MODAL MODELS
FOR BRADWARDINE'S
THEORY OF TRUTH

Greg Restall*

Philosophy Department
The University of Melbourne
restall@unimelb.edu.au

Abstract: Stephen Read has recently discussed Bradwardine's theory of truth, and defended it as an appropriate way to treat paradoxes such as the liar [13, 14]. In this paper, I discuss Read's formalisation of Bradwardine's theory of truth, and provide a class of models for this theory. The models facilitate comparison of Bradwardine's theory with contemporary theories of truth.

There are many different approaches to the logic of truth. We could agree with Tarski, that the appropriate way to formalise a truth predicate is in a hierarchy, in which the truth predicate in one language can apply only to sentences from another language. Or, we could attempt to do without type restrictions on the truth predicate. Bradwardine's theory of truth takes the second of these options: it is type-free, and admits sentences which say of themselves that they are not true to be well-formed. We could take the behaviour of the paradoxes such as the liar to motivate a revision of the basic logic of propositional inference, to allow for truth-value gaps or gluts [9, 11, 15]. On the other hand, we could take it that the paradoxes are no reason to revise our account of the basic laws of logic: a novel account of the behaviour of the truth predicate is what is required. Bradwardine's account, as elaborated by Read, takes this second option.[1] Finally,

---

[1]This is a subtle matter: Read prefers a non-classical *relevant* account of logical consequence [12], but this is not motivated by the theory of truth. Read's re-telling of Bradwardine's account of truth is not revisionary for Read's own theory of logic. And as we shall see, there is no need to reject classical logic in order to follow Bradwardine's analysis.

we could attempt to give an account of the truth predicate on its own terms, by elaborating the inference rules governing truth. Or we could attempt to define the truth predicate in terms of other notions. Contemporary theories of truth have largely eschewed this second option, in favour of attempting to retain as much of Tarski's biconditional $A \equiv T^{\ulcorner}A^{\urcorner}$, and to mount a rescue operation for a classical theory of truth by way of limiting our diet — by putting the theory on a tightened regimen from among the smorgasbord of biconditionals.[2] Like all classical theories of truth that allow us to express liar sentences, Bradwardine's theory agrees that we cannot swallow the lot. However, Bradwardine gives us guidance as to *which* T-biconditionals are to be accepted, by *defining* the truth predicate in terms of a more fundamental notion.

So, Bradwardine's theory of truth is type-free, it is non-revisionary with regard to logic, and it defines the truth predicate in terms of other notions. The T-biconditionals that the theory accepts are those that follow from more primitive notion in terms of which the truth predicate is defined.

1 THE SYNTAX

The crucial notion in terms of which truth is defined is expressed grammatically by neither a *predicate* nor an *operator*. The concept is simple: it is expressed in claims of the form 'x says that p' or 'x signifies that p'. The phrase 'says that' is syntactically and semantically a hybrid. To form a sentence, we can substitute a name or referring expression in place of 'x,' and another sentence in place of 'p.' These are some examples of sentences utilising *says that* in the intended sense.

1. The first sentence of Chapter 1 of Coffa's *Semantic Tradition* says that for better and worse, almost every philosophical development of significance since 1800 has been a response to Kant.[3]

2. "Schnee ist weiß" says that snow is white.

3. Fee's utterance "I'm happy here" says that she is happy at the University of Melbourne.

4. "This very sentence is false" says that that very sentence is false.

5. There is some sentence (say, x) which says that x is not true.

In the (1) we pick out a sentence by giving a definite description. In (2) and (4) we give a quotation name for a sentence (perhaps a particular use of a sentence), and in (3) we give a quotation of an utterance. In (5) we quantify into that place.

---

[2]For one theory of truth among many which attempts to defend this "T-biconditionals on a diet" approach, see Horwich's *Truth* [8]. For one example of an account of truth that treats the truth predicate as primitive, but attempts to give a principle for distinguishing which T-biconditionals are to be accepted, see Leitgeb's "What Truth Depends On" [10].

[3]J. Alberto Coffa, *The Semantic Tradition from Kant to Carnap* [2].

In (1) and (2), complete sentences are found in the second place. In (3) it is a complete record of an utterance in a context. In (4) the second component of 'says that' is a complete sentence "that very sentence is false", but it contains an anaphoric reference (the "that very sentence") to the first part of the larger sentence in which it is embedded. In (5) we quantify in to the second place: a variable of quantification (the 'x') occurs inside the position, while it is bound by a quantifier *outside* the position. All of these examples will count as well-formed in the theory to be elaborated.

It may be very tempting to read "says that" as relating an *agent* to a propositional content, but we will not read it in this way. Our primary sense is Bradwardine's sense of signifying. Fee did say that she was happy at the University of Melbourne. However, in our sense, she did this derivatively by producing an utterance which signified to that effect.

To say "x says that p" is to predicate a property of the object (sentence or utterance or other item) x (keeping p fixed), or to modify the sentence p (keeping x fixed). In other words, "says that p" is a predicate for each choice of p, and "x says that" is an operator, for each choice of x. The type distinction in the grammar of "says that" wears on its sleeve the distinction between *use* and *mention*. We shall define a truth predicate, and since it is a predicate, we shall *mention* the sentences (or judgements or beliefs or other objects) that we shall describe as true. On the other hand, it is not enough to merely predicate truth of objects willy nilly or in an unprincipled fashion. Of course, we must predicate truth of something in terms of its meaning: in other words, in terms of what it *says*. Here, we make the jump from mention to *use*.

In what follows we shall write "x says that p" with the shorter expression "x : p" — our syntax extends the grammar of first order logic with the clause that if t is a term and A is a sentence, then t : A is also a sentence.

Bradwardine's first contribution is the definition of truth in terms of our new conceptual machinery. To be true is to say what is the case. To be false, is to say what is not the case. To be *declarative* is to say something. Allowing for the possibility (which will later become crucial) that individual expressions or utterances can say more than one thing (or to speak more neutrally, could say that p and say that q where ⌜p⌝ and ⌜q⌝ differ) then in the account of truth, we must ask whether it suffices for truth that the expression say *something* that is the case, or if we require the stronger condition that *everything* that is says is the case. Bradwardine opts for the stronger of the two conditions. For falsity, he opts for the weaker of the two conditions: since falsity is untruth (at least, falsity is untruth for declarative objects — rocks and other objects that say nothing are neither true nor false), it suffices for falsity that the object in question says something that is not the case.

Now, the reader has almost certainly noticed that when talking of 'something' that is the case or that 'everything' that it says is the case, we have been quantifying directly into the second place of 'says that' — we have been quantifying into a position in which a sentence is *used* and not *mentioned*. This is the

second crucial piece of machinery in Bradwardine's theory: we need to make use of what has come to be called propositional quantification.

At this point we must clarify a distinction between Bradwardine's terminology and modern vocabulary: quoting Read [13, page 191] we have

> DEFINITION 1: *A true proposition is an utterance signifying only as things are.*

> DEFINITION 2: *A false proposition is an utterance signifying other than things are.*

For Read and for Bradwardine, what is called a 'proposition' is what is denoted by the referring expression in the first part of the claim of the form 'x : p.' The proposition is the utterance or sentence or other *object* of which we are predicating truth. To confuse matters, we now use 'proposition' to describe quantification into the *second* position, in which the sentence is mentioned but not used. Often we might think of the semantic value or significance of a sentence in use as a *proposition*, whether this is conceived of as a set of possible worlds, or a structured meaning or some other thing. To keep matters as clear and precise as I can, I will describe the objects of which we will predicate truth and falsity *sentences* (thinking primarily of tokens, and not types, though we may sometimes predicate truth of types if the sentences do not vary too much in significance from context to context), or *utterances* or other such things. I will avoid talking about propositions as much as possible in what follows, except when talking of what is now called propositional quantification.

Propositional quantification is relatively straightforward *syntactically*. We need only a stock of variables which may occur in sentence position, and quantifiers that bind them. In many formal frameworks, the *semantics* of propositional quantification is no more problematic: in Kripke semantics, or other model theories in which sentences are interpreted as sets of points, we allow any such set of point to be the possible denotation of a propositional variable, and quantification is modelled by varying the possible denotations of the corresponding variable in the usual manner. In proof theoretical semantics, propositional quantification is merely a degenerate form of second-order quantification in which the *predicate* variable quantified is zero-place. The relevant rules are straightforward to define. So, *formally* speaking, propositional quantification poses no difficulty. However, its interpretation is a matter of dispute.[4] Despite this difficulty, I will take propositional quantification as a given, since it is necessary for the development of Bradwardine's account of truth — at least as Read reconstructs it. Perhaps any success of the account will help in some small measure to vindicate propositional quantification by way of its fruits.

---

[4]Some theories of truth in which propositional quantification plays a role are Horwich's minimalism [8], and Grover's prosentential theory [6]. Dorothy Grover's work contains an interesting defence of the significance of propositional quantification [5, 6, 7], and this work has been taken up in Brandom's inferentialist programme [1].

Greg Restall, restall@unimelb.edu.au   OCTOBER 26, 2007

The crucial definitions of the *truth* predicate, the *falsity* predicate and the predicate of *declarativeness* are now possible:

DEFINITION [DECLARATIVENESS, TRUTH, FALSITY] For a language with ":" and propositional quantification, we may define the predicates D, T and F as follows:

» Dx is $(\exists p)(x : p)$

» Tx is $Dx \wedge (\forall p)((x : p) \supset p)$

» Fx is $(\exists p)((x : p) \wedge \neg p)$

So, an object is declarative if it says something, and it is false if it says something that is not the case. It is true otherwise: if it says something, and that everything it says is the case.[5]

Given the usual classical behaviour of the propositional connectives, and given straightforward properties of propositional quantification, we have the following immediate consequences of the definition.

LEMMA [CONSISTENCY AND BIVALENCE] $(\forall x)\neg(Tx \wedge Fx)$ *and* $(\forall x)(Dx \equiv (Tx \vee Fx))$

It is worth stepping through the proofs to see what properties of propositional quantification are required.

Proof: To show consistency, suppose that Ta. It follows that Da, and that $(\forall p)(a : p \supset p)$. If it were the case that Fa, then we would have $(\exists p)(a : p \wedge \neg p)$, so for some p, $a : p$ and $\neg p$. But by Ta, we have $a : p \supset p$, so we have a contradiction. It follows that $\neg(Ta \wedge Fa)$, and hence, $(\forall x)\neg(Tx \wedge Fx)$.

For bivalence, from right to left, notice that if Ta then Da by definition. If Fa then $(\exists p)(a : p \wedge \neg p)$ (a says something false), so it follows that $(\exists p)a : p$ (a says something), so Da. So, $(Ta \vee Fa) \supset Da$. Conversely, suppose that Da. If $\neg Fa$, then $\neg(\exists p)(a : p \wedge \neg p)$, and by elementary quantifier moves and classical connective manipulation, we have $(\forall p)(a : p \supset p)$. In other words, given Da, if $\neg Fa$ then Ta. So, if Da then $Ta \vee Fa$, and we have proved the biconditional in both directions. ∎

So, truth and falsity provide a mutually exclusive and exhaustive categorisation of the declarative objects. Truth and falsity behave as one would expect, on the basis of very weak assumptions: the presence of ":" and propositional quantification. Not only this, but we can see Bradwardine's analysis of paradoxes like the liar.

FACT [LIAR SENTENCES ARE FALSE (AND NOT TRUE)] *If* $\lambda : \neg T\lambda$, *then it follows that* $F\lambda$ (*and* $\neg T\lambda$ *as well*).

---

[5]Continually bracketing '$x : p$' will get tedious. In what follows we will read all expressions involving ':' with ':' binding most tightly. So, Tx is an abbreviation of $Dx \wedge (\forall p)(x : p \supset p)$. If I wished to say that x said that $p \supset p$, I would use brackets around the conditional like this: $x : (p \supset p)$.

Proof: The proof is straightforward. If $\lambda : \neg T\lambda$, then we have $D\lambda$, since λ says something. Now, if $T\lambda$, then it would follow that $(\forall p)(\lambda : p \supset p)$, and hence $\lambda : \neg T\lambda \supset \neg T\lambda$. But we have agreed that $\lambda : \neg T\lambda$, so it would follow (if λ were true, i. .e., if $T\lambda$) then $\neg T\lambda$. In other words, we have proved $T\lambda \supset \neg T\lambda$, which is simply $\neg T\lambda$. The liar sentence λ is not true. Since it is declarative, it must follow that it says *something* that is not the case — and hence, it follows that λ is false, as well. ∎

Since λ is false, it says something that is not the case. But what else could λ say that is not the case? The assumptions that we have made so far do not tell us. We know that $\lambda : \neg T\lambda$, and this, in fact, *is* the case (it is the case that $\neg T\lambda$). But my utterance that $\neg T\lambda$ must perforce, differ from λ itself, for my utterance may well be true, and λ is not. The definitions tell us that there *is* something else that is the λ says that is not the case, without isolating what it might be.

One candidate is that λ says that $T\lambda$. $T\lambda$ is *not* the case, and if λ declares its own truth, then this is where it goes wrong and is false. According to Read [13], Bradwardine took himself to have an *argument* for the conclusion that the liar says of itself that it is true. This distinguishes Bradwardine's account from the later account of John Buridan which *assumes* without argument that every declarative says of itself that it is true.

Bradwardine's argument for this claim is a subtle one [13]. We will not consider it here. Instead, we need to develop the theory a little more, for we have not quite arrived at Bradwardine's theory of truth. What makes Bradwardine's account distinctive is not the definitions of truth or falsity, or the treatment of the falsity of the insolubilia. Instead, it is the following postulate, which I shall call "Bradwardine's Axiom."

DEFINITION [BRADWARDINE'S AXIOM] Every proposition signifies or means contingently or necessarily everything which follows from it contingently or necessarily [13].

We may render the condition in the following way:

$$\text{If } x : p \text{ then if (if } p \text{ then } q\text{) then } x : q$$

Rearranging the conditionals, we might have another formulation

$$\text{If (if } p \text{ then } q\text{) then (if } x : p \text{ then } x : q\text{).}$$

The crucial issue in understanding Bradwardine's axiom is what form of conditional expression might be used in formulating it. If all conditionals are material, then we have

$$(p \supset q) \supset (x : p \supset x : q)$$

which we might call the *material Bradwardine's Axiom*. While perhaps initially appealing, this reduces Bradwardine's theory to one in which $x : p$ can be determined by whether or not Dx, Tx and Fx.

FACT [BRADWARDINE'S COLLAPSE] *Granted the material Bradwardine axiom, it follows that the non-declarative objects (the* x *such that* ¬Dx*) say nothing, the false declarative objects say* everything *and the true declarative objects say all and only what is the case.*

Proof: The proof turns on the behaviour of the material conditional. We have

$$q \supset (x : p \supset x : q)$$

since q entails p ⊃ q. It follows from this that if q is the case, then if x says anything (if x is declarative) then x says q. All declarative objects say everything that is the case. On the other hand, we have

$$\neg p \supset (x : p \supset x : q)$$

In other words, if there is some p that x says, that is not the case, then x says that q too. But q is arbitrary. In other words, if x is false, then x says everything.

All that remains are the non-declarative objects, but by definition, these are those that say nothing at all. It follows that in the presence of the material Bradwardine's axiom, the falsehoods say everything, the truths say all and only what is the case, and the non-declaratives say nothing. What is said collapses into this tripartite division. ∎

There are two options to consider in retaining Bradwardine's axiom without collapse. The first is to move to a kind of *relevant* implication. This is Read's preferred option, and it has the virtue of explicitly allowing for what Bradwardine draws to our attention: both necessary and contingent consequences. Relevant implication may be contingent without being material or extensional. However, in this paper the relevant approach is not the path that I will take, since relevant model theory (especially the model theory for quantified relevant logic), is rather subtle [3, 4], and the question of a natural semantics for relevant logics with propositional quantification is a completely open question.

Instead, we shall steer a simpler course, and ignore contingent consequences to concentrate on necessary ones. For the rest of this paper, I shall consider the consequences of what I shall call the *modal* Bradwardine axiom

$$\Box(p \supset q) \supset \Box(x : p \supset x : q)$$

if p *entails* q, then x : p *entails* x : q. This axiom has the virtue of allowing for a straightforward model theory. Models for necessity are well-understood, and if we take the appropriate logic of necessity to be s5 — as we shall — then they are extremely simple. A model is simply a collection of points, which we shall call 'worlds.' Sentences will either hold or not at each world, and then □A holds at a world if and only if A holds at *every* world in the model.

With this understanding of models for necessity — the simplest understanding I can imagine — then the modal Bradwardine axiom takes on another colour.

The antecedent $\Box(p \supset q)$ simply tells us that all the worlds at which p holds are worlds at which q holds. The consequent $\Box(x : p \supset x : q)$ tells us that at any world at which $x : p$ then $x : q$ too. This is *automatically* satisfied if we think of 'x : ' as a normal modal operator. If all of the p worlds are q worlds, then if from the point of view of my world, I check the x-alternative worlds (those in the accessibility relation appropriate for 'x : ') and find that they are p-worlds, it follows immediately that they are all q-worlds too. So reading 'x : ' as a modal operator will supply us with the modal Bradwardine axiom. This idea will be developed in the next section.

However, before stepping straight into the modal model theory, we should check that the immediate consequences of taking 'x : ' to be a normal modal operator are not unpalatable. The first property of a normal modal operator (closure under entailment — if p entails q then $x : p$ entails $x : q$) is, as we have seen, a reading of Bradwardine's axiom, and so it is welcome. The next characteristic principle is commuting with conjunctions:

$$(x : p) \wedge (x : q) \supset x : (p \wedge q)$$

This is not a consequence of the Bradwardine axiom on its own. It is a different principle.[6] However, this principle seems not implausible in its own right. If x says that p and it says that q, it is certainly not stretching things to take it that x has said that $p \wedge q$. It seems that taking 'x says that' to be a normal modal operator is not doing an injustice to the concept. Now let us see what can be done to model it.

2   MODAL MODELS

As we have already indicated, models for normal modal logics can be given by way of Kripke frames. In our case, we have a *quantified* modal logic, in which some of the objects in question are declarative. In this paper we will make as many simplifying decisions as possible, consistent with the goal to construct interesting models. In this case, it means taking frames for our modal logic to be constructed in the following way:

DEFINITION [BRADWARDINE FRAMES] A *Bradwardine frame* is a structure of the form $\langle W, O, D, \{R_d : d \in D\}\rangle$ where

» A non-empty set W of *worlds*

» A non-empty set O of *objects*

» A subset D of O of *declarative* objects

» A relation $R_d \subseteq W \times W$ for each $d \in D$

---

[6]It is not difficult to construct a neighbourhood frame for a non-normal logic satisfying the Bradwardine axiom but not this conjunction principle.

The simplifying assumption here is that whether or not an object is declarative is an all or nothing matter. Declarative objects do their declarative work in every world, rather than in some worlds and not others. This feature simplifies presentation, as we an define the accessibility relation for each declarative object to be a relation on the entire frame.

Given a Bradwardine frame, we can interpret sentences in our language in the usual way. Sentences hold or fail in each world, so the *extension* $[\![A]\!]$ of a sentence is the set of worlds in which it holds. Similarly, for an n-place predicate, its extension is an n-ary relation on the set O of objects, possibly varying from world to world.

DEFINITION [BRADWARDINE MODELS] An evaluation on a Bradwardine frame is provided by giving an extension to every relation in the language, and a denotation for every name. The *variables* in the language will be interpreted with the aid of an assignment α of values to variables. The value α assigns to an *objectual* variable such as x is an object $[\![x]\!]_\alpha$ in O. The value that α assigns to a propositional variable such as p is a *set* $[\![p]\!]_\alpha$ of worlds.

Given such an assignment of values to atomic expressions, we can assign values to complex expressions, relative to the choice of a world and the choice of an assignment of values to variables.

» $\mathfrak{M}, \alpha, w \Vdash Rt_1 \cdots t_n$ iff the n-tuple $\langle [\![t_1]\!]_\alpha, \ldots, [\![t_n]\!]_\alpha \rangle$ is in the extension $[\![R]\!]$ at world w.

» $\mathfrak{M}, \alpha, w \Vdash p$ iff w is in $[\![p]\!]_\alpha$.

» $\mathfrak{M}, \alpha, w \Vdash \neg A$ iff $\mathfrak{M}, \alpha, w \nVdash A$.

» $\mathfrak{M}, \alpha, w \Vdash A \wedge B$ iff $\mathfrak{M}, \alpha, w \Vdash A$ and $\mathfrak{M}, \alpha, w \Vdash A$.

» $\mathfrak{M}, \alpha, w \Vdash (\forall x)A$ iff $\mathfrak{M}, \alpha', w \Vdash A$ for every x-variant α' of α.

» $\mathfrak{M}, \alpha, w \Vdash (\forall p)A$ iff $\mathfrak{M}, \alpha', w \Vdash A$ for every p-variant α' of α.

» $\mathfrak{M}, \alpha, w \Vdash \Box A$ iff $\mathfrak{M}, \alpha, v \Vdash A$ for every world v.

These clauses are completely standard. They model constant domain quantified s5 with a universal accessibility relation, and with propositional quantification ranging over every subset of worlds. (We will call this logic $QS5^{\forall p}$.) The innovation is in the treatment of 'says that'.

» $\mathfrak{M}, \alpha, w \Vdash t : A$ iff $[\![t]\!]_\alpha \in D$ and for each v where $wR_{[\![t]\!]_\alpha}v$, we have $\mathfrak{M}, \alpha, v \Vdash A$.

If $[\![t]\!]$ is a declarative object, then 't :' functions as a normal modal operator, using the accessibility relation $R_{[\![t]\!]}$. If on the other hand, $[\![t]\!]$ is not declarative, then $t : A$ is always false.

This completes our definition of Bradwardine models.

LEMMA [TRUTH & FALSITY IN BRADWARDINE MODELS] *In any Bradwardine model we have*

» $\mathfrak{M}, \alpha, w \Vdash Dt$ *iff* $[\![t]\!]_\alpha \in D$.

» $\mathfrak{M}, \alpha, w \Vdash Tt$ *iff* $[\![t]\!]_\alpha \in D$ *and* $wR_{[\![t]\!]}w$.

» $\mathfrak{M}, \alpha, w \Vdash Ft$ *iff* $[\![t]\!]_\alpha \in D$ *and it is not the case that* $wR_{[\![t]\!]}w$.

In other words, to check for truth and falsity of an object at a world, we need check only whether or not the accessibility relation for that object is reflexive at the world.

Proof: The fact for Dt is immediate. The only way for any statement of the form $t : A$ to be true at w is for $[\![t]\!]$ to be in D, by the satisfaction clause for $t : A$. Now consider whether or not Ft, given that t is declarative. The construction utilises the set $\overline{\{w\}}$ of worlds *other* than w. This is true at every world other than w. If $R_{[\![t]\!]}$ is not reflexive at w, then if $[\![p]\!]$ is $\overline{\{w\}}$, then at w, $t : p$ holds. However, we also have $\neg p$ at w, so at w, Ft. On the other hand, if $R_{[\![t]\!]}$ is reflexive at w, then the modal operator '$t :$' satisfies the characteristic condition for reflexivity: $t : p \supset p$ for every p. In other words, we have Tt. ∎

LEMMA [BRADWARDINE MODELS ARE APPROPRIATELY NAMED] *The Bradwardine axiom holds in every Bradwardine model.*

Proof: It is immediate, given the definition of $\Box$ as truth in all worlds. If all p worlds are q worlds then if all $R_{[\![x]\!]}$-alternatives of w are p worlds (if $x : p$ holds at w) then they are also q worlds ($x : q$ holds at w too). But w is arbitrary, so in *every* world, if $x : p$ then $x : q$, and so, we have $\Box(p \supset q) \supset \Box(x : p \supset x : q)$ as desired. ∎

EXAMPLE [A MODEL WITH A LIAR] Given this definition, we can use models to illuminate a number of issues. A simple model is one in which we have two worlds a and b, a single declarative object λ for which R_λ relates a to b and vice versa, but at neither a nor b is the relation R_λ reflexive. For simplicity, we shall suppose that our language has a name λ which denotes the object λ. (No confusion should arise, since it will be clear when the name is being used in the formal language as a part of a sentence being modelled, and when it is used in our language to pick out the declarative object in the domain.)

With this very small model, we can see that Tλ is false at both a and b, since R_λ is not reflexive at either world. So, it follows that since $\neg T\lambda$ is true everywhere, $\lambda : \neg T\lambda$ holds both at a and at b. In this model, λ is a liar sentence at both world a and at world b. However, it does not follow that $\lambda : T\lambda$ anywhere. $\lambda : T\lambda$ fails at a (and at b) since, the world a R_λ-accesses the world b, and at b, Tλ fails. Similarly, b R_λ-accesses a, and at a, Tλ fails. So, $\lambda : \neg T\lambda$ — it says of

itself that it is not true, but it is not the case that $\lambda : T\lambda$. It does not say of itself that it is true.

Returning to our question asked earlier concerning liar sentences: Since $T\lambda$ is false in this model (at each world) it follows that λ must say *something* that is not the case. But what? Here, the answer is curious. At world a, $\lambda : p$ where p is a proposition true at b, but not at a. The proposition p is not the case at a, yet it is at b. So at a, λ says something that is not the case: p. However, the matter is different at b: at b, p holds. However, at b, λ no longer says that p. At b, R_λ-accesses a, and not b, so at b, λ says that $\neg p$, and at b, *this* is something that is not the case. So in both worlds, there is something that λ says that is not the case. In one world it is p, and in the other it is $\neg p$. In neither world does $\lambda : T\lambda$. We can break the argument to the conclusion that the insolubilia allowing what an object says to vary from world to world.

As we have seen, Bradwardine models have enough structure to interpret "says that" in such a way as to satisfy the Bradwardine axiom. At the very least, can use them as a tool for finding counterexamples to certain claims, such as the claim that $\lambda : T\lambda$ follows from $\lambda : \neg T\lambda$ under the Bradwardine axiom. We also have enough resources to see that the Bradwardine axiom alone does not tell us a great deal about what is *true*.

DEFINITION [EXTREME MODELS] A model in which no object is declarative is said to be *reticent*. A model in which there are declarative objects, but the relations R_d for each such object is *empty*, is said to be *verbose*.

In every reticent model, there are no truths and no falsehoods. Nothing says a word, since nothing is declarative. At the other extreme, in the verbose model, there are declaratives, but every declarative says *everything*.

Extreme models do not tell us much interesting about truth. To begin to rule them out, we must specify what more is required in order to make a model interesting. Reflecting on Tarski's biconditionals governing truth, we have a connection between a sentence used (the A) and mentioned (the claim that A is true). In our models, we posit no such connection between the sentences we use (in the language $QS5^{\forall p}$ augmented with ':') and the objects of which we predicate truth and falsity — we have made no assumptions concerning the declarative objects. This is a fair general position to take: there is no requirement in providing a theory of truth for some domain that the language in which the theory is couched must also be among the objects of study of the theory. As a general account of truth, we should make as few assumptions as possible, if we wish the theory to apply as generally as possible.

However, if we are to compare Bradwardine's account of truth with accounts inspired by Tarski, we should at least *consider* it. So, for now, consider what happens if we add a device for quotation. For every sentence A in the language, we now supply a term $\ulcorner A \urcorner$, a quotation name for the sentence A itself. Since the

sentences in our language do not vary in truth from context to context (though they do vary from *world* to *world*) it results in no problems to consider either sentence types or canonical tokens of each sentence A in our language: for different tokenings of A do not vary in truth value in any given world.[7]

So, given that we include quotation names for terms in our language, what can we say concerning what these sentences say? There are many options, but a straightforward condition is that according to our theory, the sentences say what *we* take them to say when we are developing our theory. In other words the *transparency* axiom holds.

DEFINITION [TRANSPARENCY] The transparency axiom is this: $\ulcorner A \urcorner : A$.

The sentence $\ulcorner A \urcorner$ says that A holds. From the perspective of *inside* the theory, the sentence A says what we take it to say *outside* the theory.

A straightforward consequence of transparency is T-elimination.

FACT [T-ELIMINATION] *If* $\ulcorner A \urcorner : A$ *then* $T\ulcorner A \urcorner \supset A$.

Proof: If $T\ulcorner A \urcorner$ then $(\forall p)(\ulcorner A \urcorner : p \supset p)$. Substitute A for p. We have $\ulcorner A \urcorner : A$, so A follows. In other words, $T\ulcorner A \urcorner \supset A$. ∎

We have one half of Tarski's biconditional as a consequence of the definition of truth. It follows that if we have insolubilia, among the sentences of our own language, then the *other* half of Tarski's biconditional fails for them, lest the language fall into inconsistency.

Do we have insolubilia among the sentences of our language? Not as things stand: it all depends on the expressive power of that language. We have said nothing about the primitive non-logical symbols of the language, other than ':' and □. If, for example, we have the language of arithmetic, and we postulate properties of arithmetic sufficient to prove the diagonal lemma, then it follows that there are sentences L such that, according to the background theory, L is materially equivalent to $\neg T\ulcorner L \urcorner$. In fact, the arithmetic part of the theory is not only true but *necessary*, the diagonal lemma holds necessarily and we will have

$$\Box(L \equiv \neg T\ulcorner L \urcorner)$$

Given transparency, we have $\ulcorner L \urcorner : L$, and then the Bradwardine axiom tells us that $\ulcorner L \urcorner : \neg T\ulcorner L \urcorner$. A fixed point for the context $\neg Tx$ is a sentence of *our* language that *literally* says of itself that it is not true. Since $\ulcorner L \urcorner : \neg T\ulcorner L \urcorner$ we have $\neg T\ulcorner L \urcorner$ in the usual way, and hence, we prove L, since L is equivalent to $\neg T\ulcorner L \urcorner$. However, it does *not* follow that $T\ulcorner L \urcorner$, since we have not shown that *everything* that L says is true. So, in theories of arithmetic, we can find sentences which perforce must

---

[7]For many interesting utterances, of course, we will not be able to make this idealisation. To incorporate them into this framework, we will need to incorporate more contextual features into the points of evaluation. Here, we only have worlds and values for variables, to keep thing straightforward, and to facilitate comparison with traditional theories of truth and the paradoxes.

be counterexamples to T-introduction. A fixed point for $\neg Tx$, delivered by the diagonal lemma, is a sentence for which T-introduction fails. In fact, the failure of T-introduction for L is not without pain. For our theory not only proved that $\neg T\ulcorner L\urcorner$, since L is provably equivalent (relative to the arithmetic theory, at least) to $\neg T\ulcorner L\urcorner$, we have proved L. Our theory proves a sentence (namely, L) that is, by the lights of the theory itself, not true. This is not a happy feature of the account, but it must be admitted that *every* account of the paradoxes must bite some bullet, and for this formulation of Bradwardine's theory, this is the bullet that we must bite.

Do we have to bite any worse bullets? Is the theory consistent? It is consistent in the *weak* sense, since in any verbose model with names $\ulcorner A\urcorner$ for each sentence in the language, the transparency axiom holds: we have $\ulcorner A\urcorner : A$ since in a verbose model, we have $\ulcorner A\urcorner : B$ for every sentence B. The underlying theory is consistent, but only at the cost of trivialising ':' and truth (in the verbose model, everything declarative is false and nothing is true).

So, we want models that do better than this. We want models in which transparency holds, and in which a good fragment of the language behaves completely classically and straightforwardly. Let's start with transparency:

LEMMA [TRANSPARENCY IN BRADWARDINE MODELS] $\ulcorner A\urcorner : A$ *holds at a point w in a Bradwardine model if and only if for each v, if* $wR_{[\![\ulcorner A\urcorner]\!]}v$ *then* $v \Vdash A$.

It is not straightforward to check this condition in a model, as it is not a simple condition on the accessibility relation, but rather, something linking the accessibility relation for $[\![\ulcorner A\urcorner]\!]$ to what holds at the points accessible from w.

However, for this to hold, we do not need there to *be* any points accessible from w. The condition is automatically satisfied in the *verbose* Bradwardine model, in which all accessibility relations R_d are empty. Here, all statements of the form $t : A$ are true (for declarative t), and hence $\ulcorner A\urcorner : A$ is true. The move to Bradwardine models has done nothing to rule out the verbose model. We must do more to ensure that *some* of our declarative objects manage to be true. We have seen that T-introduction fails for insolubilia, but we have seen no reason to suspect that in the pre-semantic language, without ':', that T-introduction should fail. In fact, relative to some sense of what it is for a sentence to be 'grounded', we would like the following condition to be true:

DEFINITION [GROUNDED T-INTRODUCTION] $A \supset T\ulcorner A\urcorner$ for grounded sentences A.

Our aim is to find a sense of groundedness for which grounded T-introduction can hold, along with transparency. Then we will have recovered a significant portion of Tarski's theory, in the context of a very different theory.

The next lemma provides the conditions on which the transparency axiom holds at a point on in a model.

But how can we ensure that there are enough arcs in the accessibility relation for $\ulcorner A\urcorner$ to ensure that T-introduction holds for A?

The crucial idea is this: suppose we are constructing our model bit-by-bit. We have been provided with a $QS5^{\forall p}$ model. Our job is to define the relations $R_{[\![\ulcorner A \urcorner]\!]}$ in such a way as to verify transparency and grounded T-introduction. If A is a sentence not containing ':', then its status at a world w in the model is completely independent of the behaviour of the relations $R_{[\![\ulcorner A \urcorner]\!]}$ for the semantic conditions for A do not involve those relations.[8] If we already know that A holds at w, then there is no bar to us adding an arc $vR_{[\![\ulcorner A \urcorner]\!]}w$ to w from the world v. From the point of view of v, this does not threaten $\ulcorner A \urcorner : A$, since the world w now $R_{[\![\ulcorner A \urcorner]\!]}$-accessible from v is a world at which A holds. On the other hand, if A is free of ':', then we know its status at every world, independently of what we do with the relations $R_{[\![\ulcorner B \urcorner]\!]}$, so its status is fixed at *every* world. So, add the arcs $vR_{[\![\ulcorner A \urcorner]\!]}w$ to w if and only if A holds at v. Now consider the claim $A \supset T\ulcorner A \urcorner$. Can it ever fail? Only if we have a point at which A holds and $T\ulcorner A \urcorner$ does not. But $T\ulcorner A \urcorner$ fails at a point w only when $wR_{[\![\ulcorner A \urcorner]\!]}w$ fails. But by hypothesis, if A holds at a point, we added all arcs to that point, including the arc to w from w. So, $A \supset T\ulcorner A \urcorner$ holds at every world.

This construction worked for A only because we assumed that the status of A at w did not depend on the features of any of the accessibility relations at w or elsewhere. This is the feature we wish from grounded sentences. Consider the case of an ungrounded sentence, such as L, for which it is true at w that $\ulcorner L \urcorner : \neg T\ulcorner L \urcorner$. L holds at w if we do not have $wR_{[\![\ulcorner L \urcorner]\!]}w$. We cannot reason as we did before, and say that since L holds at w, we can now add the arc $wR_{[\![\ulcorner L \urcorner]\!]}w$, and keep the transparency axiom, since adding this arc has changed the status of L at w itself. Now, L fails at w, since $T\ulcorner L \urcorner$ is true at w, and $\ulcorner L \urcorner : \neg T\ulcorner L \urcorner$. This is an ungrounded sentence: it is not *settled*. Its status at a world depends on further choices for arcs in the accessibility relation.

This is the kernel of the idea behind the limit construction. We start with a blank slate in which the relations $R_{[\![\ulcorner A \urcorner]\!]}$ are empty. We add arcs for each $R_{[\![\ulcorner A \urcorner]\!]}$ when we know that the target world has settled on A. Then repeat. Stop when you're done. Now to the details.

3   THE LIMIT CONSTRUCTION

Our starting point is a particular $QS5^{\forall p}$ model $\mathfrak{M}$, with a denotation in the domain of $\mathfrak{M}$ for each term $\ulcorner A \urcorner$ or a name in the language. It could be an extensional, classical model for arithmetic. Having more than one world is not

---

[8]It is crucial here that we understand propositional quantification not by substituting new *formulas* in the place of propositional variables, but by allowing the denotation of the variable to be given by a set of worlds. If we verified $(\forall p)(\cdots p \cdots)$ by checking every formula $(\cdots B \cdots)$, then not only would this process not terminate, but we would be forced to evaluate formulas involving ':'. Here, we simply check $(\cdots p \cdots)$ for each possible *value* of p, and a value of p is a set of points in the model, which can be checked independently of the status of the relations $R_{[\![\ulcorner A \urcorner]\!]}$. We do know, however, that we have checked the status of the formula $(\cdots B \cdots)$ where B is a formula involving ':' that B is true at some set of worlds, and we have checked this set in our evaluation. We just did not check it *under that description*.

essential to the construction. Our task is to define the $R_{[\![A]\!]}$ for each A in order to satisfy

TRANSPARENCY: $wR_{[\![A]\!]}v$ only when A is true at v.

GROUNDED T-INTRO: For *grounded* A, if A holds at w, then $wR_{[\![A]\!]}w$.

As we have seen, transparency is satisfied when $R_{[\![A]\!]}$ is empty. We will start with empty accessibility relations, add links in the accessibility graphs conservatively, in the manner we have described.

DEFINITION [THE INITIAL BRADWARDINE MODEL FOR $\mathfrak{M}$] Given a $QS5^{\forall P}$ model $\mathfrak{M}$, with a denotation in the domain of $\mathfrak{M}$ for each term $\ulcorner A \urcorner$ or a name in the language, the initial Bradwardine model for $\mathfrak{M}$ is the Bradwardine model whose declarative sentences are the collection objects in O identical to $[\![\ulcorner A \urcorner]\!]$ for some sentence A, and whose relations $R_{[\![\ulcorner A \urcorner]\!]}$ are empty. This is $Br_0(\mathfrak{M})$

We will not stay at level zero for long.

DEFINITION [DEVELOPMENT] A *development* of a model $\mathfrak{N}$ is any model $\mathfrak{N}'$ in which the relations $R_{[\![\ulcorner A \urcorner]\!]}$ are replaced by $R'_{[\![\ulcorner A \urcorner]\!]} \supseteq R_{[\![\ulcorner A \urcorner]\!]}$.

DEFINITION [SETTLED SENTENCES] A is said to be *settled at w in $\mathfrak{N}$* if and only if $(\mathfrak{N}, \alpha, w \Vdash A$ iff $\mathfrak{N}', \alpha, w \Vdash A)$ for each development $\mathfrak{N}'$ of $\mathfrak{N}$.[9]

FACT [THE ORIGINAL LANGUAGE IS SETTLED] *Each sentence in the base language — each sentence in the language not containing ':' — is settled everywhere in $\mathfrak{M}$.*

DEFINITION [SAFETY] An arc $wR_{[\![\ulcorner A \urcorner]\!]}v$ is *safe* in $\mathfrak{N}$ iff $\mathfrak{N}', \alpha, v \Vdash A$ in every development $\mathfrak{N}'$. An entire model $\mathfrak{N}$ is *safe* iff it every arc in $\mathfrak{N}$ is safe.

Safe models do not only validate transparency, they make it easy to find developments of safe models which validate transparency — you need only check transparency for *new* arcs, since if an arc is safe in a model $\mathfrak{N}$, it is also safe in every development of $\mathfrak{N}$.

DEFINITION [THE JUMP] Given a model $\mathfrak{N}$ with relations $R_{[\![\ulcorner A \urcorner]\!]}$, define $R'_{[\![\ulcorner A \urcorner]\!]}$ by setting $wR'_{[\![\ulcorner A \urcorner]\!]}v$ iff either $wR_{[\![\ulcorner A \urcorner]\!]}v$ or $\mathfrak{N}, \alpha, v \Vdash A$ and A is *settled* at v in $\mathfrak{N}$.

The development of $\mathfrak{N}$ with accessibility relations $R'_{[\![\ulcorner A \urcorner]\!]}$ is the said to be the JUMP $\mathfrak{N}'$ of $\mathfrak{N}$.

Notice that he jump opertion merely *adds* arcs to an accessibility relation. It never deletes them. The result means that the relation is *mononotone* with respect to the natural ordering on models induced by the subset order on each accessibility relation.

---

[9]This may remind the reader of the notion of super-truth in supervaluational semantics.

DEFINITION [SUCCESSOR BRADWARDINE MODELS] For each α, the successor Bradwardine model at stage $\alpha + 1$ is $Br_{\alpha+1}(\mathfrak{M}) = Br_\alpha(\mathfrak{M})'$.

FACT [SAFETY AND JUMPS] *If $\mathfrak{N}$ is safe, then so is its jump $\mathfrak{N}'$.*

Proof: Every arc in $\mathfrak{N}'$ comes either from $\mathfrak{N}$ (in which case it is safe in $\mathfrak{N}'$, as any development of $\mathfrak{N}'$ is a development of $\mathfrak{N}$, and it is safe in $\mathfrak{N}$), or it is safe in $\mathfrak{N}'$ by construction. ∎

Now, for any model $\mathfrak{M}$, its initial Bradwardine model $Br_0(\mathfrak{M})$ is safe, since each R_d is empty. So, each finite jump $Br_n(\mathfrak{M})$ is also safe. It remains only to extend the process beyond an ω sequence.

DEFINITION [LIMITS] Given a sequence $\mathfrak{M}_\lambda$ ($\lambda < \kappa$) of safe models such that $\mathfrak{M}_\lambda$ is a development of $\mathfrak{M}_\epsilon$ whenever $\lambda > \epsilon$, then the LIMIT MODEL $\mathfrak{M}_\kappa$ of the construction is defined by setting $R_{[\![\ulcorner A \urcorner]\!]\kappa} = \bigcup_{\lambda < \kappa} R_{[\![\ulcorner A \urcorner]\!]\lambda}$.

DEFINITION [THE BRADWARDINE HIERARCHY] For each ordinal α, we define the Bradwardine model $Br_\alpha(\mathfrak{M})$ as follows:

» $Br_0(\mathfrak{M})$ is $\mathfrak{M}$ with the addition of empty relations R_d for each declarative object.

» $Br_{\alpha+1}(\mathfrak{M})$ is the jump of $Br_\alpha(\mathfrak{M})$

» $Br_\lambda(\mathfrak{M})$, where λ is a limit ordinal, is the limit of the series $Br_\alpha(\mathfrak{M})$ for $\alpha < \lambda$.

In each case, $Br_\alpha(\mathfrak{M})$ differs from $Br_\beta(\mathfrak{M})$ when $\alpha > \beta$ only by having $R_d^\alpha \supseteq R_d^\beta$. In all other respects, the models are identical. As we go up the hierarchy, each sentence A says less and less at each world, for the number of $R_{[\![\ulcorner A \urcorner]\!]}$ alternatives to each world increases. Up to a point.

FACT [EVERY BRADWARDINE MODEL IS SAFE] *Each $Br_\alpha(\mathfrak{M})$ is safe.*

Proof: $Br_0(\mathfrak{M})$ is safe by definition, and its jumps are safe, by the fact we proved previously. The limit of a series of safe models is safe, since the limit model contains no arcs not contained in a model in the series. ∎

We are now a short step to our desired model.

DEFINITION [THE FIXED POINT] The fixed point Bradwardine model $Br_*(\mathfrak{M})$ is the first model $Br_\alpha(\mathfrak{M})$ that is identical to its own jump.

FACT [WE HAVE FIXED POINTS] *For any model $\mathfrak{M}$, its fixed point Bradwardine model $Br_*(\mathfrak{M})$ is defined.*

Proof: This is a standard result in limit constructions of this kind. If each jump (and limit stage) is distinct, an new arc is added at each jump. However, given an starting model, with at most κ worlds (for some cardinality κ), there are only $\kappa \times \kappa \leqslant \max(\aleph_0, \kappa)$ possible choices to be made for adding arcs in the model. So, by the time we reach a stage of cardinality *greater* than $\kappa \times \kappa$, we have reached our fixed point. ∎

Now we may investigate the properties of the fixed point model.

LEMMA [TRANSPARENCY AT THE LIMIT] *In the limit model* $\mathrm{Br}_*(\mathfrak{M})$, $\ulcorner A \urcorner : A$ *holds at every point.*

Proof: At each stage, we add an arc $w\mathrm{R}_{\llbracket \ulcorner A \urcorner \rrbracket}v$ only when it is safe to do so. In every development of the model we have A true at v. But this is all that is required to ensure that $\ulcorner A \urcorner : A$ remains true in the model. ∎

The more subtle matter is to ensure that T-introduction holds for a suitable class of sentences.

FACT [SETTLEDNESS IN A SEQUENCE] *Given a sequence* $\mathrm{Br}_\alpha(\mathfrak{M})$, *a formula (possibly containing free variables), we can keep track of when (if ever) that formula is* settled *in the sequence, relative to an assignment of values to variables α in the following way. A formula is said to be "settled at stage λ, relative to assignment α" if and only if the status of A is fixed at each world* $\mathrm{Br}_\alpha(\mathfrak{M})$, *relative to the assignment α, over every development of that model. We have:*

» *If A is ':'-free, it is settled at stage 0, for each assignment α.*

» $D\ulcorner A \urcorner$ *is settled at stage 0, for each assignment α.*

» *Boolean combinations of sentences A_i settled at level λ_i, relative to assignment α are settled at $\max_i \lambda$, relative to that assignment.*

» *If A is settled at λ, given assignment α and B is settled at γ, given α, then $\ulcorner A \urcorner : B$ is settled at $\max(\alpha, \beta)$, relative to α.*

» *If A is settled at λ_d for the x-variant α_d of α (which differs from α only by assigning the variable x the value d), then $(\forall x)A$ is settled at $\sup\{\alpha_d : d \in D\}$, relative to the assignment α.*

» *If A is settled at λ_d for the p-variant α_X of α (which differs from α only by assigning the variable p the value $X \subseteq W$), then $(\forall p)A$ is settled at $\sup\{\alpha_X : X \subseteq W\}$, relative to the assignment α.*

» *If A is settled at λ (relative to α), $T\ulcorner A \urcorner$ and $F\ulcorner A \urcorner$ are settled at $\lambda + 1$ (relative to α).*

Proof: Each case is quite straightforward. For the first case, the value of A is fixed independently of the accessibility relations for ':', so it is settled at stage zero. So is whether or not something is declarative, since this is determined by the set D, which is fixed at the outset.

   If a collection of formulas is settled at some stage, so is any Boolean combination of that set. For quantifiers, it suffices to note that if we have a model in which A is settled for each value the variable x may take, then since the value of $(\forall x)A$ depends merely the values that A can take as the denotation of x varies. If each of these instances are fixed, then so is $(\forall x)A$. The same holds for propositional quantification.

   The case of truth is straightforward. Once A is settled at stage λ, then at stage $\lambda + 1$ we add the appropriate arcs in $R_{[\![\ulcorner A \urcorner]\!]}$ to each A-world. No more $R_{[\![\ulcorner A \urcorner]\!]}$ arcs are to be added, since A is settled. This settles the points at which $T\ulcorner A \urcorner$ and $F\ulcorner A \urcorner$ are true, since this is evaluated merely by inspecting the arcs for $R_{[\![\ulcorner A \urcorner]\!]}$. ∎

With this at hand, we can prove the following corollary.

COROLLARY [T-INTRO FOR GROUNDED SENTENCES] *If* A *is settled at* some *stage* λ, *then in* $Br_*(\mathfrak{M})$, *at all points*, $A \supset T\ulcorner A \urcorner$.

This provides us with a large class of formulas for which T-intro holds in a model: those settled at some stage. If we take these to be the *grounded* sentences, we have a large domain of formulas for which T-biconditionals hold.

4   DISCUSSION

There are many questions which we might consider for future examination.

1. What *holds* in each $Br_*(\mathfrak{M})$ above and beyond what follows from Bradwardine's axiom, transparency and grounded T-introduction?

    One example of a further condition that holds is $(\forall x)(Dx \supset x : Tx)$: which we could call *Buridan's axiom* — all declaratives declare their own truth. This holds because if we add an arc for $R_{[\![x]\!]}$ from *w* to *v*, then we also added the arc from *v* to *v*. From *w* we only x-access worlds that take x to be true.

2. This raises another question: Buridan's axiom is controversial. (Bradwardine did not take it as an axiomatic principle of truth.) Is it defensible on some considerations? Are any of the extra claims that hold in our fixed-point models defensible? Or should we look for model constructions that refute questionable claims?

3. Further to the issue of what other principles hold in our models, we can ask whether any further principles concerning truth and signification have

independent merit, and of those that have merit, do they hold in fixed-point models? We have seen that T-introduction cannot be added at the point of paradox, but that T-introduction for grounded sentences can. We have used the concept of groundedness in our meta-language. Can we introduce this to the object language without paradox?

4. Other issues arise concerning the logic of fixed point models: is there an axiomatisation of the logic of a fixed-point?

5. What would be the strength of theories in this logic?

6. How does the non-modal version of the construction relate to known theories of truth?

7. What is the behaviour of *other* ungrounded sentences in the models?

8. What about other logics (ternary frames for relevant logics, etc.)?

9. Notice that we have a different accessibility relation for each declarative object. In our case, declarative objects are *sentences*. Do sentences with the same *meaning* (in some sense or other) have the same signification? In our fixed point models it does not hold that any two sentences that are true at the same set of worlds would have the same signification. Consider the one world model and a liar sentence L where we have $L \equiv \neg T\ulcorner L\urcorner$, and hence $\ulcorner L\urcorner : \neg T\ulcorner L\urcorner$. In our model, the sentence L holds (at the only world), since $\neg T\ulcorner L\urcorner$. Take a tautology $\top$. This is equivalent to L in our model: we have $L \equiv \top$. However, the accessibility relation for $\ulcorner L\urcorner$ differs from that for $\ulcorner \top\urcorner$, since we have $T\ulcorner \top\urcorner$ and we have $\neg T\ulcorner L\urcorner$. Equivalence of sentences in our model (in the sense of having the same interpretation in the set of worlds) does not suffice to force equivalence of the *signification* of those sentences, and it cannot.

   However, the fixed point construction is enough to ensure that if sentences A and B have the same interpretation in *each stage* of the construction, then the accessibility relation for $\ulcorner A\urcorner$ will equal that for $\ulcorner B\urcorner$. Is there a way of independently charactersing this kind of equivalence between declarative objects?[10]

These are more than enough questions to keep us going for the next few years. Hopefully the answers will help us relate the strength of theories of truth such as Bradwardine's with other logics that are better understood.

REFERENCES

[1] ROBERT B. BRANDOM. *Making It Explicit*. Harvard University Press, 1994.
[2] J. ALBERTO COFFA. *The Semantic Tradition from Kant to Carnap*. Cambridge University Press, 1993. Edited by Linda Wessels.

---

[10]Thanks to the referee for drawing this issue to my attention.

[3] KIT FINE. "Semantics for Quantified Relevance Logic". *Journal of Philosophical Logic*, 17(1):27–59, 1988.

[4] ROBERT GOLDBLATT AND EDWIN D. MARES. "An Alternative Semantics for Quantified Relevant Logic". *Journal of Symbolic Logic*, 71(1):163–187, 2006.

[5] DOROTHY GROVER. "Propositional quantifiers". *Journal of Philosophical Logic*, 1(2):111–136, 1972.

[6] DOROTHY GROVER. *A Prosentential Theory of Truth*. Princeton University Press, 1992.

[7] DOROTHY L. GROVER, JOSEPH L. CAMP, AND NUEL D. BELNAP. "A prosentential theory of truth". *Philosophical Studies*, 27(2):73–125, 1975.

[8] PAUL HORWICH. *Truth*. Basil Blackwell, Oxford, 1990.

[9] HANNES LEITGEB. "Truth and the Liar in De Morgan-Valued Models". *Notre Dame Journal of Formal Logic*, 40(4):496–514, 1999.

[10] HANNES LEITGEB. "What Truth Depends On". *Journal of Philosophical Logic*, 34(2):155–192, 2005.

[11] GRAHAM PRIEST. "The Logic of Paradox". *Journal of Philosophical Logic*, 8(1):219–241, 1979.

[12] STEPHEN READ. *Relevant logic: a philosophical examination of inference*. Basil Blackwell, Oxford, 1988.

[13] STEPHEN READ. "The Liar Paradox from John Buridan back to Thomas Bradwardine". *Vivarium*, 40(2):189–218, 2002.

[14] STEPHEN READ. "Symmetry and Paradox". *History and Philosophy of Logic*, 27:307–318, 2006.

[15] GREG RESTALL. "Deviant Logic and the Paradoxes of Self Reference". *Philosophical Studies*, 70(3):279–303, 1993.