Draft of November 29, 2016. Comments welcome!

Quantification and Epistemic Modality

Dilip Ninan
dilip.ninan@tufts.edu

Abstract: This paper presents a novel puzzle concerning the interaction of epistemic modals, singular terms, and quantifiers. The puzzle poses a number of problems for both *relational* and *dynamic* theories of epistemic modals. The trouble arises from the fact that both of these approaches treat being possibly thus-and-so (in the epistemic sense of *possibly*) as a trait that an object has independently of how it is thought of. But, as the puzzle reveals, this is not so: being possibly thus-and-so is a trait that an object has only relative to a 'way of thinking' of the object (cf. Quine 1953). We then formulate two theories that can incorporate this observation and thereby avoid the trouble. The first is a dynamic semantics that embeds a system of *contingent identity*; the second is a *counterpart-theoretic* version of relational semantics. We close by considering how we might decide between the resulting two approaches.

1 Introduction

This essay concerns the interaction of quantifiers, singular terms, and epistemic modals. I begin with a puzzle.

Imagine that there is a lottery with only two tickets, a blue ticket and a red ticket. The tickets are also numbered 1 through 2, but we don't know which color goes with which number. (Perhaps the number of the ticket is only printed on the front, and we can only see the colored back of each ticket.) The winner (there is only one) has been drawn, and we know that the blue ticket won. But since we don't know whether the blue ticket is ticket #1 or ticket #2, we don't know the number of the winning ticket.

We now reason as follows (in what follows, *might* is to be read epistemically):

(1) Ticket #1 is such that it might be the winning ticket. $(\lambda x.\Diamond x = w)(t_1)$

(2) Ticket #2 is such that it might be the winning ticket. $(\lambda x.\Diamond x = w)(t_2)$

(3) Those are all the tickets. $\forall x(x = t_1 \vee x = t_2)$

So:

(4) Any ticket might be the winning ticket. $\forall x \Diamond x = w$

But of course the red ticket is a ticket. Given this, it would appear to follow from (4) that:

(5) The red ticket is such that it might be the winning ticket. $(\lambda x. \Diamond x = w)(r)$

But isn't (5) false? After all, we know that the blue ticket is the winning ticket, and we can see that the blue ticket is not the red ticket. Given this, it seems like I can point at the red ticket and truly say, "That ticket is not the winner." The fact that I am in a position to say that seems to be at odds with the truth of (5). But how could (5) be false? It follows from (4), which, in turn, follows from (1)-(3), all of which appear to be true.[1]

Our puzzle is that from the apparently true (1)-(3) we can infer the apparently false (5). Note that the problem appears to arise in connection with epistemic modals *per se*. To see this, consider what happens when we re-formulate (1)-(3) replacing the epistemic *might* with the metaphysical *could have been* throughout. From the resulting claims we can infer (5′) rather than (5):

(5′) The red ticket is such that it could have been the winning ticket.

But (5′) appears to be unproblematically true. In the epistemic case, the fact that we know that the red ticket lost makes it difficult to accept that it might be the winning ticket. But in the metaphysical case, the fact that we know that the red ticket lost is no barrier to accepting the claim that it *could have been* the winning ticket.

We begin by considering what the orthodox theory of epistemic modals – *relational semantics* – says about the puzzle.

2 Relational semantics

According to the standard relational semantics for modal operators, a sentence of the form $\Diamond\phi$ is true at a possible world v just in case there is a possible world accessible from v at which ϕ is true. Where the epistemic interpretation of $\Diamond$ is at issue, 'accessibility' is 'epistemic accessibility'. Given an agent or group of agents x, we can say that possible world v' is epistemically accessible from v relative to x just in case v' is compatible with what x knows in v.[2]

On the standard 'contextualist' version of this story, when an epistemic modal is used on a particular occasion, the context of utterance determines the agent or group of agents whose knowledge is used in defining the accessibility relation relevant for interpreting the modal (as used on that occasion).[3] But this

[1] Aloni (2001, Ch. 3) and Moss (2016, §8.3) discuss similar examples.

[2] It might be better to represent epistemic possibilities using *centered worlds* (Lewis 1979, 1983), rather than (uncentered) possible worlds, but I do not think anything in what follows turns on this choice.

[3] On one influential version of this view, namely Kratzer's, this is simply an instance of the general role context plays in determining (something like) an accessibility relation. See Kratzer (1981, 1991, 2012).

version of the story has come under fire in recent years. The main controversy concerns whether the context of use is all that is needed for determining the accessibility relation. *Relativists* argue that cases in which a sentence containing an epistemic modal is used by one agent (or group) and then assessed by another agent (or group) require us to reject the idea that a single context – the context of utterance – determines an epistemic accessibility relation. Rather, such a relation is only determined given both a context of utterance and a context of *assessment*.[4]

This controversy has little bearing on our puzzle. For the puzzle was stated by considering the lottery scenario from the point of view of a group of agents who are both using and then assessing the claims in question – namely us. Thus, we are contemplating the acceptability of sentences featuring epistemic modals as they are being used and assessed by us, *qua* agents in the fictional lottery scenario. Thus, it would seem that, on anyone's view, we are the agents whose knowledge is relevant for determining the accessibility relation.

Furthermore, since it is no part of the scenario that what we know changes over the course of the story, the case can modeled using a standard Kripke model in which a single accessibility relation is specified as an element of the model, rather than, say, using a model that contains a set of contexts, each of which determines its own accessibility relation.

I propose to set out the relational semantics in some detail, even though this approach may be familiar to the reader. Some of this setup will be useful later when we turn to consider semantic theories that may be less familiar. I will proceed by informally translating the relevant English sentences into a formal language, and then considering candidate semantic theories for the formal language. But I emphasize that our interest here is, ultimately, with the semantics of the natural language expressions involved.

So we can begin by defining a formal language $\mathcal{L}$, the language of quantified modal logic with individual constants, identity, and an abstraction operator.

Definition 2.1. We start by specifying the *vocabulary of* $\mathcal{L}$, which consists of:

individual constants $a_1, a_2, a_3, ...$;

variables $x_1, x_2, x_3, ...$;

for each integer $n \geq 0$, predicates $P_1^n, P_2^n, P_3^n, ...$, including a distinguished binary predicate $=$;

the logical symbols $\neg, \wedge, \lambda, \exists, \Diamond, (,)$.

We use a, b, c, etc. to denote distinct but arbitrary individual constants; x, y, z, etc., for variables; and P, Q, R, etc. for predicates.

Definition 2.2. Individual constants and variables are the *terms of* $\mathcal{L}$.

[4]These issues are discussed in Egan *et al.* (2005), Egan (2007), Stephenson (2007), von Fintel and Gillies (2008), Dowell (2011), von Fintel and Gillies (2011), MacFarlane (2011), Schaffer (2011), Yalcin (2011), and MacFarlane (2014), among others. Related issues are discussed in Hacking (1967) and DeRose (1991).

Definition 2.3. An *atomic formula of* $\mathcal{L}$ consists of an n-ary predicate P followed by n terms.[5]

Definition 2.4. We then recursively define the notion of a *formula of* $\mathcal{L}$ as follows:

> An atomic formula is a formula of $\mathcal{L}$.
>
> If ϕ and ψ are formulas of $\mathcal{L}$, x a variable of $\mathcal{L}$, and t a term of $\mathcal{L}$, then the following are also formulas of $\mathcal{L}$: $\neg\phi$, $(\phi \wedge \psi)$, $(\lambda x.\phi)(t), \exists x\phi, \Diamond\phi$.
>
> Nothing else is a formula of $\mathcal{L}$.

The other familiar logical symbols are treated as metalinguistic abbreviations in the standard way. For example, $\forall x\phi$ is $\neg\exists x\neg\phi$ and $\Box\phi$ is $\neg\Diamond\neg\phi$. All the semantic theories we shall consider will interpret this same language $\mathcal{L}$.

Definition 2.5. The relational semantics starts with the notion of a *Kripke model* for $\mathcal{L}$, which is an n-tuple $\mathcal{M} = \langle \mathcal{W}, \mathcal{R}, \mathcal{D}, \mathcal{I}, v \rangle$ where:

> $\mathcal{W}$ is a non-empty set (whose elements are called *worlds*),[6]
>
> $\mathcal{R}$ is a binary relation (an *accessibility relation*) on $\mathcal{W}$,
>
> $\mathcal{D}$ is a non-empty set (called a *domain* whose elements are *individuals*),
>
> $\mathcal{I}$ is a function that:
>
> (i) maps each individual constant to an *individual concept*, i.e. a total function from $\mathcal{W}$ into $\mathcal{D}$;
>
> (ii) maps each n-ary predicate P (other than the identity sign) to a total function from $\mathcal{W}$ to a subset of $\mathcal{D}^n$, the set of n-ary sequences whose elements are members of $\mathcal{D}$; and
>
> (iii) maps the identity sign "=" to the constant function mapping each world v to $\{\langle o, o\rangle : o \in \mathcal{D}\}$, and
>
> v is an element of $\mathcal{W}$ (the *actual world*).

Since the relevant interpretation of the modal operators is epistemic, and since what is known is true, we confine our discussion to *reflexive* Kripke models, i.e. models whose accessibility relation is reflexive. For simplicity, our models are *constant domain* models, which means that no individuals come into or go out of existence as we move from world to world. World v is intended to represent the actual world; this element plays no role in the semantics proper, but is useful when thinking about how a model might represent a certain situation, such as the lottery scenario.

[5]Though in the case of the identity predicate, we write "$t = t'$" rather than "$= (t, t')$".

[6]Note that in a Kripke model $\mathcal{M}$, $\mathcal{W}$ is simply a non-empty set, so the 'worlds' in $\mathcal{W}$ could be anything: they *could* be possible worlds in the metaphysician's sense, but they could also be pure sets, or integers, or apples. But since our interest is with models in which $\mathcal{W}$ is a set of possible worlds in (roughly) the metaphysician's sense, I will not be careful in distinguishing the two senses of 'worlds'.

Definition 2.6. A *variable assignment* on a model $\mathcal{M}$ is a function mapping each variable x to an individual in $\mathcal{D}$. Given a variable assignment g and an object $o \in \mathcal{D}$, $g[x/o]$ is the variable assignment that is like g with the possible exception that $g[x/o](x) = o$.

Definition 2.7. *The extension of a term t at a model $\mathcal{M}$, world v and a variable assignment g* is written as $[\![t]\!]^{v,g}$, and defined as follows:[7]

$[\![t]\!]^{v,g} = g(t)$ if t is a variable,

$[\![t]\!]^{v,g} = \mathcal{I}(t)(v)$ if t is a constant.

Definition 2.8. The core of the semantics is the recursive definition of *the truth of a formula ϕ at a model $\mathcal{M}$, a world v, and a variable assignment g.* We write this as $[\![\phi]\!]^{v,g} = 1$, and define it as follows:

$[\![P(t_1, ..., t_n)]\!]^{v,g} = 1$ iff $\langle [\![t_1]\!]^{v,g}, ..., [\![t_n]\!]^{v,g} \rangle \in \mathcal{I}(P)(v)$[8]

$[\![\neg\phi]\!]^{v,g} = 1$ iff $[\![\phi]\!]^{v,g} = 0$

$[\![\phi \wedge \psi]\!]^{v,g} = 1$ iff $[\![\phi]\!]^{v,g} = 1$ and $[\![\psi]\!]^{v,g} = 1$

$[\![(\lambda x.\phi)(t)]\!]^{v,g} = 1$ iff $[\![\phi]\!]^{v,g[x/o]} = 1$, where $o = [\![t]\!]^{v,g}$

$[\![\exists x\phi]\!]^{v,g} = 1$ iff there is an individual $o \in \mathcal{D}$ such that $[\![\phi]\!]^{v,g[x/o]} = 1$

$[\![\Diamond\phi]\!]^{v,g} = 1$ iff there is a world $v' \in \mathcal{W}$ such that $v\mathcal{R}v'$ and $[\![\phi]\!]^{v',g} = 1$

We can use our triply-relative notion of truth to define a doubly-relative one in the usual way:

Definition 2.9. A formula ϕ is *true* at a model $\mathcal{M}$ and a world v of $\mathcal{M}$ iff $[\![\phi]\!]^{v,g} = 1$, for all assignments g.

Definition 2.10. An argument is *valid* just in case: for any model and world, if the premises are true at that model and world, so is the conclusion; in that case, the premises *entail* the conclusion.

As the reader will no doubt have noticed, I have already been providing translations of our target English sentences into $\mathcal{L}$. And in doing so, I have been following Frege (1892) in treating definite descriptions as terms, translating descriptions like *the red ticket* via individual constants like r. This is not an uncontroversial choice, since some philosophers may prefer to translate sentences containing such expressions in the manner of Russell (1905). However, I suspect that most of the points made in this paper would survive (in perhaps a slightly altered form) were we to have followed Russell here rather than Frege.

[7]Strictly speaking, we should write $[\![t]\!]^{\mathcal{M},v,g}$, but we suppress reference to the model for simplicity.

[8]Here I used t_1 and t_2, etc. as variables of the metalanguage ranging over terms of the object language. But I have also been using t_1 and t_2 as individual constants of the object language, as the translations of *ticket #1* and *ticket #2*. I will continue to use these in an equivocal way, leaving context to disambiguate.

It is also worth nothing that the problem raised by the lottery puzzle does not appear to turn essentially on the use of definite descriptions. For example, suppose we are tracking two criminals, Al and Bill. We know that they go by the nicknames *Big Boy* and *Fox*, but we don't know which Christian name goes with which nickname. We know that either Al or Bill (but not both) murdered the mayor, but we don't know which of them did it. We do, however, know that the murderer was Fox. This scenario will give rise to a puzzle that is structurally similar to our lottery puzzle. Thus, we do not want our solution to *depend* on a Russellian treatment of definite descriptions, since we should like that solution generalize to cases involving proper names.[9]

What does this semantic theory say about our puzzle? Let's start with this observation:

Fact 2.1. *The argument from (1)-(3) to (4) and the argument from (4) to (5) are both valid on this semantics.*

Proof. To see the validity of the inference from (1)-(3) to (4), let $\mathcal{M}$ be any Kripke model, let v be any world in $\mathcal{M}$, and let g be a variable assignment on $\mathcal{M}$. Suppose that (1)-(3) are all true at v and g. Given the definition of the universal quantifier, sentence (4) is true at v and g just in case every object o in the domain is such that $[\![\Diamond x = w]\!]^{v,g[x/o]} = 1$. Let o be any object in the domain. By the truth of (3) at v and g, it follows that either (i) o is identical to the extension of t_1 at v, or (ii) o is identical to the extension of t_2 at v. In the first case, we have $[\![\Diamond x = w]\!]^{v,g[x/o]} = 1$, given the truth of (1) at v and g, and given the clause for the abstraction operator. In the second case, we again have $[\![\Diamond x = w]\!]^{v,g[x/o]} = 1$, given the truth of (2) at v and g, and given the clause for the abstraction operator. Either way, $[\![\Diamond x = w]\!]^{v,g[x/o]} = 1$. Since o was an arbitrary object, this holds for every object in the domain, which means (4) is true at v and g.

To see that (4) entails (5), let $\mathcal{M}$ be any Kripke model, let v be any world in $\mathcal{M}$, and let g be a variable assignment on $\mathcal{M}$. Suppose that (4) is true at v and g. Then from the semantics for the universal quantifier, it follows that $[\![\Diamond x = w]\!]^{v,g[x/o]} = 1$ for every object o in the domain. Let o' be the extension of r at v. Then $[\![\Diamond x = w]\!]^{v,g[x/o']} = 1$. From this it follows that (5) is true at v and g, given the semantics for the abstraction operator. $\square$

It is worth noting that the validity of these inferences depends on the fact that the modal predications (1), (2), and (5) are *de re*. Neither inference remains valid if we replace these *de re* modal predications with their *de dicto* counterparts.[10] This explains why both inferences were formulated using the somewhat cumbersome *is such that* locution, which ensures that the relevant

[9] I'm assuming that proper names are terms, not 'disguised descriptions' that receive a Russellian analysis.

[10] A *de re* modal predication is a sentence of the form $(\lambda x.\Diamond\phi)(t)$, and the *de dicto* counterpart of such a sentence is a sentence of the form $\Diamond\phi(t/x)$, where t is free for x in ϕ. A term t is *free for* (or *substitutable for*) variable x in formula ϕ just in case no occurrence of x that is free in ϕ is in the scope of a t-quantifier or a t-abstraction operator (so no free occurrence of x in ϕ becomes a bound occurrence of t in $\phi(t/x)$).

singular terms take wide-scope over the modal operator.[11] Consider, for example, the '*de dicto* counterpart' of the inference from (4) to (5):

(4) Every ticket is such that it might be the winning ticket. $\forall x \Diamond x = w$

So:

(5*) It might be the case that the red ticket is the winning ticket. $\Diamond r = w$

Fact 2.2. *The argument from (4) to (5*) is invalid on the relational semantics.*

Proof. To see this, let $\mathcal{M} = \langle \mathcal{W}, \mathcal{R}, \mathcal{D}, \mathcal{I}, v_i \rangle$, where $\mathcal{W} = \{v_i, v_j\}$, $\mathcal{D} = \{\epsilon, \zeta\}$, $\mathcal{R}$ is the universal relation on $\mathcal{W}$, and the relevant facts about $\mathcal{I}$ can be pictured as follows:

v_i	v_j
ϵ: blue, winner	ϵ: red, loser
ζ: red, loser	ζ: blue, winner

Sentence (4) is true at v_i iff every object o in the domain is such that there a world v such that $v_i \mathcal{R} v$ and o is the winning ticket at v. This condition is satisfied. Note that there are only two objects in the domain, ϵ and ζ, and that ϵ is the winning ticket in v_i and ζ is the winning ticket in v_j. Since $\mathcal{R}$ is the universal relation on $\mathcal{W}$, this means that (4) is true at v_i. But since there is no world v in the model at which the ticket that is colored red in v wins, (5*) is false at v_i. $\square$

A similar counterexample to the inference from (1)-(3) to (4) can be given, as the reader can verify.

According to the relational theory, (1)-(3) entails (4), and (4) entails (5), which means that (1)-(3) entails (5). Sentence (3) merely says that there are no tickets other than tickets #1 and #2, a fact that is more or less stipulated in the description of the case. So, holding fixed the truth of (3), we can say that if we want to accept (1) and (2), then we must also accept (5); if, on the other hand, we want to reject (5), we must reject the conjunction of (1) and (2). Neither option looks particularly appealing. Should we, nevertheless, accept one or other of these options? The puzzle, after all, is a puzzle, and we are unlikely to find an account of it which leaves all of our initial thoughts about it intact.

It will help to imagine that we have two photographs of our two tickets. In the first, we can see the front of each ticket, which bears the number of the ticket in black-on-white. In the second photograph, we can see the colored, numberless back of each ticket.

By the description of the case, we know that the red ticket lost. So it would seem that I should be able to point at the photograph of the red ticket and say,

[11] See Yalcin (2015, §3) for a defense of the claim that in sentences like (1), (2), and (5), the initial definite description takes wide scope over the epistemic modal.

"That ticket (is such that it) cannot be the winning ticket, since we know that it lost." From the fact that that ticket is the red ticket, I can conclude that the red ticket (is such that it) cannot be the winning ticket. This suggests that (5) is false.

But it also follows from the description of the case that we don't know the number of the winning ticket. Thus, it would seem that I should be able to point at the photograph of ticket #1 and say, "That ticket (is such that it) might be the winning ticket, since for all we know, it won." From the fact that that ticket is ticket #1, I can conclude that ticket #1 (is such that it) might be the winning ticket. Since I can do the same for ticket #2, it seems as though the conjunction of (1) and (2) must be true.

The relational semanticist will have to maintain that, in one of these two cases, I helped myself to a false premise. Since the identity claims that feature in these arguments look unobjectionable, it must be that one of the initial demonstrative premises is false. If, for example, (5) is true, then when I pointed at the red ticket and said, "That ticket (is such that it) cannot be the winning ticket," I must have said something false. If, on the other hand, the conjunction of (1) and (2) is false, I said something false either when I pointed to ticket #1 and said, "That ticket (is such that it) might be the winning ticket," or when I said the same thing while pointing at ticket #2. But this is deeply implausible. Not only do these claims appear to be true, they appear to be ones that I am in a position to assert.

While these considerations might not constitute a refutation of the relational theory, they do furnish us with some motivation to look elsewhere for a satisfactory treatment of our puzzle. We can add to that motivation by examining some further problems that arise for the relational theory.

We can approach these problems by first considering how we might use our formal apparatus to represent the lottery scenario. To that end, let $\mathsf{M} = \langle \mathsf{W}, \mathsf{R}, \mathsf{D}, \mathsf{I}, \mathsf{v_1} \rangle$ be a Kripke model, where the elements of W are 'genuine' possible worlds, and $\mathsf{v_1}$ is a world that contains individuals and tickets that satisfy our description of the lottery scenario. The accessibility relation R is such that $\mathsf{R}(\mathsf{v_1})$ represents the set of worlds compatible with what we know in $\mathsf{v_1}$, i.e. in the lottery scenario. Since the only quantified sentences in which we are interested are restricted to our two tickets, we can let D consist of just those two tickets; if we call them "α" and "β", then $\mathsf{D} = \{\alpha, \beta\}$. The interpretation function I assigns appropriate extensions to the relevant non-logical expressions (e.g. $[\![w]\!]^{v,g}$ is the winning ticket in world v). (Note that our specification of I is incomplete in some important respects, as we shall presently see.) I'll refer to M as *the lottery model*, and use the sans serif font to refer to it and its elements.

World $\mathsf{v_1}$ is a world that meets the lottery description. Note that in $\mathsf{v_1}$, it will be settled which color goes with which number, although we, *qua* characters in the lottery story, will not know in $\mathsf{v_1}$ which color goes with which number. Let's suppose that the ticket numbered 1 in $\mathsf{v_1}$ is the ticket that is blue in $\mathsf{v_1}$. Since we know the blue ticket is the winning ticket, this means that the ticket numbered 1 in $\mathsf{v_1}$ is the ticket that wins in $\mathsf{v_1}$. Let α be this ticket. This means that β is the ticket numbered 2 in $\mathsf{v_1}$, the ticket that is red in $\mathsf{v_1}$, and the ticket that loses in v_1. This table summarizes the relevant facts that obtain in v_1:

v_1
α: #1, blue, winner
β: #2, red, loser

Now, given what we know in the lottery scenario, there will be two types of worlds in $\mathsf{R}(\mathsf{v}_1)$: (i) worlds like v_1 in which ticket #1 is the blue winner and ticket #2 is the red loser; and (ii) worlds in which those numbers are swapped, i.e. worlds in which ticket #2 is the blue winner and ticket #1 is the red loser. This reflects the fact that, in v_1, we do not know the number of the winning ticket. (Of course, since we know that the blue ticket is the winner, every world in $\mathsf{R}(\mathsf{v}_1)$ will be one in which the blue ticket wins and the red ticket loses.) Now let v_2 be an arbitrary world of type (ii). The following table summarizes what we've said so far about v_1 and v_2:

v_1	v_2
α: #1, blue, winner	γ: #2, blue, winner
β: #2, red, loser	δ: #1 red, loser

I have called the two tickets in v_2 "γ" and "δ" respectively. But since our model only contains two objects, α and β, there are only two options concerning the identities of γ and δ:[12]

POSSIBILITY 1: $\alpha = \gamma$, $\beta = \delta$

POSSIBILITY 2: $\alpha = \delta$, $\beta = \gamma$

If POSSIBILITY 1 obtains, the winning blue ticket #1 in v_1 is identical to the winning blue ticket #2 in v_2, and the losing red ticket #2 in v_1 is identical to the losing red ticket #1 in v_2. If, instead, POSSIBILITY 2 obtains, the situation is reversed. If POSSIBILITY 1 obtains, the tickets do not change colors as we move from v_1 to v_2; if POSSIBILITY 2 obtains, they do not change numbers as we move from v_1 to v_2.

But the important point is this: whichever POSSIBILITY obtains, the relational theory will be committed to the truth of sentences that would appear to

[12] One might wonder if my decision to include only α and β in the domain restricts our theoretical options in some important way. It seems not. For suppose that α, for example, were identical to neither γ nor δ. Then (6), which is obviously false given an epistemic reading of the modal, would be true at v_1:

(6) Ticket #1 is such that it might not be one of the tickets in the lottery at all.

be obviously false. For example, suppose that POSSIBILITY 1 obtains. Then the following are both true at v_1:

(7) Ticket #1 is such that it might be ticket #2. $(\lambda x.\Diamond x = t_2)(t_1)$

(8) Ticket #2 is such that it might be ticket #1. $(\lambda x.\Diamond x = t_1)(t_2)$

To see why this is so, think about how we would evaluate (7) for truth at v_1. We start by finding the referent of *ticket #1* at v_1. From our table, we can see that this is α. We then check to see if there is a world v compatible with what we know in v_1 such that α is identical to the ticket numbered 2 in v. Well, if POSSIBILITY 1 obtains, $\alpha = \gamma$. Since v_2 is compatible with what we know in v_1, and since γ is the ticket numbered 2 in v_2, it follows that there is a world v compatible with what we know in v_1 such that α is identical to the ticket numbered 2 in v, for v_2 is such a world. Thus, (7) is true at v_1. Parallel reasoning reveals that (8) is also true at v_1 if POSSIBILITY 1 obtains.

If instead POSSIBILITY 2 obtains, the following are both true at v_1:

(9) The blue ticket is such that it might be the red ticket. $(\lambda x.\Diamond x = r)(b)$

(10) The red ticket is such that it might be the blue ticket. $(\lambda x.\Diamond x = b)(r)$

To see why (9) comes out true at v_1 given POSSIBILITY 2, note that the referent of *the blue ticket* at v_1 is α. Given POSSIBILITY 2, $\alpha = \delta$. And since δ is the red ticket in v_2, it follows that α is the red ticket in v_2. Since v_2 is compatible with what we know in v_1, it follows that there is a world compatible with what we know in v_1 at which α is the red ticket. Thus, (9) is true at v_1. Again, parallel reasoning reveals that (10) is true at v_1, given POSSIBILITY 2.

These results are deeply implausible. Imagine looking at the photograph of the two tickets labeled "#1" and "#2." Focus on ticket #1. Is that ticket such that it might be ticket #2? How could it be? You can see ticket #2 right there beside it, and you know that the two tickets are distinct. Or consider the photograph of the colored backs of the tickets, and focus on the red ticket. Could that ticket be the blue ticket? Again it seems not, for you can see the blue ticket right there beside the red ticket, and you know that the two tickets are distinct.[13]

These observations are closely related to one due to Yalcin (2015), who points out that the relational account is bound to make some sentences of the form *The F might not be the F* true in rather ordinary contexts. But, Yalcin notes, it seems that in most normal contexts, utterances of such sentences would be infelicitous. To see the point in connection with our example, consider the following sentences:

(11) Ticket #1 is such that it might not be ticket #1. $(\lambda x.\Diamond x \neq t_1)(t_1)$

[13] Of course, if you came to suspect, for whatever reason, that the photographs had been manipulated so as to make one ticket appear as two, you might then be entitled to utter (10), for example. But I am assuming that we know that each photograph pictures two distinct tickets.

(12) The winning ticket is such that it might not be the winning ticket. $(\lambda x.\Diamond x \neq w)(w)$

If Possibility 1 obtains, (11) will be true at v_1, and if Possibility 2 obtains, (12) will be true at v_1. But, at least in the absence of special contextual set-up, these sentences appear to be infelicitous. It is not clear what explanation of this fact the relational theorist can offer.[14]

Before leaving these issues, it is worth noting again that these problems appear to be specific to epistemic modality, and do not arise in the case of metaphysical modality. In *Naming and Necessity*, Kripke drew our attention to the metaphysical counterparts of sentences like (11) and (12):

(13) The winner of the 2016 US Presidential election (is such that he) might not have been the winner of the 2016 US Presidential election.

There seems to be an asymmetry between (11)/(12), on the one hand, and (13), on the other. For as Kripke (1980) pointed out, it is easy to access a true reading of (13) in pretty much any context.[15] This reading can be brought by saying, "The winner of the 2016 election – *that man* – *he* might not have been the winner of the 2016 election." But try this with (11): "Ticket #1 – *that ticket* – *it* might not be ticket #1." There is, at the very least, an asymmetry here that the relational approach does not explain.[16]

3 Dynamic semantics

Yalcin takes his observation concerning sentences like (11) and (12) to motivate a version of *dynamic semantics*.[17] Perhaps dynamic semantics also provides a resolution to our lottery puzzle.

What is dynamic semantics? What makes a semantic theory 'dynamic'? These are not easy questions to answer in a precise and general way.[18] But for our purposes, it suffices to say this: whereas the aim of traditional 'static' semantic theories – such as the relational semantics just discussed – is to characterize the meaning of a declarative sentence in terms of what it says about the world, the basic idea behind dynamic semantics is that the meaning of a sentence consists in how it updates a body of information. Standard empirical motivations for dynamic semantics include anaphora, presupposition, and epistemic modality.[19]

[14] Aloni (2001, Ch. 3) discusses similar examples. I return to discuss the semantic/pragmatic status of sentences like these in §7.

[15] See, for example, Kripke (1980, 41, 48).

[16] As Yalcin (2015, §A.1) observes, the problem raised by (11) and (12) is also a problem for another static semantic theory of epistemic modals, namely the *domain semantics* defended in Yalcin (2007) and MacFarlane (2011).

[17] Actually, Yalcin's arguments for the dynamic approach are more extensive than this suggests; I return to some of those arguments in §7.

[18] See Rothschild and Yalcin (2015, 2016) for some recent discussion of this issue.

[19] Some classic works in the dynamic tradition are Heim (1981), Kamp (1981), Groenendijk and Stokhof (1991), and Veltman (1996). Two recent introductions to dynamic semantics are Yalcin (2012) and Nouwen *et al.* (2016).

Here we consider a dynamic theory that is a variant of the theory presented in Groenendijk *et al.* (1997), a way of combining *Dynamic Predicate Logic* (Groenendijk and Stokhof 1991) with the dynamic semantics for epistemic modals presented in Veltman (1996). Our version of this theory differs from theory of Groenendijk et al. in two main ways. First, the semantics of Groenendijk et al. is defined over a language that doesn't contain an abstraction operator, whereas our formal language $\mathcal{L}$ does. So we will need a semantics defined for formulas containing such operators. Second, Groenendijk et al. are interested in both anaphora and epistemic modality, while our interest is restricted to the latter. This difference in aim allows us to simplify their theory in a number of ways.[20]

My presentation of the dynamic theory will proceed in two steps. I shall first present the formal theory, and then move on to say something about how to draw predictions about *assertability* and *deniability* from the theory. We will then be in a position to see how the theory handles the problems we've been discussing.

The semantics is again defined for the language $\mathcal{L}$ as specified in the previous section (**Definitions 2.1-2.4**).

Definition 3.1. A *dynamic model* $\mathcal{M}$ is an n-tuple $\langle \mathcal{W}, \mathcal{D}, \mathcal{I}, v, \sigma \rangle$, where $\mathcal{W}$ and $\mathcal{D}$ are again non-empty sets, $\mathcal{I}$ is an interpretation function as defined in **Definition 2.8**, v is an element of $\mathcal{W}$, and σ is a subset of $\mathcal{W}$.

As before, world v represents the actual world. The set of worlds σ represents the set of worlds compatible with what the relevant agents know in v. These elements do not play a role in the semantics proper, but they are useful when thinking about how a model represents a particular situation, such as the lottery scenario.

The definition of a variable assignment remains the same, as does the definition of the the extension of a term t at a model $\mathcal{M}$, world v, and variable assignment g, $[\![t]\!]^{v,g}$ (**Definitions 2.6** and **2.7**). But some new definitions are needed. Let $\mathcal{M} = \langle \mathcal{W}, \mathcal{D}, \mathcal{I}, v, \sigma \rangle$ be a dynamic model. Then:

Definition 3.2. A *possibility* in $\mathcal{M}$ is a *pair* of a world and a variable assignment.

Definition 3.3. A *discourse state* in $\mathcal{M}$ is a set of possibilities in $\mathcal{M}$.

We also define some further notation:

Definition 3.4.

For any possibility $i = \langle v, g \rangle$ and term t, $i(t) = [\![t]\!]^{v,g}$.

For any possibility $i = \langle v, g \rangle$ and predicate P, $i(P) = \mathcal{I}(P(v))$.

[20] I took the idea of adapting Groenendijk *et al.* (1997) from Yalcin (2015), but his version of the theory again differs from the one presented above. On Yalcin's theory, a formula updates a set of worlds relative to a fixed variable assignment, whereas on the approach taken here, a formula updates a set of world-assignment pairs. This difference is partly what enables us to treat definite descriptions as terms; Yalcin's approach appears to require him to treat descriptions along Russellian lines. (See footnote 30.)

Definition 3.5.
For any possibility $i = \langle v, g\rangle$ and object $o \in \mathcal{D}$, $i[x/o] = \langle v, g[x/o]\rangle$.
For any state s and object $o \in \mathcal{D}$, $s[x/o] = \{i[x/o] : i \in s\}$.

Definition 3.6.
For any possibility $i = \langle v, g\rangle$ and term t, $i[x/t] = \langle v, g[x/i(t)]\rangle$.
For any state s and term t, $s[x/t] = \{i[x/t] : i \in s\}$.

Finally, we get to the core of the semantics:

Definition 3.7. Let s be any state in any dynamic model $\mathcal{M}$. The core of the semantics is the recursive definition of *the update of s by ϕ relative to $\mathcal{M}$*, which we write as $s[\phi]$. It runs as follows:

$$s[P(t_1, ..., t_n)] = \{i \in s : \langle i(t_1), ..., i(t_n)\rangle \in i(P)\}$$

$$s[\neg\phi] = s - s[\phi]$$

$$s[\phi \wedge \psi] = s[\phi][\psi]$$

$$s[(\lambda x.\phi)(t)] = \{i \in s : i[x/t] \in s[x/t][\phi]\}$$

$$s[\exists x\phi] = \{i \in s : \text{there is an } o \in D \text{ such that } i[x/o] \in s[x/o][\phi]\}$$ [21]

$$s[\Diamond\phi] = \begin{cases} s & \text{if } s[\phi] \neq \emptyset; \\ \emptyset & \text{otherwise.} \end{cases}$$ [22]

So an atomic formula ϕ updates a state s by removing all the possibilities in s at which ϕ is false.[23] This is reminiscent of Stalnaker's account of the effect an assertion has on the context set (Stalnaker 1978). To compute the update of s by a negated sentence $\neg\phi$, one first provisionally updates s with ϕ and one then *subtracts* the result from s. A conjunction $(\phi \wedge \psi)$ updates a state s by first updating it with ϕ, and then updating the result of that update with ψ.

I will not pause here to try to explain the clauses for the abstraction operator and the existential quantifier. But they will be important in what follows, and how they work should become apparent below. The clause for the possibility operator is essentially the 'test' semantics of Veltman (1996). Updating a state s with $\Diamond\phi$ is akin to testing s for consistency with ϕ: if s is consistent with ϕ, s passes the test and survives the update; if, on the other hand, s is not consistent with ϕ, s fails the test and nothing survives.

---

[21] Equivalently:

$$s[\exists x\phi] = \bigcup_{o\in\mathcal{D}} \{i \in s : i[x/o] \in s[x/o][\phi]\}$$

[22] Equivalently:

$$s[\Diamond\phi] = \{i \in s : s[\phi] \neq \emptyset\}$$

[23] Although truth and falsity are not central notions in the dynamic theory, we can say that an atomic formula $P(t_1, .., t_n)$ is true at a possibility i just in case $\langle i(t_1), .., i(t_n)\rangle \in i(P)$.

If the other logical symbols are defined in terms of these basic connectives in the usual way, we obtain the following results:[24]

Fact 3.1. *Let s be any state in any dynamic model. Then:*
$s[\phi \vee \psi] = \{i \in s : i \in s[\phi] \text{ or } i \in s[\neg\phi][\psi]\}$
$s[\phi \rightarrow \psi] = \{i \in s : i \notin s[\phi] \text{ or } i \in s[\phi][\psi]\}$
$s[\forall x\phi] = \{i \in s : \text{every } o \in D \text{ is such that } i[x/o] \in s[x/o][\phi]\}$[25]

$$s[\Box\phi] = \begin{cases} s & \text{if } s[\phi] = s; \\ \emptyset & \text{otherwise.} \end{cases}$$ [26]

The clause for the universal quantifier will be important in what follows.

Some important logical notions can be defined as follows (cf. Groenendijk *et al.* 1997, §3.1 and Yalcin 2015, 493):

Definition 3.8. For any state s in any dynamic model $\mathcal{M}$, s *supports* formula ϕ in $\mathcal{M}$ just in case $s[\phi] = s$.

Definition 3.9. An argument from premises $\phi_1, ..., \phi_n$ to conclusion ψ is *valid* just in case: for every model $\mathcal{M}$ and every state s on $\mathcal{M}$, if s supports each of the premises, then s supports the conclusion, i.e. if $s[\phi_1] = s, ..., s[\phi_n] = s$, then $s[\psi] = s$.[27]

Definition 3.10. Let s be any state in any model $\mathcal{M}$. Formula ϕ is *consistent* with s just in case $s[\phi] \neq \emptyset$. Formula ϕ is *inconsistent* with s just in case it is not consistent with s.

Definition 3.11. A formula ϕ is *consistent* (*simpliciter*) just in case there is a model $\mathcal{M}$ and a state s such that ϕ is consistent with s. A formula is *inconsistent* just in case it is not consistent.

It will be useful to take note of the following fact:

Fact 3.2. *For any state s on any dynamic model $\mathcal{M}$, if ϕ is inconsistent with s, then s supports $\neg\phi$.*

Proof. To see why this is so, imagine that ϕ is inconsistent with s, i.e. $s[\phi] = \emptyset$. Note that $s[\neg\phi] = s - s[\phi]$. Since $s[\phi] = \emptyset$, $s - s[\phi] = s - \emptyset = s$. So $s[\neg\phi] = s$, which means that s supports $\neg\phi$. $\Box$

[24] See Groenendijk *et al.* (1997, §3).

[25] Equivalently:

$$s[\forall x\phi] = \bigcap_{o \in \mathcal{D}} \{i \in s : i[x/o] \in s[x/o][\phi]\}$$

[26] Equivalently:

$$s[\Box\phi] = \{i \in s : s[\phi] = s\}$$

[27] Although *support* is a three-place relation between discourse states, formulas, and models, I often omit mention of the model for simplicity. The same goes for the notions of consistency and inconsistency defined in **Definition 3.10**.

The notion of validity just defined means that we can use this dynamic theory to generate predictions about when certain natural language inferences are valid and when they are not. But we also want our semantic theory to issue predictions about the conditions under which a sentence is assertable. For example (and to use the case at hand), we want a theory that predicts that (1) is assertable for us in the lottery scenario given what we know, but that (5) is not. It is in characterizing assertability that the notion of an *epistemic state* comes into play.

The *epistemic state of an agent or group of agents* (at a particular time t) is the set of worlds compatible with what she or they know (at t). The relations of support and consistency defined above are relations that hold between a *discourse state* and a formula, but we can use those notions to defined parallel relations that hold between an *epistemic state* and a formula. Relations of the latter kind will then play a role in the principles governing appropriate assertion and denial.

Let $\mathcal{M}$ be a dynamic model, and let σ be a set of worlds on $\mathcal{M}$.

Definition 3.12. A discourse state s *represents* σ just in case there is a variable assignment g such that $s = \{\langle v, g\rangle : v \in \sigma\}$. (So many discourse states will represent the same epistemic state.)

Definition 3.13. σ *supports* a formula ϕ just in case there is a discourse state that represents σ and that supports ϕ.[28]

Definition 3.14. Formula ϕ is *consistent* with σ just in case there is a discourse state s such that that s represents σ and ϕ is consistent with s; ϕ is *inconsistent* with σ just in case it is not consistent with σ.

Using this notion of support, we can say that a speaker in context c is in a position to assert a sentence ϕ only if her epistemic state in c – the set of worlds compatible with what she knows in c – supports ϕ. Let's say that a speaker *denies* a sentence ϕ just in case she asserts its negation. Then, in light of **Fact 3.2**, we can say that a speaker in context c is in a position to deny ϕ only if ϕ is inconsistent with his epistemic state in c. These principles serve as a bridge between the formal semantics and the pragmatic notions of assertion and denial. They can be seen as incorporating something like the 'knowledge norm' of assertion (Williamson 1996, 2000).

So, given a dynamic model $\mathcal{M}$, we have two types of states defined on $\mathcal{M}$, epistemic states (sets of worlds) and discourse states (sets of possibilities, i.e. world-assignment pairs). The latter feature centrally in the recursive semantics, the former in the principles that bridge the formal semantics with the pragmatic notions of assertion and denial. Having distinguished these two types of states, I will continue to use the unqualified term *state* to mean *discourse state*. When I want to talk about epistemic states, I will use the full term *epistemic state*. I

[28] Since our only interest (here) is in the assertability of closed sentences, it makes no difference whether we existentially or universally quantify over discourse states in **Definitions 3.13** and **3.14**.

use s, s', etc. as variables over discourse states, and σ, σ', etc. as variables over epistemic states.

In §2, we had two complaints about the relational theory. First, the relational theory forces us to choose between the truth of (1) and (2), on the one hand, and the falsity of (5) on the other. Our second complaint was that the relational theory was bound to make true certain sentences that seem obviously false (in the lottery scenario), such as (7) and (9).

The dynamic theory avoids both of these problems. The key fact about the dynamic theory is this:

Fact 3.3. *For every dynamic model $\mathcal{M}$, state s on $\mathcal{M}$, term t, variable x, and formula ϕ, if t is free for x in ϕ, then: $s[(\lambda x.\phi)(t)] = s[\phi(t/x)]$.*

(The proof of **Fact 3.3** is discussed in an appendix.) Note that an immediate consequence of this is that the update associated with a *de re* modal predication is identical to the update associated with its *de dicto* counterpart: $s[(\lambda x.\Diamond\phi)(t)] = s[\Diamond\phi(t/x)]$. This means that the epistemic state of an agent (or group thereof) will support the one just in case it supports the other. And this in turn means that a *de re* modal predication will be assertable by an agent (on a particular occasion) just in case its *de dicto* counterpart is also assertable by that agent (on that occasion).[29]

Fact 3.3 means that the dynamic theory has a simple solution to the two problems we've been discussing. Take the first problem: the fact that (1) and (2) both seem true in the lottery scenario, while (5) seems false. Or, to the point in a manner more congenial to the dynamic theorist: given what we know in the lottery scenario, it seems that we are in a position to assert (1) and (2) and to deny (5).

To see how the dynamic theory predicts this, it will help to think about how to represent the lottery scenario within the dynamic framework. So let $\mathsf{M} = \langle \mathsf{W}, \mathsf{D}, \mathsf{I}, \mathsf{v}_1, \sigma \rangle$ be the *dynamic lottery model.* Each of the elements W, D, I, and v_1 of this dynamic model is identical to the corresponding element of the Kripke lottery model defined in §2, and σ is the set of worlds compatible with what we know in v_1. So (some of) the relevant facts about M can be again displayed as follows:

v_1	v_2
α: #1, blue, winner	γ: #2, blue, winner
β: #2, red, loser	δ: #1 red, loser

World v_1 is the world we inhabit in the lottery scenario, and v_2 is an arbitrary world compatible with what we know in which ticket #2 wins. Since σ is our

[29]I should note that if metaphysical modals are added to the language, **Fact 3.3** will no longer hold in general. But terms would still remain scopeless with respect to all other operators, which is all that is needed for our purposes.

epistemic state at v_1, both v_1 and v_2 are elements of σ. Let s be a discourse state on M that represents σ. So $\mathsf{s} = \{\langle v, g\rangle : v \in \sigma\}$, for some variable assignment g on M. Let $\mathsf{i}_1 = \langle \mathsf{v}_1, g\rangle$ and $\mathsf{i}_2 = \langle \mathsf{v}_2, g\rangle$; so i_1 and i_2 are elements of s.

In light of **Fact 3.3**, we can show that s supports (1) and (2) by showing that s supports their *de dicto* counterparts, (1*) and (2*).

(1*) It might be that ticket #1 is the winning ticket. $\quad \Diamond t_1 = w$

(2*) It might be that ticket #2 is the winning ticket. $\quad \Diamond t_2 = w$

Consider (1*). By the clause for the possibility operator, we have:

$$\mathsf{s}[\Diamond t_1 = w] = \begin{cases} s & \text{if } \mathsf{s}[t_1 = w] \neq \emptyset; \\ \emptyset & \text{otherwise.} \end{cases}$$

And by the clause for identity, we have $\mathsf{s}[t_1 = w] = \{i \in \mathsf{s} : i(t_1) = i(w)\}$. So $\mathsf{s}[t_1 = w] \neq \emptyset$ just in case there is a possibility $i \in \mathsf{s}$ such that $i(t_1) = i(w)$. Since ticket #1 is the winning ticket in v_1, $\mathsf{i}_1(t_1) = \mathsf{i}_1(w)$. Thus, $\mathsf{s}[t_1 = w] \neq \emptyset$ which means that $\mathsf{s}[\Diamond t_1 = w] = \mathsf{s}$. So s supports (1*), which, in light of **Fact 3.3**, means that s supports (1). Parallel reasoning shows that, since ticket #2 is the winning ticket ticket in v_2, s supports (2*) and so supports (2). Thus, our epistemic state in the lottery scenario supports (1) and (2), which means that the dynamic theory predicts that those sentences are ones we are in a position to assert.

Now consider (5). In light of **Fact 3.3**, we can show that (5) is inconsistent with s by showing that its *de dicto* counterpart (5*) is inconsistent with s.

(5*) It might be that the red ticket is the winning ticket. $\quad \Diamond r = w$

By the clause for the possibility operator, we know that $\mathsf{s}[\Diamond r = w] = \emptyset$ if $\mathsf{s}[r = w] = \emptyset$. This, in turns, holds just in case there is no possibility $i \in s$ such that $i(r) = i(w)$. Since we, *qua* characters in the fictional lottery scenario, know that that the red ticket is not the winning ticket, there will be no world in our epistemic state σ at which the red ticket is the winning ticket. Accordingly, there will be no possibility $i \in \mathsf{s}$ such that $i(r) = i(w)$. Thus, $[\Diamond r = w] = \emptyset$, which means that $\Diamond r = w$ is inconsistent with s. Note that, in light of **Fact 3.2**, this means that s supports $\neg\Diamond r = w$. Given this and **Fact 3.3**, the dynamic theory predicts that we are in a position to deny (5) in the lottery scenario.

The dynamic theory also predicts that both (7) and (9) are inconsistent with our epistemic state in the lottery scenario, which means that their negations are assertable.

(7) Ticket #1 is such that it might be ticket #2. $\quad (\lambda x.\Diamond x = t_2)(t_1)$

(9) The blue ticket is such that it might be the red ticket. $\quad (\lambda x.\Diamond x = r)(b)$

One can see this by noting that their *de dicto* counterparts are inconsistent with our epistemic state in the lottery scenario. The dynamic theory also predicts that the Yalcin-style examples, (11) and (12), are, like their *de dicto* counterparts, inconsistent full-stop.

(11) Ticket #1 is such that it might not be ticket #1. $(\lambda x.\Diamond x \neq t_1)(t_1)$

(12) The winning ticket is such that it might not be the winning ticket. $(\lambda x.\Diamond x \neq w)(w)$[30]

All of this is good news for the dynamic theory. Now the bad news. We noted above that the dynamic theory predicts that (1) and (2) are assertable in the lottery scenario, and that (5) is not. It also predicts that (3) is assertable in that scenario (as it should). This means that the model-state pair $\langle \mathsf{M}, \mathsf{s} \rangle$ will support (1)-(3) but will fail to support (5). Now since an argument is valid just in case every state in every model that supports the premises supports the conclusion, we know that the inference from (1)-(3) to (5) fails. So this must mean that one or both of our two inferences – the inference from (1)-(3) to (4) or the inference from (4) to (5) – must fail. In fact, it turns out that both inferences fail on the dynamic theory.

We shall (for simplicity) focus our attention on the inference from (4) to (5).

Fact 3.4. *The inference from (4) to (5) is invalid on the dynamic semantics.*

Proof. A *counterexample* to an inference in the dynamic theory is a model-state pair that supports the premises but fails to support the conclusion. Consider a model $\mathcal{M} = \langle \mathcal{W}, \mathcal{D}, \mathcal{I}, u_1, \sigma \rangle$. Let u_2 be an another element of $\mathcal{W}$, let $\mathcal{D} = \{\alpha, \beta\}$, and let the relevant facts about $\mathcal{M}$ be displayed as follows:

u_1	u_2
α: #1, blue, winner	β: #2, blue, winner
β: #2, red, loser	α: #1, red, loser

Note that u_1 and u_2 are qualitatively just like v_1 and v_2 (respectively). But the transworld identities that were left open in the model M have here been settled in a particular way. Let s be a discourse state on $\mathcal{M}$ such that $s = \{i_1, i_2\}$, where $i_1 = \langle u_1, g \rangle$ and $i_2 = \langle u_2, g \rangle$, for some variable assignment g on $\mathcal{M}$.

First, note that s supports (4). To see this, it helps to take note of following fact:

Fact 3.5. *Let s be any state on any dynamic model. Then:*

$$s[\forall x \Diamond x = w] = \begin{cases} s & \textit{if for all } o \in \mathcal{D}, \textit{ there is an } i \in s \textit{ such that } o = i(w); \\ \emptyset & \textit{otherwise.} \end{cases}$$

[30] Yalcin (2015, 508) speculates that it is difficult to reconcile the infelicity of (11) and (12) with the Fregean view that definite descriptions are terms. The foregoing results appear to show otherwise. This is significant even if one is otherwise inclined to be a Russellian about definite descriptions, since (as we noted earlier) the same problems appear to arise in connection with proper names. For example, the following sentence is just as odd as (11) and (12), given an epistemic reading of *might*:

Cicero is such that he might not be Cicero.

The present approach gives the same account of this as it gives of (11) and (12).

With **Fact 3.5** in hand, we can see by consulting the above table that s supports (4). For there are only two objects in the domain, α and β. Since α is the winning ticket at u_1, there is an $i \in s$ such that $\alpha = i(w)$, for i_1 is such an i. Since β is the winning ticket at u_2, there is an $i \in s$ such that $\beta = i(w)$, for i_2 is such an i. Since there is nothing in the domain other than α and β, every object $o \in \mathcal{D}$ is such that there is a possibility $i \in s$ such that $o = i(w)$. Thus, $s[\forall x \Diamond x = w] = s$.

We can also see that s does not support (5). Given **Fact 3.3**, this can be established by showing that s does not support (5)'s *de dicto* counterpart, $\Diamond r = w$. Given the clause for the possibility operator, s supports $\Diamond r = w$ just in case there is an $i \in s$ such that $i(r) = i(w)$. But there are only two possibilities in s, i_1 and i_2, and $i_1(r) \neq i_1(w)$ and $i_2(r) \neq i_2(w)$. So s does not support $\Diamond r = w$, which, in light of **Fact 3.3**, means that s does not support (5). Thus, this model $\mathcal{M}$ and state s constitute a counterexample to the inference from (4) to (5). □

So the inferences from (1)-(3) to (4) and from (4) to (5) are invalid according to the dynamic theory. Is this a problem? On the one hand, we might think it is, since these inferences *seem* valid.[31] Just look back at the first page of this essay and run through those two inferences in your head. Doesn't each taken individually appear to be a piece of impeccable reasoning? On the other hand, we might think that the failure of these two inferences in dynamic semantics is not, on reflection, a genuine problem, since (one might think) the following cannot all be true:

- Sentences (1)-(3) are all supported by our epistemic state in the lottery scenario.
- Sentence (5) is not supported by our epistemic state in the lottery scenario.
- Sentences (1)-(3) entail sentence (5).

My view of the matter is as follows. The *apparent* validity of the inferences from (1)-(3) to (4) and from (4) to (5) is the sort of thing we would like our semantic theory (or semantic-cum-pragmatic theory) to explain, just as we want it to explain, e.g., the apparent unassertability of (5) in the lottery scenario. That explanation need not take the form of a *validation*: it needn't predict that that the inferences in question really are valid, but it should explain their apparent validity. That is, an adequate theory should go some way to accounting for the fact the inferences seem compelling, at least at first glance. For example: we shall, in due course, see an alternative dynamic theory that (like the present dynamic theory) also predicts that these inferences are invalid, but which also offers an explanation for why those inferences seem valid. The problem, to my mind, with the present dynamic theory is not simply that it predicts that these

[31] Or: they seem valid *modulo* the assumption that the terms they contain have referents.

inferences are invalid, but that it comes with no backup story about why they appear to be valid.

If that line of objection to the dynamic theory doesn't impress you, perhaps the following one will. Another way to see the problem – or another problem, depending on how you count – is that the feature of the dynamic theory that leads it to invalidate these inferences also leads it to predict that certain contradictory-sounding sentences should be assertable in the right circumstances. Given the sort of considerations that were used to motivate the dynamic theory (e.g. the infelicity of (11) and (12)), this fact ought to give its advocates pause.

Consider the following sentences:

(14) Although any ticket might be the winning ticket, the red ticket is such that it cannot be the winning ticket. $(\forall x \Diamond x = w) \wedge (\lambda x. \neg \Diamond x = w)(r)$

(15) Although the blue ticket is such that it must the be the winning ticket, no ticket is such that it must be the winning ticket. $(\lambda x. \Box x = w)(b) \wedge (\neg \exists x \Box x = w)$

These appear to be unassertable, in the sense that they are (almost?) never acceptable. In each case, the second conjunct seems to be at odds with what the first conjunct claims. But:

Fact 3.6. *Neither (14) nor (15) is inconsistent according to dynamic semantics. Indeed, there is a model-state pair that supports them both.*

Proof. The model-state pair $\langle \mathcal{M}, s \rangle$ that served as our counter-example to the inference from (4) to (5) supports both of these sentences. To see this in connection with (14), note that $s[(\forall x \Diamond x = w) \wedge (\lambda x. \neg \Diamond x = w)(r) = s$ just in case $s[(\forall x \Diamond x = w)][(\lambda x. \neg \Diamond x = w)(r)] = s$. As we noted in our proof of **Fact 3.4**, $s[(\forall x \Diamond x = w)] = s$. Thus, it remains to show that $s[(\lambda x. \neg \Diamond x = w)(r)] = s$. Given **Fact 3.3** this will hold just in case $s[\neg \Diamond r = w] = s$. And this will hold just in case $(s - s[\Diamond r = w]) = s$. This, in turn, will hold just in case $s[\Diamond r = w] = \emptyset$, which it does since there are no possibilities $i \in s$ such that $i(r) = i(w)$. Thus, s supports (14). The reader can verify that s also supports (15). □

So given the connection between support and assertion, this means that, if one's epistemic state can be represented by the above discourse state s, then one should be able to assert both (14) and (15). But this prediction seems wrong: these sentences appear to be contradictory, and so are not fit candidates for assertion no matter what one's epistemic state is like. To bring this out in connection with (14), consider this exchange:

(16) *A*: Any ticket might win.

B: No, that's not true. The organizer just told us that the red ticket lost, so the red ticket cannot be the winner.

A: Yes, that's right. Any ticket might win, but the red ticket cannot win.

It seems that B is right to think that her information contradicts what A said initially. This explains why A's response to B is so puzzling. But of course if A's epistemic state is represented by s, then A's state will support his response.

Note that, for all its faults, the relational semantics has no similar problem here, for it validates the inference from (4) to (5), and predicts that (14) and (15) are indeed contradictory.

4 A problem for everyone

Let's take stock. We complained that the relational theory validated our two inferences, and so required us to choose between the truth of (1) and (2), on the one hand, and the falsity of (5), on the other. We also noted that the relational theory suffered from an additional flaw: it is bound to predict that certain sentences that seem obviously false (in the lottery scenario) are true. We then noted that while the dynamic theory avoided both of these problems, it faces problems of its own, problems that stem from the fact that it *invalidates* the inferences in question.

Thus far, our inquiry has been somewhat negative: problems and more problems. I want to get to some constructive proposals, but before we do that, it will help to consider one further problem, one that points us in the direction of a solution. This final problem is a problem for everyone, that is, for both the relational and the dynamic semantics we've been discussing.

This problem arises when we consider the question: is (4) true/assertable in the lottery scenario? Is every ticket such that it might be the winning ticket? As Aloni (2001, Ch. 3) and Moss (2016, Ch. 8) argue, the right answer to our question seems to be: it depends. It depends on which ways of 'conceptualizing' the tickets are salient in the context of utterance.

A (somewhat contrived) example brings this out. Suppose that in the lottery scenario, Al forgets whether or not we know the number of the winning ticket. He points at the photograph on which the number of each ticket is visible and asks, "Might either of these tickets be the winning ticket?" The right answer here seems to be "yes." Since we don't know the number of the winning ticket, either of those tickets might be the winning ticket. Thus, it seems that (4) is assertable in this context. But now suppose that Barbara comes along, having forgotten whether we know the color of the winning ticket. She points at the photograph of the colored backs of the tickets, and asks, "Might *any* of these tickets be the winning ticket?" Here it seems right for us to say, "No, we know that the red ticket lost, so the red ticket can't be the winning ticket." In this context, (4) does not appear to be assertable.

Now, both the relational and the dynamic theory can accommodate two ways in which (4) might depend on the context, but neither helps to explain the foregoing observations. First, both theories can acknowledge that the context of utterance might play a role in determining the domain of the quantifier in (4). But (assuming that quantifier domains are sets of individuals) this sort of context-sensitivity is of no help here, because in both of the above contexts the same tickets are at issue (and it is known that the same tickets are at issue), and so the domain of the quantifier presumably does not change between the two contexts. Second, both theories make the truth or assertability of (4) sensitive to the epistemic state relevant in the context. In both theories, when *might* is used in a particular context, it is interpreted as quantifying over an epistemic state that is determined by that context. On the relational theory, the context accomplishes this by contributing an accessibility relation; on the dynamic theory, the context can be seen as providing the epistemic state relevant for assessing the assertability of sentences. Although neither of the formal theories we have been considering makes this sort of dependence on context explicit, such dependence is part of our informal understanding of those theories.

But in the two contexts above, the same epistemic state is at issue both times. We do not gain or lose knowledge about the outcome of the lottery when we move from talking to Al to talking to Barbara. All that changes between the two contexts is how the two tickets are being thought of. In the first context, the salient way of thinking of the two tickets is via their numbers – as *ticket #1* and as *ticket #2*. In the second context, the salient way of thinking of them is via their colors – as *the red ticket* and as *the blue ticket*. But merely changing which ways of thinking are salient doesn't change which worlds are compatible with the totality of one's knowledge.

Thus, the sort of context-sensitivity at issue does not appear to be captured by either theory, for it is not explicable in terms of quantifier domain restriction nor in terms of the general 'epistemic state-sensitivity' of epistemic modals. Quantified epistemic modal sentences are sensitive to the utterance context in a way that neither relational nor dynamic semantics predicts.

Let us examine the problem in more detail. Start with the relational theory, and recall the Kripke lottery model $\mathsf{M} = \langle \mathsf{W}, \mathsf{R}, \mathsf{D}, \mathsf{I}, \mathsf{v}_1 \rangle$. Note that the relevant features of v_1 can be depicted as follows:

v_1
α: #1, blue, winner
β: #2, red, loser

World v_1 is the world we are imagined to occupy in the lottery scenario. According to the relational semantics, (4) is true at v_1 just in case every object o in the domain is such that there is a world accessible from v_1 at which o is the winning ticket. Thus, in order to determine the truth of (4) at v_1 in $\mathcal{M}$, we need to know whether both α and β have the property of being the winning ticket at some world accessible from v_1. We know that α has this property, since $\mathsf{v}_1 \mathsf{R} \mathsf{v}_1$, and α is the winner at v_1. But what about β? If β has this property it will have to be in virtue of its being the winner at some accessible world distinct from v_1, since we know that β loses at v_1. So: is there an alternative world accessible from v_1 at which β is the winning ticket?

No matter how we answer this question, the relational theory will not be able to acknowledge the way in which the truth of (4) depends on the context. If we say there is such a world, then (4) is true in both contexts, since every object in the domain will be such that there is an accessible world at which it wins. This will validate our response to Al, but not to Barbara. If, on the other hand, we say that there is no such world we have the reverse problem: (4) will be false in both contexts, since some object in the domain – namely, β – is such that there is no accessible world at which it wins. This will validate our response to Barbara, but not to Al. So no matter how we answer this question, we cannot respect our judgments about both conversations.

It seems that what we want to say is that, when we are thinking of β as *ticket #2*, then there is an accessible world that *represents* β as being the blue winner. When, on the other hand, we are thinking of β as *the red ticket*, then no accessible world *represents* β as being the blue winner. That is, what we want to say is that, when we are thinking of β as *ticket #2*, there ought to be an accessible world like v_2 in which γ, the blue winner at that world, represents β:

v_1	v_2
α: #1, blue, winner	γ: #2, blue, winner
β: #2, red, loser	δ: #1 red, loser

But when we are thinking of β as *the red ticket*, no accessible world should represent β as being the blue winner. So when we are thinking of β as *the red ticket*, δ should *not* represent β any longer at v_2 – β's representative at v_2 should now be δ rather than γ.

The problem with the relational theory is that the only object that can represents β at an alternative world is β itself. And since β cannot be identical to γ at one time, and then identical with δ at another time (since $\gamma \neq \delta$), a world like v_2 simply either represents β as being the blue winner or it does not. How that world represents β as being never changes, and how we are thinking of β at a particular time simply never comes into the picture.

Consider now the dynamic theory, and let us ask what that theory says about sentence (4). In particular: does it predict that (4) is assertable in the lottery scenario or not? To assess this question, consider again the *dynamic* lottery model $\mathsf{M} = \langle \mathsf{W}, \mathsf{R}, \mathsf{D}, \mathsf{I}, \mathsf{v}_1, \sigma \rangle$. Let s again be equal to $\{\langle v, g \rangle : v \in \sigma\}$, where g is a some variable assignment on M. And let $\mathsf{i}_1 = \langle \mathsf{v}_1, g \rangle$. The question we need to asses is: does s support (4)?

In approaching this question, it will help to remind ourselves that it follows from **Fact 3.5** that:

$$\mathsf{s}[\forall x \Diamond x = w] = \begin{cases} \mathsf{s} & \text{if for all } o \in \mathcal{D}\text{, there is an } i \in \mathsf{s} \text{ such that } o = i(w); \\ \emptyset & \text{otherwise.} \end{cases}$$

Since our model only contains two objects, α and β, s supports (4) just in case α is the winning ticket at some possibility in s and β is the winning ticket at some possibility in s. We know that α is, and β is not, the winning ticket in i_1, so the crucial question is whether there is a possibility in s distinct from i_1 at which β wins. This is essentially equivalent question we faced above when discussing the relational theory, since it reduces to the question of whether there is a world compatible with what we know in v_1 at which β is the winning ticket.

And we face the same dilemma as we did before. Either there is such a world, in which case s supports (4), and the dynamic theory predicts that that sentence is assertable in the lottery scenario, or there is no such world, in which case the dynamic theory predicts that we are in a position to deny that sentence. If we take the former option, we make the right prediction about the correct response to Al, but the wrong prediction about the correct response to Barbara. If, on the other hand, we take the latter option, the situation is simply reversed. Whichever option we choose, the dynamic theory will not accommodate the fact that whether (4) is assertable in a context c depends on the ways of thinking of the tickets that are salient in c.

In the course of his critique of quantified modal logic, Quine once claimed that "being necessarily or possibly thus and so is not a trait of the object concerned, but depends on the manner of referring to the object" (Quine 1953, 148). The modality Quine had in mind was *analyticity*, truth in virtue of meaning. Whether or not Quine was right about analyticity, it seems to me that his remark is correct if the epistemic sense of *necessarily* and *possibly* is at issue. Whether a particular ticket is such that it might be the winning ticket depends on how that ticket is described or thought of. In my view, the source of the aforementioned problems is that neither of the standard approaches takes this Quinean insight into account.

In the remainder of the paper, we consider how each of the foregoing theories might implement this Quinean insight. As it turns out, for each theory, the modifications to the theory that are needed to accommodate the Quinean point help the theory to avoid the other problems it faces. After examining these points, we close by considering how we might choose between the resulting two theories.

5 Carnapian dynamic semantics

We need our semantics to be sensitive in some way to the 'ways of thinking' of individuals that are salient in the context. But what is a 'way of thinking' of an individual? In possible worlds treatments of attitude verbs, a way of thinking of an individual is often identified with an *individual concept*, a function from worlds to individuals. For example, if you believe that Cicero is an orator, the possible worlds theorist might represent the content of your belief with the set of worlds v such that $c(v)$ is an orator in v, where c is an individual concept that maps a world to the individual there who represents Cicero for you there. So c might map a world v to the individual called "Cicero" in v (if there is one). If you also believe that Tully is not an orator, the possible worlds theorist might represent the content of that belief with the set of worlds in which $t(v)$ is not an orator in v, where t is a function that maps a world w to the individual called "Tully" in v (if there is one).

Among other things, this allows the possible worlds theorist to say that these two beliefs are consistent, as they appear to be (at least to many philosophers). This is because while the individual concepts c and t map the actual world to the same individual, namely Cicero/Tully, they may yield different outputs when given a counterfactual world as input, if the input world v is such that the person called "Cicero" in v is distinct from the person called "Tully" in v.

In this section, we show how we can use individual concepts to model ways of thinking within the dynamic framework; in the section that follows, we consider a related proposal in the context of relational semantics.

The basic idea to be pursued in this section is that the utterance context provides a set of ways of thinking of individuals by providing an appropriate set of individual concepts. Quantifiers then range over these contextually supplied sets of individual concepts, rather than over sets of individuals. The idea that quantifiers range over sets of individual concepts is familiar from systems of *contingent identity*, and is perhaps most closely associated with Carnap (1947).[32] For that reason, I shall call the theory presented in this section *Carnapian dynamic semantics*.[33]

The models for the Carnapian theory are the dynamic models defined earlier (**Definition 3.1**). But the definition of a variable assignment changes. Let $\mathcal{M} = \langle \mathcal{W}, \mathcal{D}, \mathcal{I}, v, \sigma \rangle$ be a dynamic model.

Definition 5.1. A *variable assignment* on $\mathcal{M}$ in the Carnapian semantics is a function from variables to *individual concepts*, i.e. total functions from $\mathcal{W}$ into $\mathcal{D}$.

The definitions of a *possibility* and of a *discourse state* remain the same (**Definitions 3.2** and **3.3**). The definition of the extension of a term at a possibility changes slightly, to accommodate the fact that variables now range over individual concepts:

Definition 5.2.

[32] See also Gibbard (1975) who applies the Carnapian approach to a puzzle in the metaphysics of material constitution, a puzzle that certain structural affinities with our lottery puzzle.

[33] The theory presented in this section is quite similar, in relevant respects, to the theory presented in Aloni (2001, Ch. 3). One difference is that our language $\mathcal{L}$ contains individual constants and an abstraction operators, whereas Aloni's does not. The other differences stem in part from the fact that Aloni is interested in giving a theory of epistemic modals together with a theory of anaphora and presupposition.

$[\![t]\!]^{v,g} = g(t)(v)$, if t is a variable

$[\![t]\!]^{v,g} = \mathcal{I}(t)(v)$, if t is an individual constant

And we need to revise some of our notation. Let $\mathcal{M}$ be any dynamic model.

Definition 5.3.

For any possibility $i = \langle v, g\rangle$ and individual concept c on $\mathcal{M}$, $i[x/c] = \langle v, g[x/c]\rangle$

For any state s and individual concept c on $\mathcal{M}$, $s[x/c] = \{i[x/c] : i \in s\}$

Definition 5.4.

For any possibility $i = \langle v, g\rangle$ and term t, $i[x/t] = \langle v, g[x/\mathcal{I}(t)]\rangle$

For any state s and term t, $s[x/t] = \{i[x/t] : i \in s\}$

We want the semantics to reflect the idea that an utterance context provides a set of ways of thinking of individuals by providing a *set of individual concepts*. Let us call a non-empty set of individual concepts a *concept set*.[34] The Carnapian semantics then takes the form of a recursive definition of *the update of a state s by a formula ϕ relative to a model $\mathcal{M}$ and a concept set $\mathcal{C}$*, which we write as $s[\phi]^{\mathcal{C}}$. The relativization to a concept set only affects the clause for the existential quantifier; all other clauses, including the clause for the modal operator, remain essentially as they were in **Definition 3.7**. Here is the new clause for the existential quantifier:

$$s[\exists x\phi]^{\mathcal{C}} = \{i \in s : \text{there is a } c \in \mathcal{C} \text{ such that } i[x/c] \in s[x/c][\phi]^{\mathcal{C}}\}$$

Note that even though variables now range over intensions, predicates are still extensional in the sense that the extension of an n-ary predicate is a set of n-ary sequences of *individuals* found in $\mathcal{D}$. This is reflected in the fact that the clause for atomic formulas does not change.

We get parallel results for the defined logical symbols to the results stated in **Fact 3.1**. In particular, since $\forall x\phi$ is still defined as $\neg\exists x\neg\phi$, we have:

$$s[\forall x\phi]^{\mathcal{C}} = \{i \in s : \text{for every } c \in \mathcal{C},\ i[x/c] \in s[x/c][\phi]^{\mathcal{C}}\}$$

The definition of the logical notions remains essentially unchanged, except now the relativization to concept sets is incorporated:

Definition 5.5. For any state s and concept set $\mathcal{C}$ on any dynamic model $\mathcal{M}$, s *supports* formula ϕ relative to $\mathcal{C}$ and $\mathcal{M}$ just in case $s[\phi]^{\mathcal{C}} = s$.

Definition 5.6. An argument from premises $\phi_1, ..., \phi_n$ to conclusion ψ is *valid* just in case: for every dynamic model $\mathcal{M}$ and every state s and concept set $\mathcal{C}$ on $\mathcal{M}$, if s supports each of the premises relative to $\mathcal{C}$ and $\mathcal{M}$, then s supports the conclusion relative to $\mathcal{C}$ and $\mathcal{M}$, i.e. if $s[\phi_1]^{\mathcal{C}} = s, ..., s[\phi_n]^{\mathcal{C}} = s$, then $s[\psi]^{\mathcal{C}} = s$.

[34] Aloni (2001, Ch. 3) requires the concept sets over which quantifiers range to be what she calls *conceptual covers*. A concept set $\mathcal{C}$ (on model $\mathcal{M}$) is a *conceptual cover* just in case for each object $o \in \mathcal{D}$ and each world $w \in \mathcal{W}$, there is exactly one individual concept $c \in \mathcal{C}$ such that $c(w) = o$. I am not opposed to this constraint, but nothing I say will turn on whether it is adopted.

Definition 5.7. Let s be any state and $\mathcal{C}$ any context on any dynamic model $\mathcal{M}$. Formula ϕ is *consistent* with s relative to $\mathcal{C}$ and $\mathcal{M}$ just in case $s[\phi]^{\mathcal{C}} \neq \emptyset$. Formula ϕ is *inconsistent* with s, relative to $\mathcal{C}$ and $\mathcal{M}$, just in case it is not consistent with s relative to $\mathcal{C}$ and $\mathcal{M}$.

Definition 5.8. A formula ϕ is *consistent* (*simpliciter*) just in case there is a model $\mathcal{M}$ and a state s and context $\mathcal{C}$ on $\mathcal{M}$ such that ϕ is consistent with s relative to $\mathcal{C}$ and $\mathcal{M}$. A formula is *inconsistent* just in case it is not consistent.

The definition of what is for a discourse state to represent an epistemic state remains the same (**Definition 3.12**). The semantic-pragmatic bridge relations must be re-defined to take into account the relativity to concept sets: Let σ be any set of worlds, and $\mathcal{C}$ any concept set, on any dynamic model $\mathcal{M}$.

Definition 5.9. σ *supports* a formula ϕ relative to $\mathcal{C}$ and $\mathcal{M}$ just in case there is a discourse state that represents σ and that supports ϕ relative to $\mathcal{C}$ and $\mathcal{M}$.

Definition 5.10. Formula ϕ is *consistent* with σ relative to $\mathcal{C}$ and $\mathcal{M}$ just in case there is a discourse state s such that that s represents σ and ϕ is consistent with s relative to $\mathcal{C}$ and $\mathcal{M}$. ϕ is *inconsistent* with σ relative to $\mathcal{C}$ and $\mathcal{M}$ just in case is not consistent with σ relative to $\mathcal{C}$ and $\mathcal{M}$.

Definition 5.11. A formula ϕ is *consistent* (*simpliciter*) just in case there is a model $\mathcal{M}$, a state s, and a concept set $\mathcal{C}$ such that ϕ is consistent with s relative to $\mathcal{C}$ and $\mathcal{M}$. A formula is *inconsistent* just in case it is not consistent.

We assume that a speaker in utterance context u is in a position to assert ϕ only if her epistemic state in u supports ϕ relative to $\mathcal{C}^u$, where $\mathcal{C}^u$ is the concept set determined by u. And a speaker is in a position deny ϕ in u only if ϕ is inconsistent with her epistemic state in u relative to $\mathcal{C}^u$.

We want now to demonstrate two things about the Carnapian theory. First, we want to show that it has the resources for dealing with the objections we leveled against the dynamic theory of §3. Second, we want to make sure nothing was lost in moving from the earlier dynamic theory to the present one, i.e. we want to make sure that the new theory retains the virtues of the old one.

This second issue can be dealt with quickly. For the virtues of the earlier dynamic theory all flowed from the fact that the update associated with a *de re* modal predication was identical to the update associated with its *de dicto* counterpart (**Fact 3.3**). So we can reassure ourselves that nothing has been lost in moving to the new theory by noting that the analogue of this fact holds for the Carnapian dynamic theory:

Fact 5.1. *For every dynamic model $\mathcal{M}$, state s and concept set $\mathcal{C}$ on $\mathcal{M}$, term t, variable x, and formula ϕ, if t is free for x in ϕ, then:* $s[(\lambda x.\phi)(t)]^{\mathcal{C}} = s[\phi(t/x)]^{\mathcal{C}}$.

This means that, like the earlier dynamic theory, the Carnapian theory will not require us to choose between accepting (1) and (2), on the one hand, and denying (5), on the other. Similarly, like the earlier dynamic theory, the Carnapian theory issues appropriate verdicts about sentences like (7)-(12).

Let's turn now to see how the Carnapian theory deals with our objections to the dynamic theory of §3. Let's begin with the fact that sentences like (4) are sensitive to which ways of thinking of the tickets are salient in the context.

(4) Any ticket might be the winning ticket. $\forall x \Diamond x = w$

Consider again the dynamic lottery model $\mathsf{M} = \langle \mathsf{W}, \mathsf{D}, \mathsf{I}, \mathsf{v}_1, \sigma \rangle$, and discourse state s on M that represents σ, our epistemic state at v_1. We assume s contains i_1 and i_2, where $\mathsf{i}_1 = \langle \mathsf{v}_1, g \rangle$ and $\mathsf{i}_2 = \langle \mathsf{v}_2, g \rangle$, for some variable assignment g on M.

v_1	v_2
α: #1, blue, winner	γ: #2, blue, winner
β: #2, red, loser	δ: #1 red, loser

We want to show is that we are in a position to assert (4) in our conversation with Al, a conversation in which the salient way of thinking of the tickets is via their numbers. And we want to show that we are in a position to deny (4) in our conversation with Barbara, a conversation in which the salient way of thinking of the tickets is via their colors.

Since the salient way of thinking of the tickets in our conversation with Al is via their numbers, the Carnapian theory will say that the 'Al context' determines the set of individual concepts $\mathcal{C}_n$, where $\mathcal{C}_n$ contains (only) the following two individual concepts:

$\mathsf{I}(t_1)$, which maps each world v to the ticket numbered 1 in v, so:

$\mathsf{I}(t_1)(\mathsf{v}_1) = \alpha$, $\mathsf{I}(t_1)(\mathsf{v}_2) = \delta$

$\mathsf{I}(t_2)$, which maps each world v to the ticket numbered 2 in v, so:

$\mathsf{I}(t_2)(\mathsf{v}_1) = \beta$, $\mathsf{I}(t_2)(\mathsf{v}_2) = \gamma$

Individual concept $\mathsf{I}(t_1)$ corresponds to our way of thinking of α when we are thinking of it as *ticket #1*, and $\mathsf{I}(t_2)$ corresponds to our way of thinking of β when we are thinking of it as *ticket #2*. These individual concepts are also obviously the intensions of the terms *ticket #1* and *ticket #2*, respectively.

We want to show that s supports (4) relative to $\mathcal{C}_n$. To see this, first note the following analogue of **Fact 3.5** holds in the Carnapian theory:

Fact 5.2. *Let s be any state, and $\mathcal{C}$ any concept set, on any dyamic model. Then:*

$$s[\forall x \Diamond x = w]^{\mathcal{C}} = \begin{cases} s & \textit{if for all } c \in \mathcal{C}, \textit{ there is an } i \in s \textit{ such that } c(i) = i(w); \\ \emptyset & \textit{otherwise.} \end{cases}$$

(Where $i = \langle v, g \rangle$ is a possibility and c an individual concept, $c(i) = c(v)$.) Since $\mathcal{C}_n$ contains only two individual concepts, $\mathsf{I}(t_1)$ and $\mathsf{I}(t_2)$, if the following two conditions holds, then s will support (4) relative to C_n:

- there is an $i \in \mathsf{s}$ such that $\mathsf{I}(t_1)(i) = i(w)$; and
- there is an $i \in \mathsf{s}$ such that $\mathsf{I}(t_2)(i) = i(w)$.

Both conditions are satisfied. Since $\mathsf{I}(t_1)$ maps a possibility i to the ticket numbered 1 in i, it maps i_1 to α, which is the winning ticket in i_1. So $\mathsf{I}(t_1)(\mathsf{i}_1) = \mathsf{i}_1(w)$. Since $\mathsf{I}(t_2)$ maps a possibility i to the ticket numbered 2 in i, it maps i_2 to γ, which is the winning ticket in i_2. So $\mathsf{I}(t_2)(\mathsf{i}_2) = \mathsf{i}_2(w)$. Thus, s will support (4) relative to $\mathcal{C}_n$. Since s represents our epistemic state σ, the Carnapian theory predicts that we are in a position to assert (4) in the course of our conversation with Al.

We now want to show that the Carnapian theory predicts that we are in a position to *deny* (4) in our conversation with Barbara, despite the fact that our epistemic state remains unchanged. Since in the 'Barbara context' the salient way of thinking of the tickets is via their colors, the Carnapian theorist will say that this context determines the concept set $\mathcal{C}_c$, where $\mathcal{C}_c$ contains (only) the following two individual concepts:

$\mathsf{I}(b)$, which maps each world v to the ticket colored blue in v, so:

$\mathsf{I}(b)(\mathsf{v}_1) = \alpha$, $\mathsf{I}(b)(\mathsf{v}_2) = \gamma$

$\mathsf{I}(r)$, which maps each world v to the ticket colored red in v, so:

$\mathsf{I}(r)(\mathsf{v}_1) = \beta$, $\mathsf{I}(r)(\mathsf{v}_2) = \delta$

These two individual concepts are again identical to the intensions of *the blue ticket* and *the red ticket*, respectively.

Given **Fact 5.2**, we note that we can show that (4) is inconsistent with state s relative to $\mathcal{C}_c$ if we can show that there is a $c \in \mathcal{C}_c$ such that there is no possibility $i \in \mathsf{s}$ such that $c(i) = i(w)$. It is easy to see that $\mathsf{I}(r)$ is such an element of $\mathcal{C}_c$. Since we know that the red ticket lost, there is no world v compatible with what we know at v_1 at which $\mathsf{I}(r)(v)$ is the winning ticket. Since s represents our epistemic state at v_1, this means that there is no possibility $i \in \mathsf{s}$ at which $\mathsf{I}(r)(i) = i(w)$. (Note, for example, that $\mathsf{I}(r)(\mathsf{i}_1) = \beta$ which is the losing ticket in i_1, and $\mathsf{I}(r)(\mathsf{i}_2) = \delta$, which is the losing ticket in i_2.) Since (4) is inconsistent with s relative to $\mathcal{C}_c$, the Carnapian theorist predicts that we are in a position to deny (4) in the course of our conversation with Barbara.

So unlike the earlier dynamic theory, the Carnapian theory accommodates the unusual way in which sentences like (4) are sensitive to the context. It accommodates the fact that whether (4) should be asserted or denied depends not just on what one knows, but on what the contextually salient ways of thinking are.

What about our other objection to the earlier dynamic theory? Earlier, we complained that that theory predicted that the following inference is invalid:

(4) Every ticket is such that it might be the winning ticket. $\forall x(\Diamond x = w)$

So:

(5) The red ticket is such that it might be the winning ticket. $(\lambda x.\Diamond x = w)(r)$

And we also noted that, in invalidating this inference, that theory ended up predicting that sentences like (14) ought to be assertable in some circumstances:

(14) Although any ticket might win, the red ticket is such that it cannot win.

$(\forall x \Diamond x = w) \wedge (\lambda x.\neg\Diamond x = w)(r)$

Neither result is comfortable, since the above inference looks valid and (14) appears to be contradictory.

How do things look from the perspective of the Carnapian theory? On the Carnapian theory, the inference from (4) to (5) is still invalid (as is the inference from (1)-(3) to (4)). And there are still models $\mathcal{M}$, states s, and concept sets $\mathcal{C}$ such that s supports (14) relative to $\mathcal{C}$ and $\mathcal{M}$. So sentences like (14) are still not predicted to be inconsistent. So one might think that little progress has been made with respect to these issues in moving from the standard dynamic theory to the Carnapian one. But that would be a mistake. For the advocate of the Carnapian theory has a plausible account of our judgments here, an account that isn't available to the standard dynamic theorist.

With respect to the inference from (4) to (5), I emphasized earlier that the problem with the dynamic theory wasn't merely that it invalidated that inference. The problem was that the theory was not accompanied by any story about why there appears to be something compelling about that form of reasoning. While the Carnapian theory doesn't validate the inference, it *does* have the resources to account for the apparent validity of the inference. Similarly, while the Carnapian theory does not predict that sentences like (14) are inconsistent, it does have the resources to explain why they might seem so.

On the standard, informal understanding of validity, an inference is valid just in case whenever the premises are true, so is the conclusion. In our discussion of dynamic semantics, we have been assuming that *support* takes place of truth, so this becomes: whenever the premises are supported, so is the conclusion. But how should we interpret the *whenever* here? The definition of validity in **Definition 5.6** offers one answer: *whenever* means *for every dynamic model* $\mathcal{M}$*, state* s*, and concept set* $\mathcal{C}$. So in asking whether an inference is valid in this sense, we are asking how it fairs across all models, states, and concept sets.

But other interpretations are possible. One alternative would be to interpret *whenever* relative to a fixed model $\mathcal{M}$ and concept set $\mathcal{C}$. For example, suppose you are evaluating a particular inference in a certain context. If $\mathcal{M}$ is the model for your language, and $\mathcal{C}$ represents the ways of thinking salient in your context, then you might wonder whether the conclusion is supported by every state that supports the premises, holding fixed how the language is interpreted and which ways of thinking are salient in your context. This yields a different, 'local' notion of validity, one which can be characterized as follows:

Definition 5.12. For any dynamic model $\mathcal{M}$ and concept set $\mathcal{C}$ on $\mathcal{M}$, the inference from $\phi_1, ..., \phi_n$ to ψ is *valid-in-$\mathcal{C}$* relative to $\mathcal{M}$ iff for any state s, if s supports the premises relative to $\mathcal{C}$ and $\mathcal{M}$, it supports the conclusion relative to $\mathcal{C}$ and $\mathcal{M}$.

The Carnapian theorist can use this local notion of validity to account for our judgment that the inference from (4) to (5) is valid. To see how this would go, hold the dynamic lottery model M fixed as our intended model, and take note of two things. First, when someone utters or considers the inference from (4) to (5), they are very likely to be occupying a context in which a salient way of thinking about something is as *the red ticket*. Why? Because (5), the conclusion of the inference, explicitly mentions the red ticket. So if $\mathcal{C}$ is the concept set determined by such a context, $\mathcal{C}$ is likely to contain $\mathsf{I}(r)$, the intension of *the red ticket*.

The second point to appreciate is this:

Fact 5.3. *Let $\mathcal{C}$ be any concept set on M that contains $\mathsf{I}(r)$. Then the inference from (4) to (5) is valid-in-$\mathcal{C}$ relative to M.*

Proof. To see this, let s be an arbitrary discourse state on M, and $\mathcal{C}$ an arbitrary concept set on M such that $\mathcal{C}$ contains $\mathsf{I}(r)$. Suppose that s supports (4) relative to $\mathcal{C}$. Given **Fact 5.2**, it follows from this supposition that that every individual concept $c \in \mathcal{C}$ is such that there is a possibility $i \in s$ such that $c(i) = i(w)$. Since $\mathsf{I}(r)$ is in $\mathcal{C}$, it follows that there there is a possibility $i \in s$ such that $\mathsf{I}(r)(i) = i(w)$. From this, it follows that s supports (5*), $\Diamond r = w$, relative to $\mathcal{C}$. In light of **Fact 5.1**, this means that s supports (5) relative to $\mathcal{C}$. $\square$

If we put these two points together, we get the result that when someone utters or considers the inference from (4) to (5), they are very likely to be occupying a context relative to which that inference is valid, i.e. they are very likely to be occupying a context that determines a concept set $\mathcal{C}$ such that the inference is valid-in-$\mathcal{C}$. Thus, the Carnapian theorist can point to this fact as an explanation of why we tend to judge that the inference from (4) to (5) is valid: it is 'locally valid' relative to the contexts in which it tends to be uttered or considered. Perhaps we sometimes mistake this property for 'real' validity, or perhaps the ordinary notion of validity doesn't distinguish between the two notions of validity defined in **Definitions 5.6** and **5.12**, respectively.

The Carnapian theorist can say something similar about why (14) appears to be contradictory. Sentence (14) will ordinarily be uttered or considered in a context in which one of the salient ways of thinking of the red ticket will be as *the red ticket*, since, once again, that description appears explicitly in the relevant sentence. Such a context will then determine a concept set $\mathcal{C}$ that includes $\mathsf{I}(r)$. But (14) will be *inconsistent-in-$\mathcal{C}$*, in the sense that there is no state s such that (14) is consistent with s relative to $\mathcal{C}$, given that $\mathcal{C}$ contains $\mathsf{I}(r)$. This is the Carnapian's theorist explanation of why (14) sounds contradictory: uttering or considering it tends to give rise to a context relative to which it is inconsistent.

6 Relational counterpart theory

If the dynamic theory can escape its troubles by incorporating a system of contingent identity, can the relational theory do something similar? Can *it* employ individual concepts or some similar mechanism in order to resolve the problems it faces? I think it can. This is an important point, since it bears on the question of whether the issues we have been discussing in this essay provide a reason to favor dynamic semantics over a static alternative. The considerations of this section suggest that this turns out to be a rather delicate question, one I will re-visit in the next (and final) section of the paper.

I think the relational theory can avoid its problems if it adopts *something similar* to a system of contingent identity. *Counterpart theory*, a well-known semantics for quantified modal logic due to Lewis, is similar to a system of contingent identity.[35] Both approaches allow us to say that object o' in world v' represents object o in world v even if o is not identical to o'. And both approaches allow us to say that which object-in-a-world represents another object-in-a-world is something that might vary with the demands of the utterance context.

In the previous section, we took the Carnapian, rather than the Lewisian route. This decision was largely pragmatic: it is straightforward to upgrade dynamic semantics using individual concepts; it is less clear (to me) how a counterpart-theoretic version of dynamic semantics would go. When it comes to the relational theory, however, it seems to me that we could go either way, since both approaches have been developed in the literature (usually with another interpretation of the modal operators in mind). Here, I take the counterpart route since the resulting theory provides us with a sharper contrast to the Carnapian dynamic theory discussed above.

The models of *relational counterpart semantics* are just Kripke models (**Definition 2.5**).[36]

Definition 6.1. Given a Kripke model $\mathcal{M} = \langle \mathcal{W}, \mathcal{R}, \mathcal{D}, \mathcal{I}, v \rangle$, a *counterpart relation* $\mathcal{K}$ is a reflexive binary relation on $\mathcal{W} \times \mathcal{D}$.

We read "$\langle v, o \rangle \mathcal{K} \langle v', o' \rangle$" as *$o'$ in v' is a counterpart of o in v.*

The definitions of a variable assignment and the extension of a term remain the same as they were in the earlier relational theory (**Definitions 2.6** and

[35] Counterpart theory was first presented in Lewis (1968), as a first-order theory with a scheme for translating formulas of quantified modal logic into the language of that theory. Here we provide a *counterpart semantics*, using counterpart-theoretic ideas in interpreting our language $\mathcal{L}$.

Some of the counterpart-theoretic ideas that are relevant for our purposes can be found in Lewis (1971, 1986, §4.5). In those works, Lewis develops the idea that there is a multiplicity of counterpart relations, with context playing a role in determining which relations are relevant on a given occasion.

On the relationship between counterpart theory and systems of contingent identity, see Gibbard (1975, n. 3), Lewis (1986, 256, n. 40), and Aloni (2001, §4.2).

[36] For simplicity, we continue to use constant domain models, and so must treat counterpart relations as binary relations on $\mathcal{W} \times \mathcal{D}$. Lewis, of course, holds that the domains of any two possible worlds are disjoint, and so would treat a counterpart relation as a binary relation on $\mathcal{D}$.

2.7). In the counterpart semantics, we recursively define *the truth of a formula relative to a model, a world, a variable assignment, and a counterpart relation.* But, other than the clause for sentences of the form $\Diamond\phi$, the clauses of the recursive semantics are essentially identical to those of the relational semantics presented in **Definition 2.8**. Here is the new clause for the possibility operator (cf. Hughes and Cresswell 1996, 354):

> $[\![\Diamond\phi]\!]^{v,g,\mathcal{K}} = 1$ iff there is a $v' \in \mathcal{W}$ such that $v\mathcal{R}v'$ and $[\![\phi]\!]^{v',g',\mathcal{K}} = 1$, for some assignment g' such that, for each free variable x in ϕ, $\langle g(x), v\rangle\mathcal{K}\langle g'(x), v'\rangle$.

The definitions of truth and validity are amended as follows, to take account of the fact that truth is now relativized to a counterpart relation:

Definition 6.2. A formula ϕ is *true* at a model $\mathcal{M}$, world v, and counterpart relation $\mathcal{K}$ just in case $[\![\phi]\!]^{v,g,\mathcal{K}} = 1$, for all variable assignments g.

Definition 6.3. An argument is *valid* just in case for every model $\mathcal{M}$, world v, and counterpart relation $\mathcal{K}$: if the premises are true at $\mathcal{M}$, v, and $\mathcal{K}$, then so is the conclusion.

We had three complaints about the earlier relational theory (spread across §§2 and 4). First, it required us to choose between the truth of (1) and (2), on the one hand, and the falsity of (5), on the other. Second, it was bound to predict that certain obvious falsehoods were true in the lottery scenario. Third, it was unable to accommodate the fact that the truth of (4) appears to depends upon which ways of thinking are salient in the context of utterance. We take these issues up in order.

First, in the lottery scenario, it seems as though (1) and (2) are true and as though (5) is false. The counterpart theorist response to this begins by distinguishing two counterpart relations. There is the *number counterpart relation* that holds between ticket τ in world v and ticket τ' in world v' just in case the number τ has in world v is identical to the number τ' has in world v'. And there is the *color counterpart relation* that holds between ticket τ in world v and ticket τ' in world v' just in case the color τ has in world v is identical to the color τ' has in world v'. So consider, for example, worlds v_1 and v_2 in our lottery model M:

v_1	v_2
α: #1, blue, winner	γ: #2, blue, winner
β: #2, red, loser	δ: #1 red, loser

So γ in v_2 is a number counterpart of β in v_1, and δ in v_2 is a number counterpart of α in v_1. But the color counterpart relation swaps things around: γ in v_2 is a color counterpart of α in v_1, and δ in v_2 is a color counterpart of β in v_1.

Now, the counterpart theorist claims that, ordinarily, when we utter (1) or (2), the utterance context will deliver the *number* counterpart relation. Why?

Because those sentences pick out the tickets via their numbers, and this will tend to make that particular counterpart relation salient. And given that counterpart relation, (1) and (2) will both truth at v_1. For (1) is true at v_1 relative to the number counterpart relation just in case α, the referent of *ticket #1* at v_1, is such that there is a world v accessible from v_1 at which α has a number counterpart that is the winning ticket at v. And there is such a world, for v_1 is accessible from v_1, α is its own number counterpart at v_1, and α is the winning ticket at v_1. And (2) is also true v_1 relative to the number counterpart relation, since v_2 is accessible from v_1, γ is a number counterpart of β (the referent of *ticket #2* in v_1), and γ wins in v_2.

Now, (5) is also true at v_1 relative to the number counterpart relation, since *the red ticket* refers to β in v_1. Why then does it seem false in the lottery scenario? According to the counterpart theorist, this is because when we utter (5), the utterance context is likely to deliver the *color* counterpart relation, not the number counterpart relation. Why is this? Because that sentence refers to one of the tickets via its color. And given the color counterpart relation, (5) is false at v_1. This is because (5) is true at v_1 relative to the color counterpart relation just in case β, the referent of *the red ticket* in v_1, is such that that there is a world v accessible from v_1 at which β has a color counterpart that is the winning ticket. But there is no such world v.

To see this, let v be an arbitrary world compatible with what we know in v_1. Let ϵ in v be a color counterpart of β. Since ϵ is a color counterpart of β, these tickets must have the same color. Since β is red, ϵ must also be red. Since we know in v_1 that the red ticket did not win, we can conclude that ϵ is not the winning ticket in v, since v is compatible with what we know in v_1. Since ϵ was any color counterpart of β in v, this holds for every color counterpart of β in v. Since v was any world accessible from v_1, this holds for every world accessible from v_1. Thus, (5) is false at v_1 relative to the color counterpart relation.

Our second complaint about the earlier relational theory was that it was bound to make true some sentence that is obviously false in the lottery scenario. For example, it is bound to predict that one of (7) and (9) is true in that scenario:

(7) Ticket #1 is such that it might be ticket #2. $\quad (\lambda x.\Diamond x = t_2)(t_1)$

(9) The blue ticket is such that it might be the red ticket. $\quad (\lambda x.\Diamond x = r)(b)$

The counterpart theorist, however, can say the following about why (7), for example, appears to be false. Uttering sentence (7) will tend to make the number counterpart relation salient. Since ticket #1 in world v_1 has no number counterparts that are numbered 2, (7) is false at v_1 relative to the counterpart relation it tends to make salient. A similar account of (9) can be given.

Finally, we turn to the fact that the truth of (4) seems to depend on how the objects in the domain are conceptualized. When the salient way of thinking of the tickets is via their numbers, it seems that (4) is true or assertable in the lottery scenario; but when the salient way of thinking of the tickets is via their colors, it seems that (4) false or deniable in that scenario.

The counterpart semantics offers a straightforward explanation of this. A context in which the salient way of thinking of the tickets via their numbers is a context that delivers the number counterpart relation. And (4) is true at v_1 relative to that counterpart relation. For (4) is true at v_1 relative to the number counterpart relation just in case every object o in the domain D satisfies the following condition:

> there is a world v accessible from v_1 and an object o' such that $\langle v, o' \rangle$ is a number counterpart of $\langle \mathsf{v}_1, o \rangle$, and o' is the winning ticket in v.

There are only two objects in the domain, α and β, and we've seen that both of these object satisfy the condition in question – this was established when we demonstrated that sentences (1) and (2) are both true at v_1 relative to the number counterpart relation.

On the other hand, a context in which the salient way of thinking of the tickets is via their colors is a context that delivers the color counterpart relation. And (4) is false at v_1 relative to that counterpart relation. For (4) is false at v_1 relative to the color counterpart relation just in case there is an object o in the domain D that *fails* to satisfy the following condition:

> there is a world v accessible from v_1 and an object o' such that $\langle v, o' \rangle$ is a color counterpart of $\langle \mathsf{v}_1, o \rangle$, and o' is the winning ticket in v.

And we saw above – in the course of establishing that (5) is false at v_1 relative to the color counterpart relation – that β fails to satisfy this condition.

On the counterpart approach, the truth or falsity of (4) depends on, among other things, what counterpart relation is chosen. And the choice of counterpart relation depends on the context; in particular, on the ways of thinking that are salient in the context at the time of utterance. Note that our two inferences – the inference from (1)-(3) to (4) and from (4) to (5) – are still valid on this account. And (14) is still predicted to be contradictory. So these virtues of the earlier relational theory are retained by the counterpart semantics.

7 Which theory is right?

So we have two theories that are more-or-less adequate to the phenomena we've been discussing. One is static and employs counterpart relations; the other is dynamic and employs individual concepts. Which theory is to be preferred? I close by discussing a few considerations that bear on this matter.

The relational counterpart theory and the Carnapian dynamic theory differ on whether certain phenomena ought to be given a semantic or a pragmatic explanation. These differences may give us a way to probe the question of which theory is right.

Consider, for example, this exchange:

(17) *A*: Hey, treat those guys with respect! Anyone who might be innocent should be accorded due process, and any of those guys might be innocent.

> B: What about the guilty man? Should he be accorded due process?
>
> A: Yes, even the guilty man might be innocent, and so should be accorded due process.

If A's final utterance is true or assertable, then it presumably follows that the guilty man might not be the guilty man.[37]

Let us suppose for the moment that A's final utterance in (17) is acceptable. Then that might be trouble for the Carnapian theory. For that theory predicts that A's utterance is inconsistent (relative to any concept set). That utterance is, therefore, inconsistent with A's epistemic state. Thus, not only is A not in a position to assert that sentence, he is, in fact, in a position to deny it. Thus, the Carnapian dynamic theory appears to conflict with the claim that A's final utterance in (17) is acceptable.

The relational counterpart theorist has more room to maneuver here. Now it is part of the counterpart story that when one utters a sentence of the form, *The F is such that it might be G*, this action will ordinarily have the effect of bringing it about that one's utterance context delivers a counterpart relation that relates unique F's to unique F's. But this might be deemed merely a default rule, one that may be overridden by other pressures in the right context. Thus, the counterpart theorist can maintain that, to the extent that we judge A's final utterance in (17) to be true, this is because the imagined context is one in which that default rule has indeed been overridden. Rather than supplying a counterpart relation that preserves guilt and innocence across worlds, this context supplies a different counterpart relation. For example, perhaps the context delivers a counterpart relation $\mathcal{K}$ that preserves a person's name across worlds, i.e. $\mathcal{K}$ holds between $\langle v, o\rangle$ and $\langle v', o'\rangle$ iff the name o bears in v is the name o' bears in v'. If the participants to the above conversation do not know the name of the guilty man, then the man who is actually guilty may have a *name counterpart* in an accessible world who is innocent at that world. If A's final utterance occurs in a context that delivers a counterpart relation like that, then that utterance may well be true.

The story we've told so far on behalf of the counterpart theorist is incomplete. For ideally we should like some account of the principles that govern the selection of counterpart relation by context. Absent such details, the resulting theory is less predictive and less explanatory than we would like. But, the counterpart theorist might say, an incomplete theory is better than a false one. And the Carnapian theory is false, since it falsely predicts that A's final utterance in (17) is inconsistent.

How might the Carnapian dynamic theorist reply to this argument? First, she might challenge the datum on which it relies, namely the acceptability of A's

[37] Moss (2016, 152) offers a related example, but one which involves the adverb *probably* rather than the modal auxiliary *might*:

> Should I sell my ticket for next to nothing? Well, on the one hand, my ticket probably lost. But then again, every ticket probably lost the lottery, even the winning ticket.

final utterance in (17). Second, she might accept the conversational propriety of A's utterance, but deny that it tells us something of importance about the semantics of the sentence uttered. After all, we know that it is acceptable in some situations to point at a toy duck and say, "That's a duck," but it is not clear what consequences this has for the lexical semantics of the common noun *duck*. There is always a gap between the claim that a sentence is assertable in such-and-such circumstances and theoretical claims about the semantics of the sentence asserted. The Carnapian dynamic theorist might try to exploit that gap. Third, the Carnapian dynamic theorist might instead accept both the datum and its semantic significance, and then seek to accommodate the datum by modifying her theory in some way. I leave exploration of these options as a matter for future inquiry.

Another set of considerations arises in connection with a phenomenon that we have yet to discuss, but which plays an important role in Yalcin's case for dynamic semantics. Consider the following:

(18) Some person who is not infected is a person who might be infected. $\exists x(Ix \wedge \Diamond\neg Ix)$

(19) Every person who is not infected is a person who might be infected. $\forall x(Ix \rightarrow \Diamond\neg Ix)$

As Yalcin (2015, 485) observes, these sentences appear to be defective in the way that (11) and (12) are.[38] Yalcin also notes that this is not something that we should expect given a standard relational semantics, such as the one discussed in §2. The standard relational approach predicts that these sentences have coherent readings that we could know to be true. Furthermore, as Yalcin observes, standard dynamic theories – such as the one we discussed in §3 – predict that these sentences are indeed defective. The Carnapian dynamic theory makes a similar prediction.

Now there is much to be said about these sentences and close variants of them; Yalcin (2015) contains a detailed discussion of many of the issues they raise.[39] My interest here is in how these sentences might bear on the issue at hand, the question of which theory – the relational countepart theory or the Carnapian dynamic theory – is preferable. I will make two observations on this matter, observations which (unfortunately) pull in opposite directions.

The first consideration focuses on sentence (19). As Yalcin (2015, 497, n. 15) points out, the dynamic theory he discusses predicts that (19) is equivalent to the following:

(20) Every person is infected. $\forall x Ix$

The Carnapian dynamic theory issues the same prediction. But sentences of the form *Every F is G* typically presuppose that something is F (Cooper 1983, Beaver and Geurts 2013). Thus, (19) would also appear to presuppose that

[38] Groenendijk *et al.* (1997) and Aloni (2001, Ch. 3) discuss examples similar to (18).
[39] See also Klinedinst and Rothschild (2016).

someone is not infected. Thus, (19) will be inconsistent with any state that supports its presupposition, which presumably explains why it is defective.

A potential problem with this prediction, however, is that, despite there being something odd about (19), it may be assertable in certain situations. For example, Eric Swanson (p.c.) offers the following case:

(21) *A*: We have to give the treatment to everyone who might be infected – the treatment isn't harmful or expensive, and the infection is deadly.

B: What about the people who aren't infected? Should we figure out who they are first?

A: No, there isn't time. Every person who is not infected might be infected. And, again, everyone who might be infected should get the treatment.

Similar questions arise about the acceptability of this utterance, and about the semantic significance of the claim that it is acceptable. But if we waive such worries, we can again see that the flexibility of the relational counterpart theory may be an advantage here.

It may be that normally when one utters (19), the context delivers a counterpart relation that preserves one's status as *infected* or *not infected* across worlds. That would account for why (19) is usually bad, since in contexts like that (19) will again be equivalent to (20). But in contexts like (21), it may be that this default is overridden, and the context delivers a different counterpart relation instead, one relative to which (19) is not equivalent to (20). Again, this isn't a full account of the matter. We should like some explanation of the principles that govern the determination of counterpart relation by context. But even without such a story in place, we can see that some flexibility on this point might be desirable. If so, that looks like another point in favor of the relational counterpart theory, at least if there is the Carnapian dynamic theorist has no satisfactory response to this argument.

As I noted above, the above sentences can also be used to formulate a different argument, one that pulls in the other direction. Consider (18), for example. Yalcin calls formulas of this form $\neg\phi \wedge \Diamond\phi$ *epistemic contradictions.* Sentence (18) embeds $(\neg Ix \wedge \Diamond Ix)$, an *open* epistemic contradiction. It appears that sentences that embed open epistemic contradictions are problematic in some way; utterances of them are (generally speaking) infelicitous. In an earlier paper (Yalcin 2007), Yalcin noted that sentences that embed *closed* epistemic contradictions are also problematic. For example, the following both appear to be infelicitous:

(22) Suppose that it is not raining and it might be raining. $\mathcal{S}(\neg p \wedge \Diamond p)$

(23) If it is not raining and it might be raining, then... $((\neg p \wedge \Diamond p) \rightarrow ...)$

An advantage of the Carnapian dynamic theory is that it has the potential to offer a unified explanation of these facts, for it treats all epistemic contradictions, open and closed, as inconsistent. Thus, sentence (18) is itself inconsistent, which presumably explains why it is infelicitous. Similarly, the Carnapian semantics predicts that the antecedent of sentence (23) is inconsistent,[40] a fact which may play a role in explaining why that conditional strikes us as odd.

But it looks like the relational counterpart theorist will have to give two separate explanations here, one for sentences embedding open epistemic contradictions, another for sentences embedding closed epistemic contradictions. For the relational counterpart theorist's treatment of (18) will likely run something like this:

> When one utters (18), the context will normally deliver a counterpart relation that preserves one's infected/non-infected status across worlds. And relative to a counterpart relation like that, (18) is bound to be false.

While that might be an adequate explanation of the fact that (18) is generally infelicitous, there is no obvious way to extend this explanation to (23). For (23) contains no open modal formula, and so the truth-value of that sentence (at a context) is not sensitive to the choice of counterpart relation.

It's not that the relational counterpart theorist can say nothing about why (23) is bad. There are accounts in the literature of closed epistemic contradictions which are compatible with the relational counterpart framework.[41] The objection is that relational counterpart theorist will have to give one explanation of open epistemic contradictions and a rather different explanation of closed epistemic contradictions. The Carnapian dynamic semantics, on other hand, gives a single explanation of both species of epistemic contradiction. This would appear to be a point in favor of the dynamic approach.

Resolving the dispute between relational counterpart theory and Carnapian dynamic semantics – not to mention the more general dispute between static and dynamic semantics – is beyond the scope of this discussion. In fact, given the overall dialectic of this essay, this dispute can be seen as an intramural debate between members of a common side. For my main aim has been to argue that being possibly thus and so (in the epistemic sense of *possibly*) is not a trait that an object has in and of itself, but a trait that an object has only relative to a way of thinking of that object (*à la* relational counterpart theory), or a trait of a way of thinking itself (*à la* Carnapian dynamic semantics). This means that certain *de re* epistemic modal claims – claims like sentence (4) – are context-sensitive in a way that has not been widely-recognized. These are points that both approaches agree upon, and points that I hope will be incorporated into any adequate theory of these matters.

References

Aloni, M. (2001). *Quantification under Conceptual Covers*. Ph.D. thesis, University of Amsterdam.

[40] Treat p as a 0-place predicate.

[41] See, for example, Dorr and Hawthorne (2013). (Alternatively, one might combine the counterpart approach with the domain semantics of Yalcin (2007) and MacFarlane (2011); the result would be another static alternative.)

Beaver, D. I. and Geurts, B. (2013). Presupposition. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Fall 2013 edition.

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

Cooper, R. (1983). *Quantification and Syntactic Theory*. Reidel, Dordrecht.

DeRose, K. (1991). Epistemic possibilities. *Philosophical Review*, **100**(4), 581–605.

Dorr, C. and Hawthorne, J. (2013). Embedding epistemic modals. *Mind*, **122**(488), 867–913.

Dowell, J. (2011). A flexible contextualist account of epistemic modals. *Philosophers' Imprint*, **14**, 1–25.

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

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

von Fintel, K. and Gillies, A. (2008). CIA leaks. *Philosophical Review*, **117**(1), 77–98.

von Fintel, K. and Gillies, A. (2011). "Might" made right. In A. Egan and B. Weatherson, editors, *Epistemic Modality*, pages 108–130. Oxford University Press, Oxford.

Frege, G. (1892). Uber sinn und bedeutung. *Zeitschrift fur Philosophie und philosophische Kritik, C*, pages 25–50. English Translation: "On Sense and Meaning", in McGuinness, B. (ed), Frege: collected works, Oxford: Basil Blackwell, 157–177.

Gibbard, A. (1975). Contingent identity. *Journal of Philosophical Logic*, **4**, 187–221.

Groenendijk, J. and Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and Philosophy*, **14**, 39–100.

Groenendijk, J., Stokhof, M., and Veltman, F. (1997). Coreference and modality. In S. Lappin, editor, *Handbook of Contemporary Semantic Theory*, pages 179–216. Blackwell, Oxford.

Hacking, I. (1967). Possibility. *Philosophical Review*, **76**(2), 143–168.

Heim, I. (1981). *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst.

Hughes, G. E. and Cresswell, M. J. (1996). *A New Introduction to Modal Logic*. Routledge, London and New York.

Kamp, H. (1981). A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof, editors, *Formal Methods in the Study of Language*, volume 135, pages 277–322. Amsterdam.

Klinedinst, N. and Rothschild, D. (2016). Quantifiers, pronouns, and modals. Handout from a talk at the Pacific APA.

Kratzer, A. (1981). The notional category of modality. In H.-J. Eikmeyer and H. Rieser, editors, *Words, Worlds, and Contexts*, pages 38–74. Walter de Gruyter and Co., Berlin and New York.

Kratzer, A. (1991). Modality. In A. von Stechow and D. Wunderlich, editors, *Semantics: An International Handbook of Contemporary Research*, pages 639–650. Walter de Gruyter, Berlin.

Kratzer, A. (2012). *Modals and Conditionals*. Oxford University Press, Oxford.

Kripke, S. (1980). *Naming and Necessity*. Blackwell, Oxford.

Lewis, D. K. (1968). Counterpart theory and quantified modal logic. *Journal of Philosophy*, **65**, 113–126.

Lewis, D. K. (1971). Counterparts of persons and their bodies. *Journal of Philosophy*, **68**, 203–211.

Lewis, D. K. (1979). Attitudes *de dicto* and *de se*. *Philosophical Review*, **88**(4), 513–543. Reprinted in **?**, 133-159. Page references are to the 1983 reprint.

Lewis, D. K. (1983). Individuation by acquaintance and by stipulation. *Philosophical Review*, **92**(1), 3–32. Reprinted in **?**, 373-402. Page references are to the 1999 reprint.

Lewis, D. K. (1986). *On the Plurality of Worlds*. Blackwell, Oxford.

MacFarlane, J. (2011). Epistemic modals are assessment-sensitive. In A. Egan and B. Weatherson, editors, *Epistemic Modals*, pages 144–178. Oxford University Press, Oxford.

MacFarlane, J. (2014). *Assessment-Sensitivity: Relative Truth and its Applications*. Oxford University Press, Oxford.

Moss, S. (2016). Probabilistic knowledge. Unpublished book manuscript, University of Michigan.

Nouwen, R., Brasoveanu, A., van Eijck, J., and Visser, A. (2016). Dynamic semantics. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Fall 2016 edition.

Quine, W. (1953). Reference and modality. In *From a Logical Point of View*. Harvard University Press, Cambridge.

Rothschild, D. and Yalcin, S. (2015). On the dynamics of conversation. *Noûs*. Advance online publication.

Rothschild, D. and Yalcin, S. (2016). Three notions of dynamicness in language. *Linguistics and Philosophy*, **39**(4), 333–355.

Russell, B. (1905). On denoting. *Mind*, **14**(56), 479–493.

Schaffer, J. (2011). Perspective in taste predicates and epistemic modals. In A. Egan and B. Weatherson, editors, *Epistemic Modality*, pages 179–226. Oxford University Press, Oxford.

Stalnaker, R. (1978). Assertion. In P. Cole, editor, *Syntax and Semantics 9: Pragmatics*, pages 315–332. Academic Press, New York. Reprinted in **?**, 78-95. Page references are to the 1999 reprint.

Stephenson, T. (2007). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, **30**, 487 – 525.

Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, **25**(3), 221–261.

Williamson, T. (1996). Knowing and asserting. *Philosophical Review*, **105**(4), 489–523.

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press, Oxford.

Yalcin, S. (2007). Epistemic modals. *Mind*, **116**(464), 983–1026.

Yalcin, S. (2011). Nonfactualism about epistemic modality. In A. Egan and B. Weatherson, editors, *Epistemic Modality*, pages 295–332. Oxford University Press, Oxford.

Yalcin, S. (2012). Dynamic semantics. In D. G. Fara and G. Russell, editors, *Routledge Companion to Philosophy of Language*, pages 253–279. Routledge, New York.

Yalcin, S. (2015). Epistemic modality *de re*. *Ergo*, **2**(19), 475–527.

Appendix

Recall **Fact 3.3**, which says:

> *For every dynamic model $\mathcal{M}$, state s on $\mathcal{M}$, term t, variable x, and formula ϕ, if t is free for x in ϕ, then:* $s[(\lambda x.\phi)(t)] = s[\phi(t/x)]$.

Fact 5.1 is the parallel claim in the context of the Carnapian dynamic semantics (see §5). The proofs of **Facts 3.3** and **5.1** are essentially the same. Here we discuss the proof of **Fact 3.3**.

The proof is by induction on the complexity of formulas, where the *complexity* of a formula ϕ is the number of logical symbols in ϕ. In what follows, we provide the base case of the inductive proof, along with two parts of the induction step: the cases for negation and the possibility operator. The cases for conjunction, the abstraction operator, and the existential quantifier are omitted here for reasons of space. Note also that the applications of **Fact 3.3** made in the paper only involve instances of it in which ϕ is of the form $\Diamond\psi$ or $\neg\Diamond\psi$, where ψ is an atomic formula.

Base case. Let $P(t_1, ..., t_n)$ be an arbitrary atomic formula, let $\mathcal{M}$ be any dynamic model, s any state on $\mathcal{M}$, t any term, and x any variable. Since atomic formulas contain no quantifiers or abstraction operators, t is free for x in $P(t_1, ..., t_n)$. We want to show:

$$s[(\lambda x.P(t_1, ..., t_n)(t)] = s[P(t_1, ..., t_n)(t/x)]$$

For each of the t_i, let t'_i be t_i if t_i is not the variable x; let t'_i be t if t_i is the variable x. Then $P(t_1, ..., t_n)(t/x)$ is $P(t'_1, ..., t'_n)$. So it will suffice to show:

$$s[(\lambda x.P(t_1, ..., t_n)(t)] = s[P(t'_1, ..., t'_n)]$$

By the clause for atomic formulas, we have:

$$s[P(t'_1, .., t'_n)] = \{i \in s : \langle i(t'_1), ..., i(t'_n)\rangle \in i(P)\}.$$

So it will suffice to show:

$$s[(\lambda x.P(t_1, ..., t_n)(t)] = \{i \in s : \langle i(t'_1), ..., i(t'_n)\rangle \in i(P)\}.$$

Now by definition, for any possibility i, $i[x/t]$ is the possibility that is like i with the possible exception that $i[x/t](x) = i(t)$. So this means that, for any term t^* distinct from x, $i[x/t](t^*) = i(t^*)$. Thus, we have:

for any of the t_i, $i[x/t](t_i) = i(t'_i)$.

To see this, first suppose that t_i is not the variable x. In that case, t'_i is just t_i, so $i(t'_i) = i(t_i)$. Since (we are supposing) t_i is not x, $i[x/t](t_i) = i(t_i)$, which means that $i[x/t](t_i) = i(t'_i)$. Now suppose that t_i is x. In that case, t'_i is t, in which case $i(t'_i) = i(t)$ and $i[x/t](t_i) = i[x/t](x) = i(t)$, and we again have $i[x/t](t_i) = i(t'_i)$.

In light of this fact, and the clauses for the abstraction operator and atomic formulas, we have:

$$\begin{aligned}
&s[(\lambda x.P(t_1, ..., t_n)(t)] \\
&= \{i \in s : i[x/t] \in s[x/t][P(t_1, ..., t_n)]\} \\
&= \{i \in s : i[x/t] \in \{i' \in s[x/t] : \langle i'(t_1), ..., i'(t_n) \in i'(P)\}\} \\
&= \{i \in s : i[x/t] \in s[x/t] \text{ and } \langle i[x/t](t_1), ..., i[x/t](t_n)\rangle \in i[x/t](P)\} \\
&= \{i \in s : \langle i[x/t](t_1), ..., i[x/t](t_n)\rangle \in i(P)\}^{42} \\
&= \{i \in s : \langle i(t'_1), ..., i(t'_n)\rangle \in i(P)\}
\end{aligned}$$

[42] Note that for any possibility i and state s on any dynamic model $\mathcal{M}$, $i \in s$ iff $i[x/t] \in s[x/t]$.

which is what we needed to show.

Induction step. Our induction hypothesis is that the result holds for every formula with complexity less than n. We want to show that the result holds for every formula with complexity n. Here we provide only the cases for negation and the possibility operator.

Negation. First suppose ϕ is $\neg\psi$ for some formula ψ. Let s be any state on any dynamic model, let t be any term and x any variable such that t is free for x in $\neg\psi$. We want to show:

$$s[(\lambda x.\neg\psi)(t)] = s[(\neg\psi)(t/x)].$$

Since the free variable occurrences in $\neg\psi$ are just those that occur in ψ, $(\neg\psi)(t/x)$ is $\neg\psi(t/x)$. So we want to show:

$$s[(\lambda x.\neg\psi)(t)] = s[\neg\psi(t/x)].$$

And since t is free for x in $\neg\psi$, t is free for x in ψ.

Let s be any state on any dynamic model. By the clause for negation, the induction hypothesis (IH), and the clause for the abstraction operator, we have:

$$\begin{aligned}
&s[\neg\psi(t/x)]\\
&= s - s[\psi(t/x)]\\
&= s - s[(\lambda x.\psi)(t)] \quad \text{(by IH)}\\
&= s - \{i \in s : i[x/t] \in s[x/t][\psi]\}\\
&= \{i \in s : i[x/t] \notin s[x/t][\psi]\}.
\end{aligned}$$

Now, by the clauses for the abstraction operator and negation and the above result, we have

$$\begin{aligned}
&s[(\lambda x.\neg\psi)(t)]\\
&= \{i \in s : i[x/t] \in s[x/t][\neg\psi]\}\\
&= \{i \in s : i[x/t] \in (s[x/t] - s[x/t][\psi])\}\\
&= \{i \in s : i[x/t] \in s[x/t] \text{ and } i[x/t] \notin s[x/t][\psi]\}\\
&= \{i \in s : i[x/t] \notin s[x/t][\psi]\}\\
&= s[\neg\psi(t/x)]
\end{aligned}$$

which is what we needed to show.

Possibility operator. Suppose now that ϕ is $\Diamond\psi$, for some formula ψ. Let s be any state on any dynamic model, let t be any term and x any variable such that t is free for x in $\Diamond\psi$. We want to show:

$$s[(\lambda x.\Diamond\psi)(t)] = s[(\Diamond\psi)(t/x)].$$

Note that $(\Diamond\psi)(t/x)$ is $\Diamond\psi(t/x)$. So we want to show:

$$s[(\lambda x.\Diamond\psi)(t)] = s[\Diamond\psi(t/x)].$$

Note that since t is free for x in $\Diamond\psi$, t is free for x in ψ.

By the clause for the abstraction operator, we have:

$$\begin{aligned}&s[(\lambda x.\Diamond\psi)(t)]\\&=\{i\in s: i[x/t]\in s[x/t][\Diamond\psi]\}.\end{aligned}$$

By the clause for the possibility operator, we have claim (A):

$$\text{(A)}\quad s[x/t][\Diamond\psi]=\begin{cases}s[x/t] & \text{if } s[x/t][\psi]\neq\emptyset;\\ \emptyset & \text{otherwise.}\end{cases}$$

Suppose $s[x/t][\Diamond\psi]=s[x/t]$. Then by this supposition and the clause for the abstraction operator, we have:

$$\begin{aligned}&s[(\lambda x.\Diamond\psi)(t)]\\&=\{i\in s: i[x/t]\in s[x/t][\Diamond\psi]\}\\&=\{i\in s: i[x/t]\in s[x/t]\}\\&=s.\end{aligned}$$

Suppose now that $s[x/t][\Diamond\psi]=\emptyset$. Then by this supposition and the clause for the abstraction operator, we have:

$$\begin{aligned}&s[(\lambda x.\Diamond\psi)(t)]\\&=\{i\in s: i[x/t]\in s[x/t][\Diamond\psi]\}\\&=\{i\in s: i[x/t]\in\emptyset\}\\&=\emptyset.\end{aligned}$$

So we have:

$$s[(\lambda x.\Diamond\psi)(t)]=\begin{cases}s & \text{if } s[x/t][\Diamond\psi]=s[x/t];\\ \emptyset & s[x/t][\Diamond\psi]=\emptyset.\end{cases}$$

From this and claim (A), we have claim (B):

$$\text{(B)}\quad s[(\lambda x.\Diamond\psi)(t)]=\begin{cases}s & \text{if } s[x/t][\psi]\neq\emptyset;\\ \emptyset & \text{otherwise.}\end{cases}$$

Now by the clause for the possibility operator we have:

$$s[\Diamond\psi(t/x)]=\begin{cases}s & \text{if } s[\psi(t/x)]\neq\emptyset;\\ \emptyset & \text{otherwise.}\end{cases}$$

By the induction hypothesis, we have:

$$s[\Diamond\psi(t/x)]=\begin{cases}s & \text{if } s[(\lambda x.\psi)(t)]\neq\emptyset;\\ \emptyset & \text{otherwise.}\end{cases}$$

By the clause for the abstraction operator, we have:

$$s[\Diamond\psi(t/x)] = \begin{cases} s & \text{if } \{i \in s : i[x/t] \in s[x/t][\psi]\} \neq \emptyset; \\ \emptyset & \text{otherwise.} \end{cases}$$

Note the following:

$$\{i \in s : i[x/t] \in s[x/t][\psi]\} \neq \emptyset \text{ iff } s[x/t][\psi] \neq \emptyset.$$

To see that this is so, first suppose that $\{i \in s : i[x/t] \in s[x/t][\psi]\} \neq \emptyset$. Then there is an $i \in s$ such that $i[x/t] \in s[x/t][\psi]$. So $s[x/t][\psi] \neq \emptyset$. Now suppose that $s[x/t][\psi] \neq \emptyset$. So there is an $i[x/t] \in s[x/t][\psi]$. So there is an $i \in s$ such that $i[x/t] \in s[x/t][\psi]$.[43] So $\{i \in s : i[x/t] \in s[x/t][\psi]\} \neq \emptyset$.

So we can conclude:

$$s[\Diamond\psi(t/x)] = \begin{cases} s & \text{if } s[x/t][\psi] \neq \emptyset; \\ \emptyset & \text{otherwise.} \end{cases}$$

From this and claim (B), it follows that:

$$s[(\lambda x.\Diamond\psi)(t)] = s[\Diamond\psi(t/x)]$$

which is what we needed to show.

[43] Here we rely on the fact that if $i[x/t] \in s[x/t][\psi]$, $i[x/t] \in s[x/t]$. This follows from the fact that the dynamic theory in question has the *update property*: for any state s in any dynamic model, and any formula ϕ, $s[\phi] \subseteq s$.