but $\sigma_1 \neq \sigma$. By form 3 of the predicate abstraction rule, both σ $\psi(\tau_{\sigma_1})$ and σ $(\pi_\sigma = \tau_{\sigma_1})$ are nodes on $\mathcal{B}$. So by the substitution of identity rule σ $\psi(\pi_{\sigma_1})$ is a node on $\mathcal{B}$, and thus by the inductive hypothesis $\mathcal{M}_{\mathcal{V}}(\psi(\pi_\sigma), \sigma) = 1$. Since σ $(\pi_\sigma = \tau_{\sigma_1})$ is on $\mathcal{B}$, it follows from the basis case of the induction that $\mathcal{M}_{\mathcal{V}}((\pi_\sigma = \tau_{\sigma_1}), \sigma) = 1$, and hence that $\mathcal{M}_{\mathcal{V}}(\psi(\tau_{\sigma_1}), \sigma) = 1$. But then by part 1 of the substitution lemma and clause 12a of Section 3.4 $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_{\sigma_1}), \sigma) = 1$.

Finally, the proof of 6 is similar to that of 2 in that it divides into two cases depending on whether there is a grounded term υ_σ and a prefix σ_* such that $\sigma_*(\upsilon_\sigma = \tau_{\sigma_1})$ is a node on $\mathcal{B}$. If there is, σ $\neg\langle\lambda\alpha.\psi(\alpha)\rangle(\upsilon_\sigma)$ is a node on $\mathcal{B}$, by the substitution of identity rule, since $\mathcal{B}$ is complete. Also if $\sigma_*(\upsilon_\sigma = \tau_{\sigma_1})$ is a node on $\mathcal{B}$, $\overline{\upsilon_\sigma} = \overline{\tau_{\sigma_1}}$, by the definition of the set $\mathcal{Y}$ that is the domain $\mathcal{D}_{\mathcal{M}}$ of the categorical model $\mathcal{M}$.

So by case 4 above $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\upsilon_\sigma), \sigma) = 0$. And since $\overline{\upsilon_\sigma} = \overline{\tau_{\sigma_1}}$, $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_{\sigma_1}), \sigma) = 0$. On the other hand, if there is no grounded term υ_σ and no prefix σ_* such that $\sigma_*(\upsilon_\sigma = \tau_{\sigma_1})$ is a node on $\mathcal{B}$, then for all grounded terms υ_σ, $\overline{\upsilon_\sigma} \neq \overline{\tau_{\sigma_1}}$, by the definition of $\mathcal{Y}$. So by the definition of $\mathcal{D}$, $\mathcal{V}(\tau_{\sigma_1}) \notin \mathcal{D}(\sigma)$, and thus by 12b of Section 3.4, $\mathcal{M}_{\mathcal{V}}(\langle\lambda\alpha.\psi(\alpha)\rangle(\tau_{\sigma_1}), \sigma) = 0$. ■

Weak completeness follows immediately from the completeness lemma.

Theorem 7 *Completeness for $\mathfrak{L}$*

Let ϕ be a sentence of $\mathfrak{L}$ and Γ a finite set of sentences of $\mathfrak{L}$. If ϕ is a consequence of Γ, then ϕ is derivable from Γ. In abbreviated form, if Γ is finite and $\Gamma \vDash \phi$, then $\Gamma \vdash \phi$.

Proof Suppose ϕ is not derivable from Γ. Thus by the definition of derivability (in Section 8.1.2) no tableau beginning with the members of Γ and $\neg\phi$, all prefixed with 1, is closed. Hence every such tableau has a complete and open branch.

Let $\mathcal{B}$ be a complete and open branch of such a tableau. By the completeness lemma, $\mathcal{B}$ induces a categorical model $\mathcal{M} = \langle \mathcal{W}, @, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$ and a categorical valuation $\mathcal{V}$ (relative to $\mathcal{M}$) such that, for any formula ϕ of $\mathfrak{L}^\star$ and any prefix σ:

If σ ϕ is a node on $\mathcal{B}$, then $\mathcal{M}_{\mathcal{V}}(\phi, \sigma) = 1$; and

If σ $\neg\phi$ is a node on $\mathcal{B}$, then $\mathcal{M}_{\mathcal{V}}(\phi, \sigma) = 0$.

Since $\neg\phi$ and each member of Γ is a sentence of $\mathfrak{L}$ that appears on $\mathcal{B}$ with the prefix 1, ϕ is not a consequence of Γ. ■

Notes

1. It is so called and endorsed by Plantinga [28, p. 93], [27, p. 316] and by Stephanou [38]. Burge [1, p. 313] also endorses it. He does not use the term "serious actualism," but he presents the core idea using a nonmodal first-order schema that applies to primitive predicate and relational terms. Of this schema he says, "It expresses a deep and widely held intuition that the truth of simple singular sentences (other than those implicitly embedded in intensional contexts) is contingent on the contained singular terms' having a denotation. The pre-theoretic notion seems to be that true predications at the most basic

level express comments on topics, or attributions of properties or relations to objects: lacking a topic or object, basic predications *cannot* be true.” (Emphasis added.)

2. See the quotation from [1] in the previous footnote. I discuss serious actualism in more detail in Section 6.1.

3. More precisely, I understand actualism to be the claim that everything actually exists, not the necessitation of this claim, and not the claim that necessarily everything exists. Plantinga understands actualism in this last way. My account, unlike his, involves use of an actuality connective. I discuss actualism and argue for my account in Section 6.1.

4. Modal languages with a predicate abstraction operator can be found in Stalnaker and Thomason [35], Thomason and Stalnaker [39], Fitting and Mendelsohn [7], and Fitting [6], but none of the logics discussed by these authors are seriously actualistic. Similarly, the nonmodal logic studied by Lambert and Bencivenga in [21] contains a predicate abstraction operator, but it is not seriously actualistic. In [32] (reprinted with slight modifications and a postscript in [33]) Stalnaker presents the logic of a quantified modal language *cum* abstraction operator that is seriously actualistic. He also gives axioms and rules of inference which he tacitly assumes, but does not prove, are sound and complete with respect to his semantics. He does not consider uniform substitution. In the present article I give a set of tableau proof rules and prove that they are sound and complete with respect to validity in SAT. I also prove a uniform substitution theorem for SAT. Stalnaker’s system is compared with mine in Section 6.2.

5. I discuss the work of several authors on actualism and serious actualism in Section 6 below.

6. [19, p. 86], emphasis in the original.

7. Modal languages augmented with an actuality connective have been widely studied. (See, e.g., Hodes [13], [14], and more recently Gilbert and Mares [10].)

8. I have discussed formulas of propositional logic with a structure similar to that of (7) in [11] and [12].

9. See Hodes [13]–[15].

10. In logicians’ jargon SAT is a negative universally free logic. It is free because (as is now common) primitive predicates may have a null extension and because universal instantiation for individual quantifiers is not valid, universally free because it includes interpretations in which individual quantifiers range over the empty domain, and negative because atomic sentences containing nondenoting singular terms are always false. (For these terms, and a comprehensive account of free logic, see Lambert [20, especially pp. 124–27, 131].)

11. Subscripts can be used to ensure an infinite supply of names, predicates, and variables; superscripts to indicate the arity of predicates. In practice I will have no need for subscripts, and I will let context indicate arity. The lower case letters k through t are reserved for use in tableau proofs. (See Section 6.)

12. Clauses 4 and 5 do not allow vacuous quantification or vacuous predicate abstraction. There is no point in allowing either, and vacuous quantification yields anomalies when empty domains are allowed. For example, although $\forall x Fx \rightarrow Fa$ is false if the domain of quantification is empty, $\forall y(\forall x Fx \rightarrow Fa)$ is true.

13. A variable occurs free in a term τ only if τ *is* a variable. Clause 6 is stated so as to facilitate expansion of the class of terms to include function symbols and description operators.

14. The Fitting and Mendelsohn system to which SAT is most similar is the one they develop in [7, Chapter 11].

15. Thus an interpretation need not assign anything to an individual constant at a world w. In models in which $\mathcal{D}(w)$ is empty, this will be so for all individual constants.

16. If $\mathcal{D}(w)$ is empty, all predicates, including $=$, are assigned the null set at w.

17. The serious actualism embodied in clause 2 is exactly the same as that of Menzel and Stephanou. (See Menzel [24, clause (M2), p. 361] and the definition of a model in Stephanou [36, p. 197] and [37, p. 383].) But the actualism of clause 1 is more liberal than that of either. In his preferred system, G, Menzel limits the denotations of individual constants to objects in the domain of the real-world element of a model (see [24, clause (M1$'$), p. 361]). Stephanou does the same on the previously referenced pages. The actualism of SAT, on the other hand, allows individual constants to denote different objects at different worlds, requiring only that an object denoted at a world be a member of the domain of that world. Thus it allows, but does not require, individual constants to function like definite descriptions, referring to nothing in the actual world but to "something" in some other possible world. Although the previous sentence may sound incompatible with actualism, I argue in Section 7 that the idea behind it can be given an entirely actualistic expression.

18. Clause 1 includes the case where θ is $=$ and ϕ is $(\tau_1 = \tau_2)$.

19. Fine has shown in his "Postscript" to Prior and Fine [29, pp. 142–45] (reprinted in [5, pp. 154–56]) that when a modal logic with world-relative quantifiers, like SAT, is supplemented with Vlach [41] connectives possibilist quantifiers can be defined. (One Vlach connective is like the actuality connective, but more fine-tuned. When applied to a formula within the scope of a modal connective, it shifts evaluation back to the world in which the modal connective was evaluated, not back to the actual world. The second Vlach connective keeps track of worlds in which shifts caused by the first Vlach connective are initiated.) Possibilist quantifiers are rejected by actualists, but introducing them as primitive or defining them using the Vlach connectives at least makes them explicit. Introducing them in one of these ways is at least more straightforward and honest than allowing them to creep into the language disguised as free variables.

20. See [11], [12]. These two kinds of validity were first distinguished by Davies and Humberstone [3] and have received much discussion. For an object language that does not contain an actuality connective (or any other logical operator that makes special reference to a designated world) they coincide.

21. It may be tempting to think that the sort of difference noted in this section, in which the truth value of a sentence depends on whether a negation sign is placed inside or outside a predicate abstract, is due entirely to the fact that the sentences involved contain individual constants. But this is not true. Consider the sentences $\exists x \Diamond \langle \lambda y.\neg Sy\rangle(x)$ and $\exists x \Diamond \neg \langle \lambda y.Sy\rangle(x)$, and further specify the model under consideration so that for every world $w^* \in \mathcal{W}$ such that $w\mathcal{R}w^*$, $D(w^*)$ is disjoint from $D(w)$. The first sentence is false at w and the second true.

22. Individual constants can nevertheless be made to approximate rigid designators using suitable assumptions. (See Section 6.2 for details.)

23. In the sentence schema given in this and subsequent sections $\alpha, \alpha_1, \alpha_2, \ldots,$ are variables, $\psi(\alpha), \psi(\alpha_1), \psi(\alpha_2), \ldots,$ are formulas containing one or more free occurrences of the indicated variable but no free occurrence of any other variable, $\iota, \iota_1, \iota_2, \ldots,$ are individual constants, and $\psi(\iota), \psi(\iota_1), \psi(\iota_2), \ldots,$ the results of substituting $\iota, \iota_1, \iota_2, \ldots,$ for each free occurrence of the indicated variable in $\psi(\alpha), \psi(\alpha_1), \psi(\alpha_2), \ldots.$

24. The label (ND), for necessity of distinctness, is borrowed from Stalnaker [32]. Necessity of distinctness is valid in SAT even though none of the versions of necessity of identity listed here are.

25. This approach to identity, including the logical equivalence of $\mathcal{E}(\iota)$ and $(\iota = \iota)$, is endorsed by Burge [1].

26. See [24, especially pp. 359–64].

27. Kripke notes a closely related distinction in [19, p. 90].

28. An example of such an argument is
(Gibb) $(g = l), \Box(\exists x(x = l) \rightarrow (l = l)) \nvDash \Box(\exists x(x = l) \rightarrow (g = l))$.
I call this argument (Gibb) because it plays a pivotal role in Gibbard's argument (see [9, p. 200 ff]) for contingent identity. Gibbard's reasons for favoring a contingent identity logic are very different from mine. I will not consider his work here except to note that a crucial step in his argument is illuminated by the use of predicate abstraction.

It is important for Gibbard's purposes that (Gibb) not be a valid argument. He explains its invalidity as follows:

> [The conclusion] follows from [the premises] by Leibniz's law, then, only if the context
> $$\Box(\exists x(x = l) \rightarrow (\ldots = l)) \tag{7}$$
> attributes a property. We can block the inference to [the conclusion], then, simply by denying that the context (7) attributes a property.

Gibbard does not make use of predicate abstraction, but by its use we can see more clearly what he means by "attributing a property." The property he denies (7) attributes is expressible in my notation as

$$\big\langle \lambda y.\Box(\exists x(x = l) \rightarrow (y = l))\big\rangle \tag{7*}$$

And in view of (Sub1) the following variant of (Gibb)

$$(g = l), \big\langle \lambda y.\Box(\exists x(x = l) \rightarrow (y = l))\big\rangle(l) \vDash \big\langle \lambda y.\Box(\exists x(x = l) \rightarrow (y = l))\big\rangle(g)$$

is a valid argument. So from my perspective Gibbard's criticism of (Gibb) amounts to denying that (7*) "attributes a property."

29. Indeed similar reasoning shows that $\exists\alpha_1\Box\langle\lambda\alpha_2.\phi(\alpha_2)\rangle(\alpha_1) \rightarrow \Box\exists\alpha_1\langle\lambda\alpha_2.\phi(\alpha_2)\rangle(\alpha_1)$ is valid.

30. Kripke [19] gives a counterexample to the $\Box\forall$-form of the Converse Barcan Formula. It assumes that *false* predications using monadic predicates can be made of objects at worlds in which those objects do not exist. And the semantics of [19] also allows the kind of counterexample given above to (CBF), an instance of the $\exists\Diamond$-form of the Converse Barcan Formula, in which *true* predications using monadic predicates are made of objects at worlds in which those objects do not exist. Yet in this same paper Kripke expresses misgivings about this latter allowance, as the passage quoted in Section 1 above shows.

31. Consider a model with just two worlds, w and w', in which $\mathcal{R}$ is reflexive, $w\mathcal{R}w'$, $\mathcal{D}(w) = \{0\}$, and $\mathcal{D}(w') = \{1\}$. If the predicate F is assigned $\{0\}$ at w and $\{1\}$ at w', then at $w\Box\forall x\langle\lambda y.Fy\rangle(x)$ is true but $\forall x\Box\langle\lambda y.Fy\rangle(x)$ is false.

32. The model given in the previous footnote, modified only by changing the extension of F at w from $\{0\}$ to the null set, makes $\Diamond\exists x\langle\lambda y.Fy\rangle(x)$ true but $\exists x\Diamond\langle\lambda y.Fy\rangle(x)$ false at w. And the model of the previous footnote, modified by changing $\mathcal{D}(w')$ from $\{1\}$ to $\{0, 1\}$ and the extension of F at w' from $\{1\}$ to $\{0\}$, makes $\forall x\Box\langle\lambda y.Fy\rangle(x)$ true but $\Box\forall x\langle\lambda y.Fy\rangle(x)$ false at w.

33. All the results given in Sections 5.1–5.4 for SAT hold as well for SAK, SAB, SAS4, and SAS5.

34. In systems without predicate abstraction both of these sentences count as instances of the Converse Barcan Formula. By an argument similar to the one given for (CBF) in Section 4.4, the first sentence is valid in SAK and hence in SAT, SAB, SAS4, and SAS5. But the second sentence is invalid even in SAS5 and hence in the other systems.

35. As I mentioned in an earlier footnote, Stalnaker [31] noticed this problem in a slightly different context. He gives an interesting example on pp. 335–336, but there is no uniform substitution theorem in [31] or his later [32]. I discuss [32] in Section 6.2.

36. By a primitive predicate abstract I mean one like $\langle\lambda y.Fy\rangle$, in which the formula following the dot is atomic.

37. Where $n = 2$, again define $\text{Ext}_{\mathcal{M}_w}(\theta)$ as $\mathcal{I}(\theta, w)$. Hence $\text{Ext}_{\mathcal{M}_w}(\theta) \subseteq \mathcal{D}(w) \times \mathcal{D}(w))$. Define $\text{Ext}_{\mathcal{M}_w}(\langle\lambda\alpha_1, \alpha_2.\psi\rangle)$ as $\{\langle\mathcal{V}(\alpha_1), \mathcal{V}(\alpha_2)\rangle \mid M_{\mathcal{V}}(\langle\lambda\alpha_1, \alpha_2.\psi\rangle(\alpha_1, \alpha_2), w) = 1\}$, where $\mathcal{V}$ ranges over all valuations relative to $\mathcal{M}$. It follows that $\text{Ext}_{\mathcal{M}_w}(\langle\lambda\alpha_1, \alpha_2.\psi\rangle) \subseteq \mathcal{D}(w) \times \mathcal{D}(w)$.

38. This definition of logical equivalence applies to all formulas. If ψ and ψ' are sentences, it reduces to the definition given in Section 3.6.

39. Since SAT bans vacuous quantifiers, the result of replacing an occurrence of ψ by ψ' in ϕ may not be a formula. (Let ψ, ψ', and ϕ be $Fx \vee \neg Fx$, $Fa \vee \neg Fa$, and $\forall x(Fx \vee \neg Fx)$.) Hence the limitation of Theorem 2 to those cases where ϕ' is a formula.

40. It is sufficient to consider just these cases since all the SAT logical operators can be defined using $\neg$, $\wedge$, $\Box$, A, $\forall$, and λ.

41. See Plantinga [26]–[28], Fine [4], Jager [16], [17], Menzel [24], and Stephanou [36], [37].

42. Stalnaker [32] is reprinted with minor changes in [33].

43. See [28, pp. 91–92], [27, p. 314]. Plantinga apparently does not use the term "actualism" in [26], but on p. 149 he states the denial of what he calls actualism in the previously cited passages, and he argues against it on pp. 149–53.

44. The quantifiers of varying domain logics are sometimes called *actualist* quantifiers, while those of constant domain logics are called *possibilist* quantifiers. (Fitting and Mendelsohn use the terms in this way in [7, pp. 95, 101].) *World bound* and *world independent* would be more accurate terms.

45. See [26, pp. 136, 149].

46. See [26, p. 133].

47. See [26, p. 151].

48. For example, $\neg\langle\lambda y.\Diamond\neg\mathcal{E}(y)\rangle(s)$, $\langle\lambda y.\neg\Diamond\neg\mathcal{E}(y)\rangle(s)$, and $\neg\Diamond\neg\mathcal{E}(s)$ are also false in $\mathcal{M}^{\mathcal{A}}_{\mathcal{I}}$.

49. In a recent paper [25] Percival discusses many issues concerning actualism and the use of predicate abstraction. Space prevents me from going into details, but I suggest that using a distinction I drew earlier in Section 6.1 may help us understand his position. In [25, p. 414, footnote 45] he seems to endorse the view that "... some existing object is such as to exist contingently" and also that "(s)omewhat confusingly ... it is metaphysically possible for there to be objects that do not exist." If "exist" is construed as $\mathcal{E}(x)$ in the former sentence but as $\mathsf{A}\mathcal{E}(x)$ in the latter, they become $\exists x(\mathcal{E}(x) \wedge \neg\Box\mathcal{E}(x))$ and $\Diamond\exists x\neg\mathsf{A}\mathcal{E}(x)$, respectively. Neither of these is valid in SAT, but both are true (i.e., true at the actual world element of the intended model). (If on the other hand, "exist" is construed as $\mathsf{A}\mathcal{E}(x)$ in the former and as $\mathcal{E}(x)$ in the latter, both of the results—$\exists x(\mathsf{A}\mathcal{E}(x) \wedge \neg\Box\mathsf{A}\mathcal{E}(x))$ and $\Diamond\exists x\neg\mathcal{E}(x)$—are the negations of valid sentences.) Perhaps the distinction between, roughly, existence and actual existence is the key to understanding Percival's position.

50. See [27, p. 316]. In [28, p. 93] he expresses it as "...no object x has any property in any world in which x doesn't exist," not as the necessitation of this claim. Since the discussion of serious actualism in [27, p. 316] is more detailed than that in [28, p. 93], and since [27] and [28] were both published in the same volume, it is reasonable to assume that what he intends in [28] is the necessitation of what he says there.

51. For the formula Rxy, $\neg_r Rxy$ can be expressed in SAT as $\langle\lambda z,w.(\mathcal{E}(z) \wedge \mathcal{E}(w)\wedge\neg Rzw)\rangle(x,y)$ and $\Box_r Rxy$ as $\langle\lambda z,w.(\mathcal{E}(z)\wedge\mathcal{E}(w)\wedge Rzw\wedge\Box((\mathcal{E}(z)\wedge\mathcal{E}(w)) \rightarrow Rzw))\rangle(x,y)$, and similarly for formulas, atomic or complex, with free occurrences of three or more variables.

52. The only complication is that SAT must be augmented by infinitely many position variants of each of its variables, and, for each SAT-model-plus-valuation, each position variant of a variable must be assigned the same object as the variable itself.

53. See [8, pp. 278–80].

54. For SAK $(d = d)$ would be added. The sentences in this list that are existential generalizations of conjunctions are similar to, but in general stronger than, the sentences used by Hintikka in formulating a rule of universal instantiation for Q3. The latter sentences are conjunctions of existential generalizations. (See Garson [8, pp. 279–80] for details.)

55. For M_1, M^*, and the Principle of Compossibility, see Chihara [2, pp. 220–24].

56. Stalnaker [32] is reprinted in [33], with a change in the axiomatization and a brief postscript responding to a point made by Williamson. Subsequent references are to [33]. Further work by Stalnaker on modal languages with an abstraction operator can be found in [31], [40], and [34]. His work has been criticized by Williamson, who espouses necessitism, the view that necessarily everything is necessarily something. (In my notation: $\Box\forall x\Box\mathcal{E}(x)$.) For Williamson's work see [42] and the references contained therein.

57. See [33, p. 146].

58. Stalnaker [33, p. 157] says his results hold if the underlying propositional modal logic is K, D, T, KB, B, K4, S4, or S5. The similarities and differences noted in the text between SAT and Stal-T also hold when the underlying modal logic is any of those listed.

59. See [33, p. 157].

60. In the postscript to the version of the paper that appears in [33], Stalnaker considers some effects of adding the actuality connective. For ease of comparison I restrict my attention here to formulas that do not contain it.

61. In Stal-T quantifiers apply directly to unary predicates, whether simple or complex, so $\forall S$ and $\exists S$ are well formed.

62. Although Menzel's ontology includes properties and relations, and he uses them in defining the model in question, that model consists only of pure sets. In [30] Ray argues that pure sets alone suffice for defining such a model. I believe Ray is correct, but I will not address the issue here.

63. If in constructing such a model we attempt to represent "all of reality," including both the physical world and the entire set-theoretic universe, cardinality problems will of course arise. But I can see no barrier to constructing models of this kind that, while not completely comprehensive, are large enough to be interesting.

64. I have discussed related matters elsewhere. In [11] I argue that sentences of the form $(\mathsf{A}p \rightarrow p)$ (where p is not itself a necessary truth) are contingent yet true a priori. In [12] I argue that such sentences are also synthetic. (See especially [11, p. 447] and [12, Section 2.3].)

65. Namely, SAK, SAB, SAS4, and SAS5.

66. The individual constants and variables of $\mathfrak{L}$, as defined in Section 2.1, may be indexed with numerical subscripts as a way of ensuring that their supply is unlimited. Strictly speaking, the new grounded terms of $\mathfrak{L}^\star$ should have this feature as well. In practice,

one seldom needs more than a few such symbols of either language and thus seldom needs these purely indexical subscripts. When such a need arises, indexical and tableau-node-prefix subscripts can be distinguished by placing the latter in parentheses after the former. Thus $b_{257(1.2.@.2.1)}$ counts as a grounded name, and $p_{3(1.1.1)}$ as a parameter. In general, a grounded term can be represented as $\tau_{\kappa(\sigma)}$, where τ is one of the lower case letters a through t, κ is an indexical subscript, and σ is a tableau prefix.

67. Subscripting an individual constant τ with a prefix is similar to applying Kaplan's [18] operator *dthat* to τ.

68. For more on tableau rules for K, T, B, S4, and S5 see [7, pp. 51–55].

69. Thus τ_σ is either a parameter or a grounded name.

70. Form 3 can be simplified so as to introduce only the second conclusion stated here. The first conclusion can be derived using form 2 and the substitutivity of identity rule (see Section 8.2.5).

71. Form 3 need not be taken as primitive. Its conclusion can be obtained using the substitutivity of identity rule (see Section 8.2.5) and form 2 of the negated predicate abstraction rule.

72. The blank space above the horizontal line indicates that this rule has no premise.

73. Although atomic formulas other than identities do not contain parentheses (see Section 2.2), the use of parentheses in the metalinguistic representation of atomic formulas enhances readability.

74. It is important to notice that forms 1 and 2 apply to the case where θ is the identity predicate. If ι is an individual constant, $\theta(\ldots\iota\ldots)$ is $(\iota = \ldots)$ or $(\ldots = \iota)$. Similarly, if τ_{σ_1} is a grounded term, $\theta(\ldots\tau_{\sigma_1}\ldots)$ is $(\tau_{\sigma_1} = \ldots)$ or $(\ldots = \tau_{\sigma_1})$.

75. This form does not apply to parameters, since the result of removing the tableau-prefix subscript from a parameter is not a symbol of either $\mathfrak{L}$ or $\mathfrak{L}^\star$.

76. As in the case of forms 1 and 2, it is important to notice that form 3 applies where θ is the identity predicate, ι_σ is a grounded name, and $\theta(\iota_\sigma)$ is thus $(\iota_\sigma = \ldots)$ or $(\ldots = \iota_\sigma)$.

77. This tableau suggests a two-world SAS5 model that falsifies $\exists x\Diamond\neg Fx \rightarrow \Diamond\exists x\neg Fx$. The domain of one world contains a single object, and the predicate F holds of that object at that world. The domain of the other world is empty. Alternatively the second world can contain a distinct object, with F also holding of that object at that world.

78. The model $\mathcal{M}$ and valuation $\mathcal{V}$ given here apply to the language $\mathfrak{L}$. Since grounded terms (i.e., grounded names and parameters) are not symbols of $\mathfrak{L}$ they are assigned no values by either $\mathcal{M}$ or $\mathcal{V}$. But if we add the stipulation that $\mathcal{V}$ assigns $\{p_1, d_1, q_{1.1}, r_{1.1}\}$ to $p_1, d_1, q_{1.1}$, and $r_{1.1}$, and we leave $\mathcal{M}$ unchanged, the result is a model and a valuation of $\mathfrak{L}^\star$ that falsify *faux* (UI).

79. Possibilists may complain that line 2 begs the question against them because it assumes that the object denoted by $p_@$ is an actual object. But if quantifiers are given a possibilist interpretation (Act3) is valid, a result that possibilists presumably do not want.

80. The definitions and proofs concerning soundness are similar to those of Fitting and Mendelsohn [7, pp. 57–60, 121–24, 152–53].

81. As specified in Section 8.1.1, all grounded terms of $\mathfrak{L}^\star$ are treated semantically like variables. Given a model $\mathcal{M}$ and a valuation $\mathcal{V}$ relative to $\mathcal{M}$, $\mathcal{V}$ assigns to each grounded term a single, model-wide denotation from the domain of the model, $\mathcal{D}_{\mathcal{M}}$. Clauses 3 and 4 reflect this, since together they require that any term grounded with σ designates, at each world of the model, a member of $\mathcal{D}(\Theta(\sigma))$. And of course $\mathcal{D}(\Theta(\sigma)) \subseteq \mathcal{D}_{\mathcal{M}}$.
    Note that $\mathcal{I}(\iota, \Theta(\sigma))$ may be defined even if ι_σ does not appear in S.

82. For example, clause 3 requires that the parameter p_1 designates, throughout the model, an object from the domain of world $\Theta(1)$. Clause 4 requires that the grounded name $c_{1.1}$ designates, throughout the model, an object from the domain of world $\Theta(1.1)$. By contrast, if no occurrence of the individual constant c has a tableau prefix as a subscript, clause 4 does not require that c be given an assignment by $\mathcal{M}$ *cum* $\mathcal{V}$. Individual constants are not required to denote at every, or even any, world.

83. $\mathcal{I}$ does not, however, assign values to the grounded terms of $\mathfrak{L}^\star$ or to the variables of $\mathfrak{L}$. These are given their values by $\mathcal{V}$.

84. It is important to bear in mind in what follows that individual constants and grounded names are treated very differently in the semantics of $\mathfrak{L}^\star$. In the categorical model and valuation ($\mathcal{M}$ and $\mathcal{V}$) just defined, this difference may manifest itself in initially surprising ways. For example, if the grounded names b_1, $b_{1.1}$, and $b_{1.1.1}$ appear on a branch, each will have been introduced by application of a tableau rule to the individual constant b. Yet it may turn out that, under $\mathcal{M}$ and $\mathcal{V}$, $\overline{b_1} \neq \overline{b_{1.1}} \neq \overline{b_{1.1.1}}$.

85. Strictly speaking, Sections 3.2 and 3.4 present only the semantics of $\mathfrak{L}$. $\mathfrak{L}^\star$, introduced in Section 8.1.1, is $\mathfrak{L}$ plus grounded names and parameters, which are treated like variables.

References

[1] Burge, T., "Truth and singular terms," *Noûs*, vol. 8 (1974), pp. 309–25. MR 1047931. 274, 275, 277

[2] Chihara, C. S., *The Worlds of Possibility: Modal Realism and the Semantics of Modal Logic*, Oxford University Press, New York, 1998. Zbl 0901.03005. MR 1695518. 280

[3] Davies, M., and L. Humberstone, "Two notions of necessity," *Philosophical Studies*, vol. 38 (1980), pp. 1–30. MR 0582007. DOI 10.1007/BF00354523. 276

[4] Fine, K., "Model theory for modal logic, III: Existence and predication," *Journal of Philosophical Logic*, vol. 10 (1981), pp. 293–307. Zbl 0464.03017. MR 0628091. DOI 10.1007/BF00293421. 279

[5] Fine, K., *Modality and Tense: Philosophical Papers*, Oxford University Press, Oxford, 2005. Zbl 1082.03002. 276

[6] Fitting, M., *Types, Tableaus, and Gödel's God*, vol. 12 of *Studia Logica Library*, Kluwer, Dordrecht, 2002. Zbl 1038.03001. MR 2041528. DOI 10.1007/978-94-010-0411-4. 275

[7] Fitting, M., and R. L. Mendelsohn, *First-Order Modal Logic*, vol. 277 of *Synthese Library*, Kluwer, Dordrecht, 1998. Zbl 1025.03001. MR 1956131. DOI 10.1007/978-94-011-5292-1. 238, 255, 256, 267, 275, 276, 279, 281, 282

[8] Garson, J. W., "Quantification in modal logic," pp. 267–323 in *Handbook of Philosophical Logic, Vol. 3*, edited by D. M. Gabbay and F. Guenther, Kluwer, Dordrecht, 2001. Zbl 1003.03519. MR 1884635. 252, 280

[9] Gibbard, A., "Contingent identity," *Journal of Philosophical Logic*, vol. 4 (1975), pp. 187–221. Zbl 0317.02028. MR 0505172. DOI 10.1007/BF00693273. 277

[10] Gilbert, D. R., and E. D. Mares, "Completeness results for some two-dimensional logics of actuality," *Review of Symbolic Logic*, vol. 5 (2012), pp. 239–58. Zbl 1260.03042. MR 2924361. DOI 10.1017/S1755020311000384. 275

[11] Hanson, W. H., "Actuality, necessity, and logical truth," *Philosophical Studies*, vol. 130 (2006), pp. 437–59. MR 2263638. DOI 10.1007/s11098-004-5750-8. 255, 275, 276, 280

[12] Hanson, W. H., "Logical truth in modal languages: Reply to Nelson and Zalta," *Philosophical Studies*, vol. 167 (2014), pp. 327–39. MR 3148196. DOI 10.1007/s11098-012-0088-0. 255, 275, 276, 280

[13] Hodes, H. T., "Axioms for actuality," *Journal of Philosophical Logic*, vol. 13 (1984), pp. 27–34. Zbl 0538.03015. MR 0737516. DOI 10.1007/BF00297575. 275

[14] Hodes, H. T., "On modal logics which enrich first-order S5," *Journal of Philosophical Logic*, vol. 13 (1984), pp. 423–54. Zbl 0553.03008. MR 0765995. DOI 10.1007/BF00247714. 275

[15] Hodes, H. T., "Some theorems on the expressive limitations of modal languages," *Journal of Philosophical Logic*, vol. 13 (1984), pp. 13–26. Zbl 0538.03014. MR 0737515. DOI 10.1007/BF00297574. 275

[16] Jager, T., "An actualistic semantics for quantified modal logic," *Notre Dame Journal of Formal Logic*, vol. 23 (1982), pp. 335–49. MR 0658802. 250, 279

[17] Jager, T., "De re and de dicto," *Notre Dame Journal of Formal Logic*, vol. 29 (1988), pp. 81–90. Zbl 0664.03005. MR 0932696. DOI 10.1305/ndjfl/1093637772. 250, 279

[18] Kaplan, D., "On the logic of demonstratives," *Journal of Philosophical Logic*, vol. 8 (1979), pp. 81–98. Zbl 0415.03004. MR 0527404. DOI 10.1007/BF00258420. 281

[19] Kripke, S. A., "Semantical considerations on modal logic," *Acta Philosophica Fennica*, vol. 16 (1963), pp. 83–94. Zbl 0131.00602. MR 0170800. 241, 245, 275, 277, 278

[20] Lambert, K., *Free Logic: Selected Essays*, Cambridge University Press, Cambridge, 2003. Zbl 1026.03003. MR 1950119. DOI 10.1017/CBO9781139165068. 275

[21] Lambert, K., and E. Bencivenga, "A free logic with simple and complex predicates," *Notre Dame Journal of Formal Logic*, vol. 27 (1986), pp. 247–56. Zbl 0614.03010. MR 0842151. DOI 10.1305/ndjfl/1093636615. 275

[22] Mates, B., *Elementary Logic*, 2nd edition, Oxford University Press, New York, 1972. MR 0321678. 247

[23] Menzel, C., "Actualism, ontological commitment, and possible world semantics," *Synthese*, vol. 85 (1990), pp. 355–89. MR 1087292. DOI 10.1007/BF00484834. 249, 254

[24] Menzel, C., "The true modal logic," *Journal of Philosophical Logic*, vol. 20 (1991), pp. 331–74. Zbl 0748.03015. MR 1136703. DOI 10.1007/BF00249434. 242, 276, 277, 279

[25] Percival, P., "Predicate abstraction, the limits of quantification, and the modality of existence," *Philosophical Studies*, vol. 156 (2011), pp. 389–416. MR 2852018. DOI 10.1007/s11098-010-9609-x. 279

[26] Plantinga, A., *The Nature of Necessity*, Oxford University Press, Oxford, 1974. MR 0460074. 249, 279

[27] Plantinga, A., "Replies to my colleagues," pp. 313–96 in *Alvin Plantinga*, edited by J. Tomberlin and P. van Inwagen, Springer, Dordrecht, 1985. 274, 279

[28] Plantinga, A., "Self-profile," pp. 3–97 in *Alvin Plantinga*, edited by J. Tomberlin and

P. van Inwagen, Springer, Dordrecht, 1985. 274, 279

[29] Prior, A. N., and K. Fine, *Worlds, Times and Selves*, Duckworth, London, 1977. 276

[30] Ray, G., "Ontology-free modal semantics," *Journal of Philosophical Logic*, vol. 25 (1996), pp. 333–61. Zbl 0860.03018. MR 1400149. DOI 10.1007/BF00249664. 280

[31] Stalnaker, R. C., "Complex predicates," *The Monist*, vol. 60 (1977), pp. 327–39. 248, 278, 280

[32] Stalnaker, R. C., "The interaction of modality with quantification and identity," pp. 12–28 in *Modality, Morality, and Belief: Essays in Honor of Ruth Barcan Marcus*, edited by W. Sinnott-Armstrong, D. Raffman, and N. Asher, Cambridge University Press, Cambridge, 1995. 244, 248, 252, 275, 277, 278, 279, 280

[33] Stalnaker, R. C., *Ways a World Might Be*, Oxford University Press, Oxford, 2003. 275, 279, 280

[34] Stalnaker, R. C., *Mere Possibilities: Metaphysical Foundations of Modal Semantics*, Princeton University Press, Princeton, 2012. 280

[35] Stalnaker, R. C., and R. H. Thomason, "Abstraction in first-order modal logic," *Theoria*, vol. 34 (1968), pp. 203–7. MR 0337500. DOI 10.1111/j.1755-2567.1968.tb00351.x. 275

[36] Stephanou, Y., "Investigations into quantified modal logic," *Notre Dame Journal of Formal Logic*, vol. 43 (2002), pp. 193–220. Zbl 1050.03017. MR 2034746. DOI 10.1305/ndjfl/1074396306. 251, 276, 279

[37] Stephanou, Y., "First-order modal logic with an 'actually' operator," *Notre Dame Journal of Formal Logic*, vol. 46 (2005), pp. 381–405. Zbl 1092.03010. MR 2183050. DOI 10.1305/ndjfl/1134397658. 251, 276, 279

[38] Stephanou, Y., "Serious actualism," *Philosophical Review*, vol. 116 (2007), pp. 219–50. 251, 274

[39] Thomason, R. H., and R. C. Stalnaker, "Modality and reference," *Noûs*, vol. 2 (1968), pp. 359–72. 275

[40] Thomson, J., and A. Byrne, *Content and Modality: Themes from the Philosophy of Robert Stalnaker*, Oxford University Press, Oxford, 2006. 280

[41] Vlach, F., "'Now' and 'then': A formal study in the logic of tense and anaphora," Ph.D. dissertation, University of California, Los Angeles, 1973. 276

[42] Williamson, T., *Modal Logic as Metaphysics*, Oxford University Press, Oxford, 2013. Zbl 1294.03008. MR 3202317. DOI 10.1093/acprof:oso/9780199552078.001.0001. 280

Acknowledgments

I am grateful to Jack Woods, Christopher Menzel, and an anonymous referee for many helpful suggestions.

Department of Philosophy
University of Minnesota
Minneapolis, Minnesota 55455
USA
whanson@umn.edu