

Quantification and Epistemic Modality*

Dilip Ninan
dilip.ninan@tufts.edu

Abstract: This paper presents a puzzle concerning the interaction of epistemic modals, singular terms, and quantifiers. The puzzle poses a number of problems for both static and dynamic theories of epistemic modals. The trouble arises because neither approach takes into account the fact 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 one that object possesses only relative to a way of thinking of the domain of quantification. I consider two theories that implement this insight: a static version of counterpart theory and a dynamic system of contingent identity. I then consider a variant on the initial puzzle that helps us to choose between the two theories. The variant also sheds light on how the phenomenon discussed in this essay relates to Frege’s Puzzle about attitude ascriptions.

Keywords: quantifiers, epistemic modals, dynamic semantics, counterpart theory, conceptual covers

1 Introduction

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 each ticket is only printed on its front, and we can only see the colored backs of the tickets.) 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.

*Earlier versions of this material were presented at a conference at MIT in honor of Robert Stalnaker’s retirement, at Institut Jean Nicod in Paris, at the University of Chicago Workshop on Subjectivity in Language and Thought, and at a conference in Lima at Pontificia Universidad Católica del Perú. For helpful comments and discussion, thanks to these audiences and to Fabrizio Cariani, Tom Donaldson, Paul Egré, Adam Elga, Itamar Francez, Melissa Fusco, Anastasia Giannakidou, Ephraim Glick, Chris Kennedy, Justin Khoo, Daniel Lassiter, John MacFarlane, Matt Mandelkern, Alda Mari, Carlotta Pavese, François Recanati, Adam Sennet, Philippe Schlenker, James Shaw, Robert Stalnaker, Shane Steinert-Threlkeld, Isidora Stojanovic, Scott Soames, Eric Swanson, Zoltán Gendler Szabó, Dan Waxman, Brian Weatherson, Malte Willer, and Seth Yalcin. Thanks also to the editors of *The Philosophical Review*, and to Sarah Moss for providing detailed comments on an earlier draft.

We now reason as follows (in what follows, ‘might’ is to be read as an epistemic modal):

(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. So it seems like I should be able to 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

The present essay uses this puzzle to investigate the meanings of quantifiers, singular terms, and epistemic modals. I begin in §2 by arguing that the puzzle poses various problems for both standard static and dynamic theories of epistemic modals. That is, the puzzle poses a problem for almost every extant theory of epistemic modals, when those theories are combined with otherwise plausible assumptions about quantifiers and variables.

In §3, I argue that the key to solving the puzzle is the following claim, which (for reasons I shall explain) I dub ‘the Quinean insight’:

Whether an object satisfies an epistemically modalized predicate in a given context c depends on how the domain of quantification in c is thought of in c .

Or, to put it in Fregean terms, whether an object is possibly thus-and-so (in the epistemic sense of ‘possibly’) depends on the *mode of presentation* under which the domain is thought of. To illustrate this point, suppose that ticket #1 is in fact the blue ticket and that ticket #2 is in fact the red ticket (though we, *qua* characters in the lottery scenario, do not know this). Let α be the former ticket, and let β be the latter ticket. Does β satisfy the predicate ‘is such that it might

¹Scenarios with this structure (though not puzzles of precisely this sort) were first discussed in Aloni 1997, 2001 and Gerbrandt 1997, 1998, with Aloni 2001 being the most extensive development. Moss (2018, §7.5) discusses related examples involving the adverb *probably*.

For other discussions of quantified epistemic modality, see Groenendijk et al. 1996, Beaver 2001, Yalcin 2015, Klinedinst and Rothschild 2016, and Mandelkern 2017, Ch. 1. Lennertz 2015 is also relevant.

be the winning ticket'? The idea behind the Quinean insight is that there just isn't a univocal answer to this question. The best we can do is to say that, in contexts in which we think of the tickets as a collection of numbered tickets, it satisfies that predicate, but in contexts in which we think of the tickets as a collection of colored tickets, it does not satisfy that predicate.

In §4, I argue that properly accommodating the Quinean insight requires adopting a non-standard theory of *transworld representation*, such as *counterpart theory* or a *contingent identity system*. I first consider a static version of counterpart theory, and then consider the dynamic system of contingent identity presented in Aloni 2001. Both theories accommodate the insight and solve our puzzle.

I then compare these two theories in §5 by considering a variant on the above lottery scenario. The variant appears to provide an argument for the counterpart-theoretic approach over any contingent identity system, Aloni's included. The variant is also interesting because it suggests that, although the phenomenon motivating the Quinean insight is very similar to Frege's Puzzle about attitude ascriptions, the two differ in important respects.

I close in §6 with some remarks about how our discussion bears on the choice between static and dynamic semantics.^{2,3}

2 Static and dynamic semantics

2.1 Static semantics

The reasoning that leads from (1)–(3) to (5) is intuitively valid – that's at least part of the reason why the puzzle *is* a puzzle.⁴ But the validity of that reasoning also follows from relatively modest assumptions about the semantics of the expressions involved. Consider, for example, the inference from (1)–(3) to (4). Let c be a context of utterance relative to which (1)–(3) are all true, and let o be an arbitrary ticket. Given (3), we know that o must be identical to either ticket #1 or ticket #2. Sentence (1) tells us that ticket #1 has a certain property, namely whatever property it is that 'is such that it might be the winning ticket' expresses at c . Sentence (2) tells us that ticket #2 has this property as well. Since o just is ticket #1 or ticket #2, o must also have this property, that is, o has the property that 'is such that it might be the winning ticket' expresses at c . Since o was an arbitrary ticket, it follows that every ticket

²In what follows, I follow Frege (1892) in treating definite descriptions as singular terms, translating them into our formal language using individual constants. But I suspect that most of the points made in this paper would go through (perhaps in a slightly altered form) were we to have followed Russell (1905) rather than Frege on this point. Note also that since parallel puzzles can be formulated using proper names rather than definite descriptions, we do not want our solution to depend on a Russellian treatment of definite descriptions.

³For simplicity, our discussion ignores any complications that might be introduced by 'exocentric' readings of epistemic modals, that is, readings of epistemic modals that do not concern the speaker's body of information.

⁴The argument from (4) to (5) is arguably enthymematic, with the suppressed premise being 'There is a unique red ticket.' I gloss over this subtlety in what follows.

has this property. From this it follows that (4) is true at c . If an inference is valid just in case it preserves *truth at a context* (Kaplan 1989), then the first inference is valid.

This argument assumes that sentences containing epistemic modals can be evaluated for truth relative to a context of utterance. While this is in some sense the orthodox view, it has been the subject of some controversy in the recent literature.⁵ Relativists and expressivists about epistemic modals, for example, join forces in denying it. Fortunately, the argument for the validity of this inference doesn't depend on this assumption. The validity of this inference – along with the validity of the inference from (4) to (5) – follows from some rather minimal assumptions about the semantics of the expressions involved, assumptions that are compatible with a wide variety of approaches to the semantics of epistemic modals.

Assume that we have a semantic theory that recursively defines a notion of truth relative to a *point of evaluation*. A point of evaluation is an n -ary sequence that might include things like a context of utterance, a possible world, a time, and so on (Kaplan 1989). A point of evaluation might also include one or more parameters needed specifically for stating the truth conditions of epistemically modalized sentences, such as a conversational background (Kratzer 1991), an accessibility relation (Ninan 2016), an information state, (Yalcin 2007, MacFarlane 2014), or a judge (Stephenson 2007). The only assumption about points of evaluation we make is that each point contains a variable assignment g . So the axioms and theorems of our theory will look something like this:

$$\llbracket \phi \rrbracket^{c,w,t,g,\dots} = 1 \text{ iff...}$$

where $\langle c, w, t, g, \dots \rangle$ is any point of evaluation.

We assume that if an inference preserves *truth relative to a point of evaluation*, then that inference is valid. Then we can establish the validity of our two inferences without making any specific assumptions about the semantics of epistemic modals. That is, the validity of those inferences follows from standard assumptions about the meanings of the relevant *non-modal vocabulary*. For standard assumptions about identity, disjunction, the abstraction operator (λx), and the universal quantifier allow us to essentially reproduce the “contextualist” reasoning given above.⁶

To see this, let e, g be any point of evaluation, where g is a variable assignment, and e consists of the other elements of the point of evaluation, whatever these might be (context, world, time, etc.). And let o be any ticket in the domain of discourse. Suppose that (1)–(3) are all true relative to e, g . Now given standard assumptions about the universal quantifier, disjunction, and the identity predicate, the truth of (3) at e, g ensures that:

$$o \text{ is either identical to } \llbracket t_1 \rrbracket^{e,g} \text{ or to } \llbracket t_2 \rrbracket^{e,g}.$$

⁵These issues are discussed in Egan et al. 2005, Egan 2007, Stephenson 2007, Yalcin 2007, 2011, von Fintel and Gillies 2008, 2011, Dowell 2011, MacFarlane 2011, 2014, and Schaffer 2011, among others. Related issues are discussed in Hacking 1967 and DeRose 1991.

⁶The relevant assumptions are made explicit in the appendix.

And given the standard semantics for the abstraction operator, λx , the truth of (1) and (2) at e, g ensures that:

both $\llbracket t_1 \rrbracket^{e,g}$ and $\llbracket t_2 \rrbracket^{e,g}$ satisfy the open sentence $\diamond x = w$ relative to e, g .⁷

From the two displayed claims, it follows that:

o satisfies the open sentence $\diamond x = w$ relative to e, g .

Since o was an arbitrary ticket in the domain, this holds for every ticket in the domain. But if every ticket in the domain satisfies the open sentence $\diamond x = w$ at e, g , then (4) is true at e, g . Since e, g was any point of evaluation, it follows that the inference from (1)–(3) to (4) is valid by the assumed criterion. The reader can verify that the inference from (4) to (5) also comes out valid given our assumptions.⁸

Again, we have not assumed anything about the semantics of epistemic modals *per se*, except that, whatever that semantics is, it can be stated within our broad framework. So the assumptions needed to establish the validity of our two inferences all concern the workings of the non-modal fragment of the language. Thus, our setup is compatible with a wide range of approaches to the semantics of epistemic modals, such as contextualism, relativism, and expressivism. Our setup is also compatible with a standard relational semantics, as well as with the “domain semantics” of Yalcin 2007 and MacFarlane 2014. In short, the problem here is a problem for any static – that is, broadly truth-conditional – theory that embraces our setup.⁹

Before considering how an advocate of this framework might try to respond to this problem, it is worth noting that standard semantic theories predict that our two inferences are valid only on the assumption that (1), (2), and (5) are (syntactically) *de re* modal predications. Neither inference remains valid according to standard theories 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 it might be’ locution, which

⁷If ϕ is an open sentence with one free variable x , then an object o satisfies ϕ at e, g just in case $\llbracket \phi \rrbracket^{e,g[x/o]} = 1$. Here and elsewhere, $g[x/o]$ is the variable assignment g' such that: (i) $g'(x) = o$, and (ii) for all variables y distinct from x , $g'(y) = g(y)$.

⁸For simplicity, I have been assuming that we have a “constant domain” semantics. The inference from (1)–(3) to (4) remains valid on a “varying domain” semantics. But the inference from (4) to (5) is not valid in the varying domain setting, since r may denote something (relative to the evaluation world) that is not in the domain of the evaluation world. But as we observed in footnote 4, the inference from (4) to (5) is arguably enthymematic, with the suppressed premise being ‘There is exactly one red ticket,’ which will entail $\exists x(x = r)$. With that premise added, the resulting inference will be valid even in the varying domain setting.

⁹The static theory of Mandelkern 2017, Ch. 1 might invalidate our two inferences, depending on what assumptions are made about how natural language definite descriptions work. But that view embraces a non-standard treatment of the logical connectives, and so is not consistent with the broad framework discussed above. The objections to dynamic semantics that I make in §§2.2–3 seem to carry over to Mandelkern’s theory; see Mandelkern 2017, 51, n. 16.

¹⁰Where a is an individual constant, a *de re* modal predication is a sentence of the form $(\lambda x.\diamond\phi)(a)$, and the *de dicto* counterpart of such a sentence is a sentence of the form $\diamond\phi(a/x)$, where $\phi(a/x)$ is the result of replacing all free occurrences of x in ϕ with a .

ensures that the relevant singular terms take wide-scope over the modal operator (Yalcin 2015, §3). Consider, for example, the “*de dicto* counterpart” of the inference from (4) to (5):

(4) Any ticket 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. $\diamond r = w$

Unlike the argument from (4) to (5), this argument is *invalid* on standard versions of quantified modal logic (constant or varying domain) that permit non-rigid terms.¹¹

Now: how should an advocate of this framework respond to our puzzle? If (1)–(3) entail (5), then we appear forced to choose between the truth of the former and the falsity of the latter. Note that sentence (3) merely says that there are no tickets other than tickets #1 and #2, a fact which is stipulated in the description of the case. So holding fixed the truth of (3), we really face a choice between accepting (1) and (2), on the one hand, and rejecting (5) on the other. 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.

The problem with the first option (accepting (5)) is that anyone who accepts the foregoing reasoning in favor of (5) would appear to be forced to accept the truth of the following sentences:

(6) The losing ticket is such that we might discover that it is the winning ticket.

(7) The red ticket is such that we might discover that it is the blue ticket.¹²

For we can easily construct parallel arguments for each of (6) and (7).¹³ If the right response to the argument for (5) is to accept that sentence, then presumably the right response to these parallel arguments is to accept their

¹¹To see this, consider a model with two worlds, v and v' , each with domain $\{\alpha, \beta\}$. Let the accessibility relation be universal. Let ‘the winning ticket’ denote α in v and β in v' , and let ‘the red ticket’ denote β in v and α in v' . Then (4) is true at world v , since each object in the domain of v wins at some world accessible from v . But (5*) is false at v , since there is no world v' accessible from v at which the ticket that is red in v' wins. A related argument shows that the *de dicto* counterpart of the inference from (1)–(3) to (4) is also invalid given standard assumptions. See Garson 2006, §13.6 for relevant discussion.

¹²Sentences (6) and (7) are similar to some examples discussed by Aloni (2001, Ch. 3) and Yalcin (2015):

(a) The biggest flea might be the smallest flea. (Aloni 2001, 104)

(b) The winner is a person who might not be the winner. (Yalcin 2015, 482)

¹³We obtain an argument for sentence (6) if we simply replace ‘it might be’ with ‘we might discover that it is’ in (1), (2), and (4). We obtain an argument for (7) if we replace ‘it might be’ with ‘we might discover that it is’ and replace the ‘the winning ticket’ with ‘the blue ticket’ in (1), (2), and (4).

conclusions. But these sentences are even more implausible than (5). How, for example, could the losing ticket be such that we might discover that it is the winning ticket? Whichever ticket is the losing ticket, and whatever course our inquiry takes, we can be sure that the losing ticket will not be discovered to be the winning ticket.

That leaves the option of rejecting (5) along with the conjunction of (1) and (2). Perhaps we could say that exactly one of (1) and (2) is true, but we don't know which.¹⁴ As a result, we don't know (1) and we don't know (2). This option is not much better. For one thing, it does nothing to avoid other implausible features of the static approach. For example, if we assume that epistemic modals are quantifiers over the possibilities compatible with what we know, any static theory of the sort we've been discussing is going to have to admit that one of the following is true in the lottery scenario:

- (6) The losing ticket is such that we might discover that it is the winning ticket.
- (8) Ticket #1 is such that we might discover that it is ticket #2.

Which of these is predicted to be true depends on what the relevant transworld identity facts are, but one or other will end up being true no matter what those facts turn out to be. To see this, note that the worlds compatible with what we know are, qualitatively speaking, of two kinds. In the following table, v_1 is an arbitrary world of the first kind, and v_2 is an arbitrary world of the second kind:

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

Table 1: Two ticket case: qualitative types

There are two possibilities for the transworld identity facts: either (i) $\alpha = \gamma$ and $\beta = \delta$, or (ii) $\alpha = \delta$ and $\beta = \gamma$. If the first possibility obtains, then (8) is true at both v_1 and v_2 ; if the second obtains, then (6) is true at both worlds.¹⁵ But surely if we know anything, we know that these two sentences are false.

¹⁴Thanks to Fabrizio Cariani and Carlotta Pavese for pressing me to think about this possibility.

¹⁵To see, for example, that if the first possibility obtains, then (8) is true at v_1 , note that (8) is true at v_1 just in case the referent of 'ticket #1' at v_1 is such that there is a world accessible from v_1 at which it is the referent of 'ticket #2'. Since the referent of 'ticket #1' at v_1 is α , (8) will be true at v_1 just in case there is a world accessible from v_1 at which α is the referent of 'ticket #2'. If the first possibility obtains, then $\alpha = \gamma$, and so (8) will be true at v_1 just in case there is a world accessible from v_1 at which γ is the referent of 'ticket #2'. Since γ is the referent of 'ticket #2' at v_2 , and since v_2 is accessible from v_1 , it follows that if the first possibility obtains, (8) is true at v_1 .

There is one further objection to static semantics that I will consider, but I postpone discussion of it until §3, since it is also a problem for the theory to be discussed in §2.2. But what we’ve said thus far should be enough to motivate looking elsewhere for a solution to our puzzle.

2.2 Dynamic semantics

While static semantics continues to be the dominant approach to meaning in semantics and the philosophy of language, much recent work on epistemic modality has focussed instead on *dynamic semantics*.¹⁶ In dynamic semantics, sentences cannot in general be evaluated for truth with respect to a point of evaluation, at least not if a point of evaluation contains only the usual suspects: context, world, time, variable assignment, etc.. Instead of placing a condition on a point of evaluation, the meaning of a sentence in dynamic semantics is understood as its capacity to update a *state of information*. The meaning of a sentence is a “context change potential” or a function from states of information to states of information.

Yalcin (2015) argues in favor of a dynamic approach to quantifiers and epistemic modals on the basis of (among other things) sentences similar to (6). Here is one of Yalcin’s examples:

(9) The winner is a person who might not be the winner. (Yalcin 2015, 482)

$$(\lambda x. \diamond x \neq w)(w)$$

As Yalcin points out, if ‘the winner’ is not rigid over the set of possibilities that ‘might’ quantifies over, then standard relational semantic theories predict that (9) will be true whenever we do not know the identity of the winner. In contrast, dynamic theories will tend to predict that (9) is incoherent, as it arguably is.

Exactly how a dynamic theory yields this prediction depends somewhat on the precise details of the dynamic theory in question. In what follows, we consider a dynamic theory that yields that prediction by making a *de re* modal predication like (9) equivalent to its *de dicto* counterpart:

(10) It might be that the winner is not the winner.

$$\diamond w \neq w$$

Since the latter is obviously contradictory, this will predict that (9) is contradictory as well.

This approach has interesting consequences for our puzzle. For a theory that makes a *de re* modal predication equivalent to its *de dicto* counterpart will likely *invalidate* our two inferences, given our earlier observation that the *de*

¹⁶Dynamic semantics was first introduced in Kamp 1981 and Heim 1982, with Karttunen 1969 and Stalnaker 1974, 1978 providing some important conceptual background. Veltman (1996) presented the first dynamic theory of epistemic modals (though see Stalnaker 1970 for an important precedent). Veltman’s approach has been developed in the subsequent literature; see, for example, Groenendijk et al. 1996, Aloni 2001, Beaver 2001, Willer 2013, and Yalcin 2012, 2015.

dicto counterparts of those inferences fail in standard static frameworks (§2.1). And this prediction is indeed borne out: our two inferences are invalid according to the dynamic theory we shall consider. This allows the dynamic approach to avoid the problems facing the static approach that we’ve just considered. But, as we shall see, this apparent virtue of dynamic semantics leads it into problems of its own.

The particular dynamic theory I have in mind is modeled on the theory of Groenendijk et al. (1996), which (among other things) shows how to add quantifiers to Veltman’s dynamic theory of epistemic modals (Veltman 1996). For the sake of simplicity, I will simply sketch the outlines of a modified version of the theory of Groenendijk et al., and then state some facts about the theory that are relevant to our puzzle. A more thorough presentation of the theory can be found in the appendix.

In dynamic semantics, a *state of information* is typically represented by a set of possibilities, where a *possibility* is a pair of a possible world and a variable assignment. A state of information s is said to *support* a formula ϕ just in case updating s with ϕ simply returns s : in symbols, $s[\phi] = s$. Where ϕ contains no modal operators, a state of information will support ϕ just in case ϕ is true at every possibility in the state.¹⁷ I will assume that an agent’s state of information – the set of possible worlds compatible with what he or she knows – supports a formula ϕ , then she is (epistemically speaking) in a position to assert ϕ .¹⁸ An inference is *valid* just in case any state that supports the premises also supports the conclusion.

Three facts are relevant for understanding what dynamic semantics says about our puzzle. The first concerns the interpretation of modal formulas. In dynamic semantics, a modal formula $\Diamond\phi$ “tests” a body of information s for compatibility with ϕ . If s is compatible with ϕ , s passes the test, and the updating procedure returns s unchanged; if s is not compatible with ϕ , s fails the test, and the updating procedure “crashes” the context and returns the empty set:

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

Where ϕ contains no modal operators, this procedure simply amounts to checking whether there is a possibility in s at which ϕ is true. Consider, in particular, a formula of the form $\Diamond a = w$, where a is any individual constant. Updating

¹⁷Although sentences cannot in general be evaluated for truth with respect to a possibility, sentences drawn from the non-modal fragment of the language can be so evaluated (in the particular dynamic theory under discussion).

¹⁸Strictly speaking, the notion of *support* is defined as a relation between a formula and a set of *possibilities*, which are *pairs* of a possible world and a variable assignment. But we can define a corresponding relation between sets of *possible worlds* and formulas as follows: a set σ of worlds *supports* a formula ϕ just in case there is a variable assignment g such that $\{\langle v, g \rangle : v \in \sigma\}$ supports ϕ . (One shouldn’t read much significance into the fact that we *existentially* quantify over variable assignments in this definition, since all of the sentences in which we are interested are closed.)

a state of information s with this formula will return s if there is a possibility i in s at which the extension of a at i is the winning ticket at i ; otherwise, it returns the empty set:

$$\mathbf{Fact\ 1.} \quad s[\Diamond a = w] = \begin{cases} s & \text{if there is an } i \in s \text{ such that } i(a) = i(w); \\ \emptyset & \text{otherwise.} \end{cases}$$

(For any individual constant a and any possibility i , $i(a)$ is the extension of a at i .)

The second fact concerns the relationship between *de re* modal predications and their *de dicto* counterparts. Our version of dynamic semantics makes this connection as tight as possible, holding that a *de re* modal predication is equivalent to its *de dicto* counterpart, in the sense that they are associated with the same update potential:

$$\mathbf{Fact\ 2.} \quad \text{For any state } s, s[(\lambda x.\Diamond\phi)(a)] = s[\Diamond\phi(a/x)].^{19}$$

This means, for example, that Yalcin’s sentence (9), $(\lambda x.\Diamond x \neq w)(w)$ is equivalent to its *de dicto* counterpart $\Diamond w \neq w$. Since the latter is inconsistent (in the technical sense that updating any state with it will yield the empty set), so is the former.²⁰

Now, given **Facts 1** and **2**, the dynamic theory predicts that, in the lottery scenario, we are in a position to assert (1) and (2). To see this, let s_l be the set of possibilities compatible with what we know in the lottery scenario. Then since it is compatible with what we know in that scenario that ticket #1 is the winning ticket, s_l contains a possibility i such that $i(t_1) = i(w)$. From this and **Fact 1**, it follows that s_l supports $\Diamond t_1 = w$. And from this and **Fact 2**, it follows that s_l supports (1). A parallel argument shows that s_l supports (2),

¹⁹Groenendijk et al. (1996) do not define their semantics over a language that contains an abstraction operator, and so we have had to extend their approach to such a language. The approach to this matter taken here (see the appendix for details) yields the equivalence stated above. But an alternative approach is available: one can treat the abstraction operator as a defined symbol, so that $(\lambda x.\Diamond\phi)(a)$ abbreviates $\exists x(x = a \wedge \Diamond\phi)$ (compare Aloni 2001, Ch. 3). If this option is taken, a *de re* modal predication will entail its *de dicto* counterpart, but the reverse will not be true (so **Fact 2** fails). The distinction between these approaches has some bearing on the issues being discussed in this essay, since the approach I have chosen to adopt invalidates *both* of our two inferences, the inference from (1)–(3) to (4) and the inference from (4) to (5). On the alternative approach, on the other hand, the former inference comes out *valid*, though the latter inference remains invalid. My choice of approach was dictated by two considerations. First, in order to wring predictions concerning assertability out of the alternative approach, one needs to make decisions about the relevant transworld identity facts, decisions which look arbitrary in the present context. Second, my objections to the dynamic theory concern the inference from (4) to (5), and so those objections go through no matter which of these two accounts is adopted.

²⁰Yalcin’s own approach is slightly different. He treats sentences containing definite descriptions along Russellian lines, and does not employ a formal language that contains an abstraction operator. He speculates that a Fregean approach to definite descriptions is incompatible with taking a dynamic approach to the infelicity of sentences like (9) (Yalcin 2015, 508, 518). The above result seems to show otherwise.

given that it is compatible with what we know that ticket #2 is the winning ticket.

Since we know that the red ticket is the losing ticket, there is no possibility i in our information state s_l such that $i(r) = i(w)$, that is, no possibility at which the red ticket is the winning ticket. So given **Fact 1**, it follows that $s_l[\diamond r = w] = \emptyset$. From **Fact 2**, it follows that $s_l[(\lambda x.\diamond x = w)(r)] = \emptyset$. So our state of information in the lottery scenario does not support (5). In fact, we can say something stronger: s_l supports the *negation* of (5). To see this, note that updates of negated formulas are defined as follows:

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

Since $s_l[(\lambda x.\diamond x = w)(r)] = \emptyset$, it follows that $s_l[\neg(\lambda x.\diamond x = w)(r)] = s_l$. In other words, our information state supports the negation of (5), which means we are in a position to *deny* (assert the negation of) (5).

And as one would expect, dynamic semantics predicts that s_l supports (3), given that we know that ticket #1 and ticket #2 are the only tickets. So unlike static theories, the dynamic approach predicts that we are in position to assert (1)–(3) and deny (5). So dynamic semantics avoids our principal objection to static theories. All this is good news. Now for the bad news.

The first piece of bad news for dynamic semantics concerns what it says about the validity of our two inferences. Since an inference is valid just in case any state that supports the premises supports the conclusion, it follows from the fact that s_l supports (1)–(3) but not (5) that the inference from the former to the latter is invalid. That means that either the inference from (1)–(3) to (4) fails or the inference from (4) to (5) fails. On the present version of dynamic semantics, both inferences fail, though I will focus in what follows on the inference from (4) to (5).²¹

To appreciate why the inference in question is predicted to be invalid, we first need to say something about the dynamic account of the quantified sentence (4) ('Any ticket might be the winning ticket'). Given a domain of objects \mathcal{D} , dynamic semantics associates (4) with the following update:

Fact 3.

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

So a state of information s supports (4) just in case, for every object o in the domain, there is a possibility $i \in s$ at which o is the winning ticket.

To obtain a counterexample to the inference from (4) to (5), let \mathcal{D} be a domain whose only elements are our two tickets, α and β . And consider a state s that consists of two possibilities, i and i' , which can be depicted as follows:

²¹See footnote 19.

i	i'
α : #1, blue, winner	α : #1, red, loser
β : #2, red, loser	β : #2, blue, winner

Table 2: Two ticket case: rigid numbers

Given **Fact 3**, we can see that s supports (4), since α wins in i , β wins in i' , and α and β exhaust the domain. But since the red ticket wins at no possibility in s , it follows that s does not support $\diamond r = w$. Given **Fact 2**, this means that s does not support (5). So the inference from (4) to (5) is invalid on this approach. (As I noted earlier, this is not surprising given that: (i) the dynamic approach makes a *de re* modal predication equivalent to its *de dicto* counterpart, and (ii) the inference from (4) to the *de dicto* counterpart of (5) fails in standard semantic systems (§2.1).)

But why should we think that this is bad news? After all, one might think that this is simply the cost of being able to accept (1) and (2) while also rejecting (5). One reason to think that this is bad news is simply that the inference in question really does seem valid, and dynamic semantics offers no explanation of this fact. Just think about that inference again. If *any* ticket might be the winning ticket, then surely the red ticket might be the winning ticket, given that the red ticket is a ticket. Conversely, if we agree that the red ticket can't be the winning ticket, then surely it is not the case that *any* ticket might be the winning ticket, for the red ticket must be a counterexample to that generalization. If we are told that these forms of reasoning are faulty, we are owed some explanation of why they nevertheless appear so compelling.

Now, one might again complain that I am simply demanding too much. After all (one might think) the following can't all be true:

- Sentences (1)–(3) are assertable in the lottery scenario.
- The negation of (5) is assertable in the lottery scenario.
- Sentences (1)–(3) entail (4), and (4) entails (5).

But note that I am not demanding that a theory make good on all of these claims. In asking for an explanation of the *apparent* validity of an inference, I am not demanding that the explanation in question take the form of a *validation*: the explanation needn't predict that that the inference in question really is valid. For example: we shall, in due course, see an alternative dynamic theory that (like the present dynamic theory) also predicts that the inference from (4) to (5) is invalid, but which also offers an explanation for why that inference nevertheless *seems* valid (§4.2). The problem with the present dynamic theory is not simply that it predicts that the inference is invalid, but that it comes with no backup story about why it appears 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 are often used to motivate dynamic theories of this sort (for example, the infelicity of (6)), this fact ought to give its advocates pause.

Consider the following sentences:

- (11) Although any ticket might be the winning ticket, its false that the red ticket might be the winning ticket. $(\forall x \diamond x = w) \wedge \neg((\lambda x. \diamond x = w)(r))$
- (12) Although the blue ticket must be the winning ticket, no ticket is such that it must be the winning ticket. $(\lambda x. \Box x = w)(b) \wedge (\neg \exists x \Box x = w)$

Each of these appears to be contradictory. In each case, the second conjunct seems to be at odds with what the first conjunct claims. In dynamic semantics, a sentence ϕ can be said to be *inconsistent* just in case for every state s , $s[\phi] = \emptyset$. Given this definition, neither (11) nor (12) is inconsistent according to dynamic semantics. Indeed, the state of information s that we used to demonstrate the invalidity of the (4)–(5) inference – the one depicted in Table 2 above – supports them both.²² That suggests that (11) and (12) ought to be assertable in certain circumstances. That they are not is a problem for dynamic semantics.

It is worth noting that, despite all their problems, the static theories discussed earlier predict that sentences (11) and (12) are contradictions, as indeed they seem to be.

We seem to be stuck. Static semantics fails because it validates our two inferences. Dynamic semantics runs into trouble because it *invalidates* them. But since a theory must do one or the other, it might seem as though *any* theory is going to be impaled on one of the horns of this dilemma.

3 Context and the Quinean insight

I propose that we shift our gaze slightly for a moment, and examine a feature of quantified sentences like (4), ‘Any ticket might be the winning ticket’, that we have yet to discuss. As we shall see, getting clearer about such sentences will help us to see where the solution to our puzzle lies.

We have been considering the truth/assertability of (1)–(3) and of (5), and we have been considering the validity of our two inferences, both of which involve (4). But we have yet to ask directly about the truth/assertability of (4). Is that sentence true/assertable in the lottery scenario? Is every ticket such that it might be the winning ticket? As Aloni (2001, Ch. 3) and Moss (2018, §7.5)

²²Note that $\Box\phi$ is defined as $\neg\diamond\neg\phi$, and that the semantics of conjunction runs like this:

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

argue, the right answer to our question seems to be: it depends. It depends on which *way of thinking* or *mode of presentation* of the tickets is salient in the context in which this question is posed.

A (somewhat contrived) example brings this out. Imagine that in the lottery scenario we have two photographs, each of which depicts 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. Suppose now that our friend 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, ‘Is it true that *any* of these tickets might 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), ‘Any ticket might be the winning ticket,’ 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, ‘Is it true that *any* of these tickets might 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.

This feature of (4) poses a problem for both the static and dynamic theories we have been discussing, but it also points the way to a solution to our puzzle.²³ To see the problem, note that while both standard static and dynamic theories can accommodate two ways in which (4) might depend on the context, 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 (see Moss 2018, §7.5). Second, both theories can make the truth or assertability of (4) sensitive to the epistemic state relevant in the context. Both theories may allow that when ‘might’ is used in a particular context, it is interpreted as quantifying over an epistemic state that is determined by that context. For example, on a static contextualist semantics, the context might accomplish this by contributing an accessibility relation. On the dynamic theory, the context can be seen as providing the state of information relevant for assessing the assertability of sentences.

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 *ticket #2*. In the second context, the salient way of thinking of them is via

²³Moss (2018, 158, n. 14) points out that the context-sensitivity of sentences like (4) poses a difficulty for the dynamic theory of Yalcin 2015, a theory which is relevantly similar to the dynamic theory we’ve been discussing.

their colors – as *the red ticket and the blue ticket*. But merely changing which way of thinking of the tickets is salient in the context 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 theory predicts.

The relevant sort of context-sensitivity exhibited by (4) appears to support the observation that I called the ‘Quinean insight’ in §1:

Whether an object satisfies an epistemically modalized predicate in a given context c depends on how the domain of quantification in c is thought of in c .

I called this the ‘Quinean insight’ because it is reminiscent of Quine's remark 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 something like this is correct when the epistemic sense of ‘necessarily’ and ‘possibly’ is at issue.²⁴

To see how the context-sensitivity of (4) supports the Quinean insight, suppose, for the sake of concreteness (and without loss of generality), that ticket #1 is in fact the blue ticket, and that ticket #2 is in fact the red ticket, though we, *qua* characters in the lottery scenario, do not know these facts. Let α be ticket #1 (the blue ticket) and let β be ticket #2 (the red ticket). In that case, (4) is assertable at a context c just in case both α and β satisfy the predicate ‘is such that it might be the winning ticket’ in c . Since (4) is assertable in the “Al-context,” both α and β must satisfy that predicate in the Al-context. Since (4) is not assertable in the “Barbara-context,” at least one of α and β must fail to satisfy that predicate in the Barbara-context. But since the only relevant difference between those two contexts is how the tickets are being thought of, it must that: (i) both α and β satisfy that predicate when the collection of tickets is thought of as a collection of *numbered* tickets, and (ii) one of α and β fails to satisfy that predicate when the collection of tickets is thought of as a collection of *colored* tickets. And presumably it is β that fails to satisfy the predicate in question when we think of the tickets via their colors. For when we are thinking of the tickets via their colors, it is the red ticket that appears to falsify (4), and β is the red ticket. Thus, it seems that whether β satisfies the predicate ‘is such that it might be the winning ticket’ depends on how we are thinking of tickets.

As I shall argue in §4, the Quinean insight is the key to resolving the lottery puzzle. But before I do that, a few remarks about the observation are in order. The Quinean insight by itself is perhaps not so surprising, and has been discussed in the literature before:

²⁴For a useful discussion of Quine's views, see Burgess 1997.

‘...intensional properties, such as *perhaps being the culprit*, do not properly apply to individuals *simpliciter*, but depend on the perspective under which these individuals are conceived.’ (Aloni 2001, 105)

‘...a thing does not have epistemically modal properties in abstraction from the way that it is specified.’ (Yalcin 2015, 519)

But it seems to me that there is still more to say about this idea. Indeed, it seems to me that neither the precise nature of the Quinean insight nor its theoretical import has been properly understood. Two points stand out.

First, note that I have formulated the Quinean insight by saying that whether an object satisfies an epistemically modalized predicate depends on how the *domain of quantification* is thought of. In contrast, the other authors (Quine included), take the observation to be that whether an object is possibly thus-and-so depends on how *the object* is thought of. This might seem like a small difference, and for the moment this difference doesn’t actually matter that much. But it turns out that it does make a difference in situations in which we have two ways of thinking of the domain, but lack appropriate modes of presentation for the objects in the domain. One such situation is considered in §5.

Second, *contra* Yalcin 2015, standard dynamic theories of the sort discussed in §2.2 do not adequately account for the Quinean observation, as Moss (2018, 158, n. 14) points out. As we noted earlier, standard dynamic semantics doesn’t account for the fact that whether (4) is assertable depends on how the domain of quantification is thought of. Furthermore, recall the dynamic account of (4):

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

In the context of dynamic semantics, we can say that an object *o* satisfies the open sentence $\diamond x = w$ relative to a state of information *s* just in case there is an *i* \in *s* such that $o = i(w)$. Then the dynamic account says that (4) is supported by a state *s* just in case every object in the domain satisfies the open sentence $\diamond x = w$. But this way of putting it presumes that whether an object satisfies that open sentence can be settled independently of how the domain is thought of. But this is just what we’ve been denying. So when we look in detail at the dynamic account of quantified sentences like (4), we see that it doesn’t fully accommodate the Quinean insight.²⁵

²⁵Having said that, one reason for thinking that standard dynamic semantics *does* accommodate the Quinean insight, at least to some extent, is the fact that it predicts that epistemic *de re* modal predications create *opaque contexts* even for the term occurring outside the scope of the modal. To see this, suppose again that ticket #2 is the red ticket. In that case, (5) can be obtained from (2) by replacing ‘ticket #2’ with the co-designating term ‘the red ticket,’ and this substitution takes places outside of the scope of the modal. Nevertheless, according to the dynamic theory, our state of information in the lottery scenario *s_l* supports (2), but does not support (5) – indeed, it supports the negation of (5). This is presumably what opacity amounts to in the dynamic setting.

In contrast, my view is that properly accommodating the Quinean insight requires adopting a non-standard theory of *transworld representation*, and that this is so whether one favors static semantics or dynamic semantics. This idea is explored in the remainder of the paper. In §4.1, I examine a static version of counterpart theory, showing how it accommodates the Quinean insight and helps to solve the lottery puzzle. In §4.2, I consider the dynamic contingent identity system developed by Aloni (2001, Ch. 3), arguing that it also has the resources to deal with our puzzle. I then compare these two theories in §5.

4 Counterparts and conceptual covers

4.1 Static counterpart theory

The connection between the Quinean insight and the issue of transworld representation arises as follows. Consider again β , the object that we are supposing is actually picked out by ‘ticket #2’ and ‘the red ticket.’ Given that ‘might’ is understood as a quantifier over possibilities, the question of whether β satisfies the predicate ‘is such that it might be the winning ticket’ reduces to the following question:

- (a) Is there a world compatible with what we know which represents β as being the winning ticket?

Now the answer to this question is presumably intimately connected to the answer to another question:

- (b) If v is compatible with what we know, which object in v represents β there?

In my view, part of the problem with standard static and dynamic theories is that the way they answer (b) entails that there is a univocal answer to (a), that is, an answer to (a) that is independent of how the tickets are thought of. For the answer they give to (b) is: β itself is the ticket that represents β in any world compatible with what we know. So if σ is the set of worlds compatible with what we know, σ either simply contains a world in which β is the winning ticket or it fails to contain such a world. And this, in turn, means that σ either contains a world that represents β as being the winning ticket, or it fails to contain such a world. Thus, (a) is assumed to have a univocal answer, an answer that obtains independently of how the tickets are being thought of.

Since we want the answer to (a) to vary depending on how we are thinking of the tickets, I propose that we adopt an account of transworld representation that also allows the answer to (b) to vary depending on how we are thinking of the tickets. *Counterpart theory* offers a simple way of doing this.²⁶

But despite this prediction about opacity, it becomes apparent that the dynamic theory doesn’t *fully* accommodate the Quinean insight when we examine quantified sentences like (4).

²⁶Counterpart theory was first introduced in Lewis 1968. Relevant for us is Lewis’s later idea that there is a multiplicity of counterpart relations, with context playing a role in determining which relations are relevant on a given occasion (Lewis 1971, 1986). Santorio

According to counterpart theory, the object that represents β in another world need not be identical to β itself. Rather, it need only be a *counterpart* of β , that is, an object that is similar to β in a certain way, with the relevant sort of similarity being determined by the utterance context. If we assume that how the domain of quantification is being thought of is the relevant feature of the utterance context that determines how counterparts are selected, we will have an account that allows the answer to (b) to vary depending on how we are thinking of the tickets. And this, in turn, will allow the answer to (a) to vary depending on how we are thinking of the tickets.

Consider, for example, a context in which we are thinking of the tickets via their numbers. In a context like that, the dimension of similarity relevant for choosing counterparts is *having the same number as*. Thus, we assume that such a context will deliver the *number counterpart relation*, where a ticket o in a world v is a number counterpart of o' in v' just in case, for any number n , o is numbered n in v just in case o' is numbered n in v' . Since β is numbered 2 in the actual world, each of β 's number counterparts will be numbered 2 in their respective worlds. Thus, in a context like this, the answer to question (b) is: the ticket numbered 2 in v represents β in v . Furthermore, since we don't know the number of the winning ticket, there will be a world v compatible with what we know at which the ticket numbered 2 in v will be the winning ticket in v . Thus, in a context like this – a context in which we are thinking of the tickets via their numbers – such a world v will represent β as being the winning ticket there. And this, in turn, means that β will satisfy the predicate 'is such that it might be the winning ticket' in contexts like this.

But in a context in which we are thinking of the tickets via their colors, the context will deliver the *color counterpart relation*, a counterpart relation that preserves a ticket's color, as opposed to its number, across worlds (a ticket o in a world v is a color counterpart of o' in v' just in case the color of o in v is identical to the color of o' in v'). This is because in a context in which we are thinking of tickets via their colors, the dimension of similarity relevant for choosing counterparts is *having the same color as*. Thus, in a context like this, the answer to question (b) is: the ticket colored red in v represents β in v (since β is the red ticket in the actual world). And since we know that the red ticket lost, there will be no world v compatible with what we know at which the ticket colored red in v is the winning ticket in v . Thus, in a context like this – a context in which we are thinking of the tickets via their colors – no world v will represent β as being the winning ticket there, which means that β will *not* satisfy the predicate 'is such that it might be the winning ticket' in contexts like this.

This more or less immediately accounts for the fact that the truth of sentence (4) varies according to how we are thinking of the tickets. In our conversation with Al, the tickets are being thought of via their numbers, since we are attending to the photograph of the numbered fronts of the tickets. We can assume,

(2012) uses counterpart relations in his theory of epistemic modals, though the constructions he considers are all syntactically *de dicto*.

then, that this context delivers the number counterpart relation. Since each of our two tickets, α and β , has a number counterpart that wins in some world compatible with what we know, (4) is true in the context of our conversation with Al. But in our conversation with Barbara, the tickets are being thought of via their colors, since we are attending to the photograph of the colored backs of the tickets. We can assume, then, that this context delivers the color counterpart relation. But it is *not* the case that each ticket has a color counterpart that wins in some world compatible with what we know. For as we just observed, β has no such color counterpart. Thus, (4) is predicted to be false in the context of our conversation with Barbara.

What does the counterpart approach say about the rest of our puzzle? To answer this, it will help to flesh out the story a bit more. A *model* in counterpart semantics is a triple $\mathcal{M} = \langle \mathcal{W}, \mathcal{D}, \mathcal{I} \rangle$ consisting of a non-empty set \mathcal{W} of worlds, a non-empty domain \mathcal{D} of objects, and interpretation function \mathcal{I} that assigns appropriate intensions to the non-logical symbols of the language.²⁷ An *accessibility relation* \mathcal{R} on \mathcal{M} is a reflexive binary relation on \mathcal{W} .²⁸ A *counterpart relation* \mathcal{K} is a reflexive binary relation on $\mathcal{W} \times \mathcal{D}$. The counterpart semantics takes the form of a recursive definition of truth relative to a model and a point of evaluation, where a point of evaluation consists of a world, a variable assignment, an accessibility relation, and a counterpart relation. The accessibility relation and counterpart relation can be thought of as being determined by the context of utterance. The recursive clauses are all standard, save for the clause for the modal operator, which runs as follows (compare Hughes and Cresswell 1996, 354):

$$\begin{aligned} \llbracket \Diamond \phi \rrbracket^{v,g,\mathcal{R},\mathcal{K}} = 1 \text{ iff there is a } v' \in \mathcal{W} \text{ such that } v\mathcal{R}v' \text{ and } \llbracket \phi \rrbracket^{v',g',\mathcal{R},\mathcal{K}} = 1, \\ \text{for some assignment } g' \text{ such that, for each free variable } x \text{ in } \phi, \\ \langle g(x), v \rangle \mathcal{K} \langle g'(x), v' \rangle. \end{aligned} \quad ^{29}$$

We assume that an argument is valid just in case it preserves *truth relative to a model and a point of evaluation*.

Now as I mentioned in §2.1, any semantics that makes certain standard assumptions about the non-modal fragment of the language will predict that our two inferences – the inference from (1)–(3) to (4) and the inference from (4) to (5) – are valid. Since what is distinctive about counterpart theory is its treatment of the modal operator, the counterpart semantics, as I am envisioning it, validates all of the relevant assumptions about the non-modal fragment of

²⁷For more details on counterpart semantics, see the appendix.

²⁸Since the relevant modality is epistemic, we restrict our attention to reflexive accessibility relations, reflecting the fact that what is known is true.

²⁹Although we've adopted a relational semantics here, the counterpart approach is easily adapted to other static theories of epistemic modality. For example, to obtain a counterpart-theoretic version of the domain semantics of Yalcin 2007 and MacFarlane 2014, we simply substitute a state of information s in for the accessibility relation \mathcal{R} in our points of evaluation, and then formulate the semantic clause for the modal operator as follows:

$$\begin{aligned} \llbracket \Diamond \phi \rrbracket^{v,g,s,\mathcal{K}} = 1 \text{ iff there is a } v' \in \mathcal{W} \text{ such that } v' \in s \text{ and } \llbracket \phi \rrbracket^{v',g',s,\mathcal{K}} = 1, \text{ for} \\ \text{some assignment } g' \text{ such that, for each free variable } x \text{ in } \phi, \langle g(x), v \rangle \mathcal{K} \langle g'(x), v' \rangle. \end{aligned}$$

the language. Thus, the counterpart approach predicts that our two inferences are valid. So why, then, does it seem that (1)–(3) are true, while (5) is false?

The counterpart explanation of this adverts to the two counterpart relations we distinguished earlier, the number counterpart relation and the color counterpart relation. 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.³⁰ Since both of our two tickets have number counterparts that win in some possible world compatible with what we know, (1) and (2) will both be true in such a context.

Now (5) will also be true relative to a context that delivers the number counterpart relation. 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 not likely to deliver the number counterpart relation; instead, it is much more likely to deliver the *color* counterpart relation. Why is this? Because that sentence refers to one of the tickets via its color. And since the red ticket has no color counterparts that win in a possible world compatible with what we know, (5) will be false at such a context. Thus, (1) and (2) seem true in the lottery scenario because they are true in the contexts in which they are most naturally evaluated; (5) seems false in the lottery scenario because it is false in the context in which it is most naturally evaluated.

Thus, counterpart theory offers a simple and pleasing solution to the lottery puzzle. Note also that, like the earlier static theory and unlike the earlier dynamic theory, the counterpart approach predicts that sentences like (11) and (12) are contradictory, for there is no model and point of evaluation relative to which they are true.

(11) Although any ticket might be the winning ticket, its false that the red ticket might be the winning ticket. $(\forall x \diamond x = w) \wedge \neg((\lambda x. \diamond x = w)(r))$

(12) Although the blue ticket must be the winning ticket, no ticket is such that it must be the winning ticket. $(\lambda x. \Box x = w)(b) \wedge (\neg \exists x \Box x = w)$ ³¹

4.2 Dynamic conceptual covers

I want now to consider what I take to be the most salient alternative to the counterpart approach: the dynamic system of *contingent identity* developed in

³⁰Moss (2018, §7.5) argues that linguistic material can affect the contextually-salient ways of thinking of something, thus affecting the interpretation of sentences containing epistemically modalized predicates. She does not couch the observation in terms of counterpart relations, however.

³¹Lewis’s original version of counterpart theory faces a number of problems (see, for example, Fara and Williamson 2005), and has subsequently been revised by a number of authors (Dorr 2010, Kment 2012, Russell 2013, Bacon 2014). Do parallel problems arise for the foregoing counterpart-theoretic treatment of epistemic modality? If so, are parallel solutions available? These are important questions for assessing the ultimate tenability of the theory outlined above, but ones I leave as a matter for future research.

Aloni 2001, Ch. 3.³² There are two reasons for examining this alternative system. First, doing so will give us a better sense of the theoretical options for dealing with these issues. Second, the project of attempting to choose between the resulting two theories helps to illuminate the phenomenon that underlies our lottery puzzle (§5).

Recall Quine’s remark: ‘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.’ We have thus far been interpreting this as implying that being possibly thus and so (in the epistemic sense of ‘possibly’) is not a trait of an object *simpliciter*, but a trait that an object has relative to a way of thinking. But Quine’s remark can be interpreted in a slightly different way: being possibly thus and so is not a trait of an object at all, but, rather, a trait of a *way of thinking* of an object.

Contingent identity systems are most naturally thought of as interpreting the Quinean insight in this second way. Unlike counterpart theory, contingent identity systems typically do not change the semantics for the modal operator; rather, they alter the standard semantics for quantifiers and variables. Rather than taking quantifiers to range over objects, contingent identity systems take them to range over *ways of thinking* of objects (Moss 2018, §7.5). In the possible worlds setting, ways of thinking of objects are typically represented by *individual concepts*, or functions from worlds to individuals. So, for example, the way of thinking of β that we employ when we think of it as *the red ticket* would be represented by the individual concept c_r that maps each world to the unique red ticket in that world.³³

In the lottery scenario, we have two different ways of thinking about the domain: we can think of it as a collection of numbered tickets or as a collection of colored tickets. Note that each of these ways of thinking is a way of thinking about a *set* of tickets. But, at least in the present case, it is plausible to suppose that each of these ways of thinking of the set “decomposes” into a (two-membered) set of ways of thinking, one for each object in the set. When we are thinking of $\{\alpha, \beta\}$ as a collection of colored tickets, we are thinking of α as the red ticket and β as the blue ticket. And when we are thinking of $\{\alpha, \beta\}$ as a collection of numbered tickets, we are thinking of ticket α as ticket #2 and ticket β as ticket #1.

This suggests that we can represent a way of thinking of the domain as a set of ways of thinking, one for each object in the domain. And if we represent ways of thinking of objects using individual concepts, then each way of thinking of the domain corresponds to a *set of individual concepts*, one individual concept for each element of the domain. Aloni’s proposal is precisely to use such sets to represent ways of thinking of the domain. More precisely, and to use her

³²For other work on contingent identity systems, see Carnap 1947, Bressan 1972, Gibbard 1975, Fitting 2004, and Holliday and Perry 2014.

³³This is not to say that ways of thinking *are* individual concepts, but merely that they can be so represented for certain purposes. This shouldn’t be too controversial in the present case, since the ways of thinking in question are (a) *descriptive* modes of presentation, and (b) modes of presentation that correspond to non-rigid individual concepts.

terminology, what we have in the lottery scenario are two different *conceptual covers*. We first give a formal definition of the notion of a conceptual cover, and then illustrate the notion using our lottery scenario.

The definition of a conceptual cover is relative to a model, and a model in dynamic semantics with conceptual covers is again a triple $\langle \mathcal{W}, \mathcal{D}, \mathcal{I} \rangle$ where these elements are defined as in counterpart semantics. An *individual concept* on a model is a function from \mathcal{W} into \mathcal{D} . Then we have:

Given a model \mathcal{M} , a set of individual concepts \mathcal{C} is a *conceptual cover* on \mathcal{M} iff:

- (i) for each object $o \in \mathcal{D}$ and each world $v \in \mathcal{W}$, there is a concept $c \in \mathcal{C}$ such that $c(v) = o$, and (*existence*)
- (ii) for each object $o \in \mathcal{D}$ and each world $v \in \mathcal{W}$, there is at most one concept $c \in \mathcal{C}$ such that $c(v) = o$. (*uniqueness*)

To illustrate the notion of a conceptual cover,, it will help to consider a model $\mathcal{M} = \langle \mathcal{W}, \mathcal{D}, \mathcal{I} \rangle$, which contains two objects in \mathcal{D} and two worlds in \mathcal{W} , and which we can depict as follows:

v	v'
#1, blue, winner	#1, red, loser
#2, red, loser	#2, blue, winner

Table 3: Model to illustrate conceptual covers

Now, there are two relatively “natural” sets of individual concepts over this model. The first contains two individual concepts, c_1 and c_2 , where c_1 maps a world u to the ticket numbered 1 in u , and c_2 maps a world u to the ticket numbered 2 in u . The second also contains two individual concepts, c_b and c_r , where c_b maps a world u to the ticket colored blue in u , and c_r maps a world u to the ticket colored red in u . The first set of individual concepts preserves the number of each ticket across worlds, whereas the second preserves color. Each of these sets of individual concepts – $\{c_1, c_2\}$ and $\{c_b, c_r\}$ – is a conceptual cover on \mathcal{M} .

To see a set of individual concepts that is *not* a conceptual cover, consider $\{c_1, c_b\}$. Since $c_1(v) = c_b(v)$ in the model depicted above, this set of concepts violates both existence (since no concept maps v to the red ticket numbered 2 in v) and uniqueness (since both concepts in the set map v to the same individual).

To understand how conceptual covers help to solve the lottery puzzle, we need to revise the dynamic theory discussed in §2.2.³⁴ There are two principal differences between the dynamic theory discussed earlier and the conceptual

³⁴Again, we sketch the conceptual covers theory in outline in the text, leaving a more detailed presentation for the appendix.

covers theory. First, variable assignments are now understood as functions from variables to *individual concepts*, which are functions from worlds to individuals. Second, quantifiers range over a contextually provided contextual cover \mathcal{C} . This means that the semantics recursively defines *the update of a state s with a formula ϕ relative to a conceptual cover \mathcal{C}* , written $s[\phi]^{\mathcal{C}}$. We think of \mathcal{C} as representing the contextually given way of thinking of the domain. But other than that, everything else remains essentially the same as it was in the theory discussed in §2.2.

This semantics associates the following update with (4):

Fact 4.

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

(Where $i = \langle v, g \rangle$ is a possibility and c an individual concept, $c(i) = c(v)$.) So a state of information s supports (4) relative to a conceptual cover \mathcal{C} just in case, for every individual concept $c \in \mathcal{C}$, there is a possibility $i \in s$ at which $c(i)$ is the winning ticket in i .

This account immediately predicts the distinctive sort of context-sensitivity exhibited by (4). In our conversation with Al, the tickets are being thought of via their numbers. This means that the context will deliver the “number” conceptual cover $\{c_1, c_2\}$. Relative to this conceptual cover, our state of information s_l will support (4). This is because each concept c in this cover maps some possibility $i \in s_l$ to the winning ticket in s_l . This follows straightforwardly from two facts. First, in some possibilities i compatible with what we know, the ticket numbered 1 in i wins, whereas in others i' , the ticket numbered 2 in i' wins. Second, c_1 maps a possibility i to the ticket numbered 1 in i , and c_2 maps i to the ticket numbered 2 in i .

But in our conversation with Barbara, the tickets are being thought of via their colors. This means that the context will deliver the “color” conceptual cover $\{c_b, c_r\}$. And relative to this conceptual cover, our state of information s does *not* support (4). This is because c_r maps no possibility $i \in s_l$ to the winning ticket in i . For c_r maps each possibility i to the ticket colored red in i ; but since we know that the red ticket lost, $c_r(i)$ is the losing ticket in i (for any $i \in s_l$). Thus, not only does s_l not support (4) relative to this cover, it supports its negation, which reflects the fact that the correct answer to Barbara’s question is ‘no.’

What does this account have to say about the rest of our puzzle? To answer this, we need to develop the theory in a bit more detail.

We begin by revising the definition of *support* slightly. Let us say that a state of information s *supports ϕ relative to a conceptual cover \mathcal{C}* just in case $s[\phi]^{\mathcal{C}} = s$.

In the conceptual covers theory, the equivalence between an epistemic *de re* modal predication and its *de dicto* counterpart is retained:

Fact 5. For any state s and conceptual cover \mathcal{C} , $s[(\lambda x.\diamond\phi)(a)]^{\mathcal{C}} = s[\diamond\phi(a/x)]^{\mathcal{C}}$.³⁵

Since simple *de dicto* epistemic modal claims are not sensitive to the choice of conceptual cover, neither are simple *de re* modal predications like (1), (2), and (5).³⁶ So like the dynamic theory of §2.2, it follows from the conceptual covers theory that our state of information s_l supports (1)–(3) (with respect to any cover) and fails to support (5) (with respect to any cover). Indeed, as before, s_l supports the negation (5) (with respect to any cover). Note that this means that according to the conceptual covers approach (and unlike the counterpart approach), simple *de re* modal predications are not sensitive to how we think of the domain of quantification.

So the conceptual covers approach can explain why (1)–(3) seem acceptable in the lottery scenario, and why (5) does not. But what does it have to say about the validity of our two inferences? Recall that the problems with the dynamic account of §2.2 stemmed from the fact that it failed to validate those inferences. At first glance, it might seem that little progress on this issue has been made, for it turns out that neither of our two inferences is valid on the conceptual covers account – at least not in the strict sense of *valid* (a caveat I shall explain shortly). An inference is *valid* on the conceptual covers theory just in case for every state s and every conceptual cover \mathcal{C} , if s supports the premises relative to \mathcal{C} , then it supports the conclusion relative to \mathcal{C} . Neither of our two inferences is valid in this sense.

To see that the inference from (4) to (5) fails, recall that, in our discussion of the context sensitivity of (4), we established that our state of information s_l in the lottery scenario supports (4) relative to the “number” conceptual cover. Since (5) is simply equivalent to its *de dicto* counterpart $\diamond r = w$ (**Fact 5**), and since there is no possibility $i \in s_l$ at which the red ticket is the winning ticket, s_l will not support (5) relative to any conceptual cover; *a fortiori*, it does not support it relative to the number conceptual cover. Thus, s_l supports (4) relative to the number conceptual cover, but s_l does not support (5) relative to that conceptual cover, which means that (4) does not entail (5) on this approach.

But the conceptual covers theorist has more room to maneuver here. For although the conceptual covers theory entails that these two inferences are invalid, it can offer a plausible account of our judgments here, an account that isn’t available to the standard dynamic theorist.

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 the place of truth, so this becomes: whenever the premises are supported, so is the conclusion.

³⁵Like Groenendijk et al. (1996), Aloni defines her semantics over a language that doesn’t contain an abstraction operator. But she represents English sentences like (5) using a formula with a wide-scope existential quantifier: $\exists x(x = r \wedge \diamond x = w)$. Unlike the version of the conceptual covers theory adopted here (see appendix for details), Aloni’s official version predicts that the inference from (1)–(3) to (4) is valid (in the strict sense). But this difference has little impact on our discussion (see footnote 19).

³⁶A *de dicto* epistemic modal claim $\diamond\phi$ or *de re* modal predication $(\lambda x.\diamond\phi)(a)$ is ‘simple’ in the relevant sense if ϕ contains no quantifier.

But how should we interpret the ‘whenever’ here? The definition of validity we imputed to the conceptual covers theory above offers one answer: ‘whenever’ means ‘for every state s and conceptual cover \mathcal{C} .’ So in asking whether an inference is valid in this sense, we are asking how it fares across all states and conceptual covers.

But other interpretations are possible. One alternative would be to interpret ‘whenever’ relative to a fixed conceptual cover. For example, suppose you are evaluating a particular inference in a certain context. If \mathcal{C} represents the way of thinking of the domain that is salient in your context, you might wonder whether the conclusion is supported by every state that supports the premises, holding fixed that \mathcal{C} represents the way of thinking of the domain that is salient in your context. This yields a different, “local” notion of validity, one which can be characterized as follows:

For any conceptual cover \mathcal{C} , an inference is *valid-in- \mathcal{C}* iff for any state s , if s supports the premises relative to \mathcal{C} , it supports the conclusion relative to \mathcal{C} .

The conceptual covers 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, 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 that delivers a conceptual cover that contains c_r , the intension of *the red ticket*. Why? Because (5), the conclusion of the inference, explicitly mentions the red ticket, thus making it natural to think of the domain as a collection of colored tickets. The second point to appreciate is this. Let \mathcal{C} be any conceptual cover on the model that contains c_r . Then the inference from (4) to (5) is valid-in- \mathcal{C} .³⁷

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, that is, they are very likely to be occupying a context that determines a conceptual cover \mathcal{C} such that the inference is valid-in- \mathcal{C} . Thus, the conceptual covers 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 that we’ve defined in the conceptual covers

³⁷To see this, let s be an information state, and let \mathcal{C} be any conceptual cover that contains c_r . Suppose that s supports (4) relative to \mathcal{C} . Given **Fact 4**, it follows from this supposition 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 $c_r \in \mathcal{C}$, it follows that there is a possibility $i \in s$ such that $c_r(i) = i(w)$. Since $c_r(i)$ is the red ticket in i , $c_r(i) = i(r)$. It follows that there is a possibility $i \in s$ such that $i(r) = i(w)$. So $s[\diamond r = w]^{\mathcal{C}} = s$, which, given **Fact 5**, means that $s[(\lambda x. \diamond x = w)(r)]^{\mathcal{C}} = s$. So s supports (5) relative to \mathcal{C} .

theory.^{38,39}

So the conceptual covers theorist has a plausible account of our puzzle. Sentences (1)–(3) seem assertable in the lottery scenario because they are supported by our state of information in that scenario (no matter what cover is chosen). Sentence (5) seems like something we are in a position to deny in the lottery scenario because our state of information supports its negation in that scenario (no matter what cover is chosen). The two inferences are not in fact valid (in the strict sense), but each inference *seems* valid because, for each inference, when one evaluates that inference, one is very likely to be occupying a context that determines a conceptual cover \mathcal{C} such that the inference is valid-in- \mathcal{C} .

5 Quantification without identification

So we have two solutions to our lottery puzzle, one provided by static counterpart theory, another by the dynamic conceptual covers theory. How should we choose between them? I propose to investigate this question by examining a variant on our initial lottery case. I think that the variation case provides an argument in favor of counterpart theory over the conceptual covers approach. But the variation case has a somewhat broader interest as well, for it helps us to situate the “lottery phenomenon” with respect to a more familiar phenomenon in the philosophy of language, namely Frege’s Puzzle about attitude ascriptions. I begin by introducing the variation case by returning to an issue I alluded to earlier, namely the question of how the Quinean insight ought to be formulated.

5.1 The one hundred ticket case

As I noted at the end of §3, I formulate the Quinean insight by saying that whether an object satisfies an epistemically modalized predicate depends on how the *domain of quantification* is thought of. In contrast, other authors (Quine included) take the observation to be that whether an object satisfies such a predicate depends on how *the object* is thought of. The latter way of

³⁸The conceptual covers theorist can also say something similar about why (11) appears to be contradictory.

(11) Although any ticket might be the winning ticket, its false that the red ticket might be the winning ticket. $(\forall x \diamond x = w) \wedge \neg((\lambda x. \diamond x = w)(r))$

Sentence (11) will ordinarily be uttered or considered in a context which delivers a conceptual cover \mathcal{C} that includes the intension of ‘the red ticket,’ since that definite description occurs explicitly in the sentence. But (11) will be *inconsistent-in- \mathcal{C}* , in the sense that for any state s , $s[(11)]^{\mathcal{C}} = \emptyset$, given that \mathcal{C} contains c_r . This is the conceptual covers theorist’s explanation of why (11) sounds contradictory: uttering or considering it tends to give rise to a context relative to which it is inconsistent.

³⁹If this explanation seems far-fetched to you, note that something like it might be needed in less exotic theoretical contexts. For when we combine the standard view that quantifier domains are contextually restricted with the idea that an inference is valid just in case it preserves *truth at a context*, universal instantiation appears to fail. All men might be mortal, and Socrates might be a man, but if Socrates isn’t in the domain of ‘all’ in this context, then these premises do not guarantee that Socrates is mortal. See Gauker 1997 and Szabó 1998 for discussion.

putting the insight might seem more natural, and it works perfectly fine as long as we stick to cases like our lottery case, cases in which each way of thinking of the domain can be decomposed into a set of ways of thinking of objects, one for each object in the domain. But we often possess a way of thinking of a group without that way of thinking decomposing into a set of ways of thinking of the objects in the group.

For example, I might have two ways of thinking about a certain group of men: as the members of the Boston Red Sox and as the people on a certain bus. I might believe that every member of the Boston Red Sox has a beard without believing that every person on that bus has a beard. I could have these two ways of thinking of this group without either of these ways of thinking decomposing into a set of ways of thinking of individual men, one for each man. I may not know any of the names of the Red Sox players, nor be in a position to pick them out by their position (I may know little about baseball). And I may not be able to see all of the people on the bus, and so may not possess, for some of the people on the bus, a visual mode of presentation of those people. A way of thinking of a group is not in general reducible to a group of ways of thinking, one for each element in the group.

That is fairly unsurprising. What's interesting, however, is that the dependence of epistemically modalized predicates on ways of thinking arises even in cases in which we do not possess a mode of presentation for each object in the domain.⁴⁰ Here is an example. Imagine that there are one hundred tickets in a certain room. Exactly one ticket has a tiny and very valuable diamond attached to it; let's call this ticket 'the winning ticket.' The back of each ticket is either completely red or completely blue; on the front of each ticket there is either a black circle on an otherwise white background or a black triangle on an otherwise white background. Let us call the former 'circular tickets' and the latter 'triangular tickets.' We know that there are twenty-five tickets of each color-shape combination: we know that there are twenty-five blue triangular tickets, twenty-five blue circular tickets, twenty-five red triangular tickets, and twenty-five red circular tickets. We also know that the winning ticket is blue. But since we don't know whether that ticket is a circular ticket or a triangular ticket, we don't know which shape is depicted on the front the winning ticket.

Imagine that you are shown two photographs. In the first, you can see the backs of the tickets organized by color into two piles (so you see one pile of red tickets, and one pile of blue tickets). In the second, you can see the fronts of the tickets organized by depicted shape into two piles (so you see one pile of circular tickets, and one pile of triangular tickets). If you were to look at the second of these photographs, it would be natural enough for you to point to the pile of triangular tickets and say, 'Any of those might be the winning ticket.' And you could point at the pile of circular tickets and say the same thing.

So this scenario gives rise to a puzzle very similar to our initial one, for we again appear to be able to reason from apparently true premises to an apparently false conclusion:

⁴⁰Thanks to Sarah Moss for pressing me to think about cases like this.

- (13) Any circular ticket might be the winning ticket. $\forall x(Cx \rightarrow \Diamond x = w)$
- (14) Any triangular ticket might be the winning ticket. $\forall x(Tx \rightarrow \Diamond x = w)$
- (15) Every ticket is circular or triangular. $\forall x(Cx \vee Tx)$
- So:
- (4) Any ticket might be the winning ticket. $\forall x\Diamond x = w$
- So:
- (16) Any red ticket might be the winning ticket. $\forall x(Rx \rightarrow \Diamond x = w)$

Again, this inference is valid given standard assumptions about the workings of the non-modal vocabulary. And again the premises seem true and the conclusion false. The conclusion appears to be false since we know that a blue ticket won. Note that you could, for example, point at the pile of red tickets in the first photograph and say, ‘The winner can’t be in that pile.’⁴¹

Now, in this scenario, it seems to me that the quantified sentence (4) is just as context-sensitive as it was in the original lottery scenario. When we’re looking at the photograph in which the shapes of the tickets are visible, it is natural to say that any of the tickets might be the winning ticket. But when we’re looking at the photograph of the colored backs of the tickets, it seems natural to deny this, on the grounds that the winner cannot be found among the red tickets.

Now note that (4) is true/assertable in a context c just in case each ticket o satisfies the predicate ‘is such that it might be the winning ticket’ in c . So given that (4) varies in truth/assertability across these two contexts, and that the contexts only differ in how the tickets are being thought of, it presumably follows that there is a ticket o and modes of presentation m_1 and m_2 such that o satisfies the predicate in question relative to m_1 but not relative to m_2 . The mode of presentation m_1 should somehow be made salient by the ‘shape context’ (since every ticket should satisfy that predicate in the shape context), and the mode of presentation m_2 should somehow be made salient in the ‘color context’ (since o is, by hypothesis, one of the objects that fails to satisfy that predicate in the color context). But what is o and what are m_1 and m_2 ?

Now o is presumably a red ticket. For suppose o were a blue ticket. Then o would satisfy the relevant predicate in the shape context, since every ticket satisfies that predicate in the shape context. But o would also satisfy that predicate in the color context, since, given what we know, any of the blue tickets

⁴¹ According to the dynamic theories discussed in this essay, the inference from (4) to (16) is not, strictly speaking, valid (though the dynamic conceptual covers theorist may be in a position to offer an explanation of why it seems valid). This is of some interest, since it suggests that, according to dynamic semantics, the restrictor of a universal quantifier is not a ‘downward entailing’ environment. This may bring dynamic semantics into conflict with a well-known account of the conditions under which *negative polarity items* (for example, ‘every,’ ‘any’) are licensed (Ladusaw 1979, von Stechow 1999). This may bear on the ultimate tenability of dynamic semantics, but I must leave this as an issue for future research.

could be the winner. So if o is a ticket that satisfies the relevant predicate in the shape context, but not in the color context, o must be a red ticket.

But what are the two modes of presentation m_1 and m_2 such that o satisfies the predicate ‘is such that it might be the winning ticket’ relative to m_1 but not relative to m_2 ? It seems to me that we can’t assume that either m_1 or m_2 is a mode of presentation of o , since we (*qua* characters in the lottery scenario) may not possess *any* (relevant) mode of presentation of o . For imagine that we saw the two photographs quite quickly, or even that we never saw them at all, the setup of the case having simply been explained to us in words. In either of those cases, we might now have no way of “mentally singling out” any particular red ticket, o included. We may have no *demonstrative* mode of presentation of o : it might be that neither our current nor our past perceptual experience puts us in a position to think of o as ‘that red ticket.’ And we may have no *descriptive* mode of presentation of o ; we may not, for example, be in a position to think of o via some description like ‘the slightly bent red ticket on the left side of the pile.’ So we might not possess any “red ticket” mode of presentation of o , that is, we might not possess a mode of presentation of o relative to which the belief ‘ o is a red ticket’ is trivial. Similarly, we might not possess any relevant ‘shape’ mode of presentation of o . To see this, note that o is either a triangular ticket or a circular ticket. But, again, since we might not be in a position to mentally single out any of the triangular tickets *as* triangular tickets nor any of the circular tickets *as* circular tickets, we might not possess either a “triangular ticket” mode of presentation of o nor a “circular ticket” mode of presentation of o . So it seems that we cannot assume that either m_1 or m_2 is a mode of presentation of o .

It seems that the natural thing to say at this point is that m_1 and m_2 are not modes of presentation of o , but, rather, modes of presentation of the *entire domain of tickets*. Mode m_1 is something like *the collection of triangular and circular tickets*, and mode m_2 is something like *the collection of red and blue tickets*. On this approach, whether o , an arbitrary red ticket, satisfies the predicate in question depends on how we are thinking about *the tickets* as a whole, rather than on how we are thinking about o itself. This is why I formulate the Quinean insight in terms of how the *domain*, rather than the *object*, is thought of.

5.2 Lottery cases and Frege cases

Our main interest in the one hundred ticket case concerns how it might bear on the choice between the two theories discussed in §4. But before I turn to that, I want to pause here to say something about how the lottery cases we’ve been discussing relate to a type of case more familiar to philosophers of language: Frege cases. Although the two types of cases are obviously related, the one hundred ticket case suggest that phenomenon under investigation in this essay differs in important ways from standard Frege cases.⁴²

⁴²See Moss (2018, §7.5) for a discussion of some relevant similarities between sentences containing the adverb ‘probably’ and attitude ascriptions of various kinds.

By a “standard Frege case,” I mean a case like the following. Imagine that Sue fails to realize that Cicero is Tully. In part because of her ignorance of this identity claim, she comes to believe that Cicero denounced Catiline, while at the same time failing to believe that Tully denounced Catiline. Now, in a case like this, it is natural (though not inevitable) to suppose that whether the predicate

‘is such that Sue believes that he denounced Catiline’

applies to Cicero/Tully in a context c depends on which of Sue’s modes of presentation of Cicero/Tully is salient in c . If, in c , we are focused on Sue’s “Cicero” mode of presentation, it will be natural to think that Cicero/Tully satisfies this predicate in c . But if, in c , we are instead focussed on Sue’s “Tully” mode of presentation, it will be natural to think that Cicero/Tully fails to satisfy this predicate in c .⁴³

Our original two-ticket lottery case is similar to the Cicero-Tully case in many respects. For example, in that scenario, we (*qua* characters in the lottery scenario) fail to realize that the ticket #2 is the red ticket. In part because of our ignorance of this identity claim, it is compatible with what we know that ticket #2 is the winning ticket, but not compatible with what we know that the red ticket is the winning ticket. And it is natural to suppose that whether the predicate

‘is such that it might be the winning ticket’

applies to this ticket in a context c depends on which of these modes of presentation is salient in c .

In both of these cases, the relevant subject possesses two modes of presentation, m_1 and m_2 , of an object o , but the subject fails to realize that m_1 and m_2 are two ways of thinking of the same object. (Sue fails to realize that Cicero is Tully, we fail to realize that the red ticket is ticket #2.) Furthermore, in each case, whether a certain predicate applies to o in a context c depends on which of these ways of thinking, m_1 or m_2 , is salient in c .

But the one hundred ticket case does not appear to have this structure. For as we have just been arguing, in the one hundred ticket case, there need not be any ticket o for which we have two modes of presentation, m_1 and m_2 , such that: (i) we fail to realize that m_1 and m_2 are two ways of thinking of o ; and (ii) whether o satisfies the relevant predicate varies according to which of m_1 and m_2 is chosen. For in the one hundred ticket case, there may be no ticket such that we possess two (relevant) ways of thinking of *it*.

It might be replied that the one hundred ticket case really does have the structure of a Frege case, but that the relevant object o about whose identity we are confused is not an individual ticket, but rather a *group* of tickets, and that the relevant m_1 and m_2 are, accordingly, modes of a presentation of that group. But which group could this be? We appear to possess only one mode

⁴³Views of this sort are fairly common in the literature on attitude ascriptions. See, for example, Crimmins and Perry 1989 and Heim 1992.

of presentation of the set of red tickets, for example, since neither the set of circular tickets nor the set of triangular tickets coincides in extension with the set of red tickets. By the setup of the case, for each of the two shapes, there are red tickets bearing that shape.

We do, of course, possess two ways of thinking of the entire domain of tickets. And I argued above that whether the predicate ‘is such that it might be the winning ticket’ applies to any particular red ticket depends on which of these ways of thinking is salient in the context. But note that, unlike in standard Frege cases, we (the relevant subjects) *know* that these are two ways of thinking of the same thing: in the one hundred ticket scenario, we (*qua* characters in that scenario) know that the collection of red and blue tickets is identical to the collection of circular and triangular tickets. We are not like Sue who *doesn’t know* that Cicero is Tully. This is an unusual feature of the one hundred ticket case: whether the relevant predicate applies to a red ticket is sensitive to which mode of presentation is chosen, even though the relevant subjects (namely us) are not ignorant that these two modes of presentation are modes of presentation of the same thing.

Given the Frege-like structure of the original two ticket lottery case, it would have been natural to assume that the phenomenon being investigated in this essay is simply one aspect of a phenomenon that is very familiar to philosophers, namely Frege’s Puzzle. But the one hundred ticket case suggests that this assumption may not be quite right. While the phenomenon we’re discussing is clearly related in some way to Frege’s Puzzle, there also appear to be clear differences. How exactly to understand the relationship between these two phenomena is an interesting question, but one I leave as a matter for future inquiry.

5.3 Counterparts and conceptual covers again

How do cases like the one hundred ticket case bear on the choice between the two theories presented in §4, the static counterpart theory and the dynamic conceptual covers theory? It seems to me that the counterpart story extends straightforwardly to the one hundred ticket case, but that the conceptual covers approach faces a challenge here. This yields a *prima facie* argument for the former and against latter. Our focus here will be on how each theory predicts the fact that whether (4) is true/assertable depends on how the tickets are being thought of.

What we want is a theory that predicts that (4) is true/assertable when we’re thinking of the tickets as a collection of circular and triangular tickets, but false/deniable when we’re thinking of the tickets as a collection of blue and red tickets. According to counterpart theory, “shape contexts” will deliver the *shape counterpart* relation, where $\langle v, o \rangle$ bears the shape counterpart relation to $\langle v', o' \rangle$ just in case the shape depicted on o in v is identical to the shape depicted on o' in v' . And sentence (4) is true in the one hundred ticket scenario in a context that delivers the shape counterpart relation. To see this, let o be an arbitrary ticket in the world v in which that scenario obtains. Ticket o is either circular or triangular in v . Since we (*qua* characters in the lottery scenario)

don't know the shape of the winning ticket, there are worlds v' compatible with what we know in which the winning ticket o' is circular and worlds v'' compatible with what we know in which the winning ticket o'' is triangular. If o is circular in v , then $\langle v', o' \rangle$ will be a shape counterpart of $\langle v, o \rangle$ that wins in v' . If o is triangular in v , $\langle v'', o'' \rangle$ will be a shape counterpart of $\langle v, o \rangle$ that wins in v'' . So either way, $\langle v, o \rangle$ will have a winning shape counterpart. Since o was an arbitrary ticket in v , this holds for them all, which means that (4) is true in v , given the shape counterpart relation.

“Color contexts,” on the other hand, will deliver the *color counterpart* relation, defined as before: $\langle v, o \rangle$ bears the color counterpart relation to $\langle v', o' \rangle$ just in the color of o in v is identical to the color of o' in v' . Sentence (4) is *not* true in the one hundred ticket scenario in a context that delivers the color counterpart relation. To see this, consider any red ticket o in the world v of the scenario. Since we (*qua* characters in the lottery scenario) know that the winning ticket is not red, there is no world compatible with what we know whose winning ticket is red. It follows that $\langle v, o \rangle$ has no color counterpart that wins in some epistemically accessible world, which means that (4) is false in v , given the color counterpart relation.

So the counterpart story extends to this case straightforwardly. What about the conceptual covers theory? Now, the conceptual cover theorist would presumably begin her story about the variability of (4) by saying something like this:

‘Let s_h be our state of information in the one hundred ticket scenario. When we're thinking of the one hundred tickets as collection of circular and triangular tickets, the context will deliver a *shape-preserving conceptual cover*, where \mathcal{C} is a shape-preserving conceptual cover just in case for each concept $c \in \mathcal{C}$ and any pair of worlds v, v' in s_h , the shape of $c(v)$ in v is identical to the shape of $c(v')$ in v' . And relative to such a cover, our state of information s_h supports (4).’

The conceptual covers theorist will continue this story by saying something about how s_h does not support (4) relative to the conceptual cover delivered by “color contexts.” But we can stop here because we already have a problem.

The problem begins with the observation that there are *many* conceptual covers that are shape-preserving in the above sense. To get a sense of why this is so, think of a pair of worlds v and v' which are both compatible with our information state s_h in the one hundred ticket scenario. World v will contain fifty triangular tickets, $\alpha_1, \dots, \alpha_{50}$, and fifty circular tickets, $\beta_1, \dots, \beta_{50}$. World v' will also contain fifty triangular tickets, $\alpha'_1, \dots, \alpha'_{50}$, and fifty circular tickets, $\beta'_1, \dots, \beta'_{50}$:

v	v'
α_1	α'_1
.	.
.	.
.	.
α_{50}	α'_{50}
β_1	β'_1
.	.
.	.
.	.
β_{50}	β'_{50}

Table 4: Two worlds in the one hundred ticket case

Now a conceptual cover over s_h will determine a way of drawing one hundred lines on the above table, where each line connects exactly one element on the left-side of the dividing line with exactly one element on the right-side of the dividing line.⁴⁴ A *shape-preserving* conceptual cover will then correspond to a set of one hundred lines meeting this constraint plus the additional constraint that no line connects a triangular α_i on the left-side of the line to a circular β'_k on the right-side of the line. It's easy to see that there are many ways of drawing lines on this table meeting these two constraints. And each such set of lines will correspond to a different shape-preserving conceptual cover over s_h .

Now the fact that there are many shape-preserving conceptual covers over s_h might not matter that much if the theory ended up predicting that s_h supported (4) relative to *every* shape-preserving cover. But this turns out not to be the case. Note, for example, that some shape-preserving conceptual covers will consist of the following:

- one individual concept c_t such that, for any world v :

$$c_t(v) = \begin{cases} \text{the winning ticket in } v & \text{if a triangular ticket wins in } v; \\ \text{a triangular ticket in } v & \text{otherwise;} \end{cases}$$

- forty-nine individual concepts c such that c maps each world v to a losing triangular ticket in v ; and

⁴⁴Let \mathcal{C} be a conceptual cover over s_h . Then \mathcal{C} determines the following way of drawing one hundred lines on the above table. For each object o on the left-side of the dividing line, draw a line between o and an object o' on the right-side of the dividing line just in case there is a $c \in \mathcal{C}$ such that $c(v) = o$ and $c(v') = o'$.

- fifty individual concepts c such that c maps each world v to a circular ticket in v .

Of the fifty “triangular” individual concepts in a cover like this, only one – namely c_t – will ever be the winning ticket in an accessible world. The other forty-nine triangular individual concepts will map each world v to a losing triangular ticket in v . But recall **Fact 4**:

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

At least forty-nine of the individual concepts in the conceptual cover under consideration are such that there is no $i \in s_h$ such that $c(i)$ is the winning ticket in i . It follows then that s_h does not support (4) relative to this shape preserving cover (indeed, it follows that s_h supports its negation relative to this cover).

Now, it is important to note that there *are* shape-preserving conceptual covers relative to which s_h supports (4).⁴⁵ So whether the conceptual covers theory predicts that (4) is assertable in a shape context depends on whether that context delivers a ‘good’ cover or a ‘bad’ cover, that is, on whether it delivers a cover relative to which s_h supports (4) or one relative to which s_h does not support (4). The challenge for the conceptual covers theorist is to tell a story about this context – about us, the tickets, and how we’re thinking of them – that makes it plausible to assume that the context will deliver a conceptual cover relative to which s_h supports (4).

I don’t see how the challenge can be met, at least not without substantially revising the conceptual covers theory. For the only clear constraint that this context offers is that the relevant conceptual cover be shape-preserving. If this is right, then it would be natural to assume that, rather than delivering a *unique* conceptual cover, this context delivers a *set* of admissible covers. A cover will be admissible in this context just in case it is shape-preserving. It would then be tempting to assess the assertability of a sentence by supervaluating over admissible covers, so that (4) is assertable only if s_h supports (4) relative to every admissible cover. But of course this approach doesn’t yield the desired result, for it tells us that (4) is not assertable in the shape context, since we know that there are shape-preserving covers relative to which s_h does not support (4).

Could the conceptual covers theory be revised so as to avoid this problem? Perhaps. But until a satisfactory revision that avoids this problem is spelled out, we have a *prima facie* reason to prefer a counterpart-theoretic approach to these matters.

⁴⁵At least this is true on the assumption that s_h contains at least one hundred worlds. Assuming knowledge of all other facts except those pertaining to the lottery, there are only two qualitatively distinct types of worlds in s_h : worlds in which a blue triangular ticket wins and worlds in which a blue circular ticket wins. But given the way we’ve set things up, s_h will presumably contain many “haecceistically distinct” worlds of each qualitative type.

It's worth observing that the problem here appears to arise because the relevant way of thinking of the domain does not decompose into a set of ways of thinking, one for each ticket in the domain. To see this, imagine that the triangular tickets were numbered 1 through 50, and that the circular tickets were numbered 51 through 100, and that we knew these facts. Then the context would presumably deliver a shape-preserving conceptual cover that also preserved a ticket's number across worlds. And if we didn't know the number of the winning ticket, our information state s_h would support (4) relative to a conceptual cover like that. But in this version of the case, our "shape and number" way of thinking of the domain decomposes into a set of ways of thinking, one for each ticket in the domain: *the triangular ticket numbered 1, the triangular ticket numbered 2, etc.*

It's also worth pointing out that the challenge facing the dynamic conceptual covers approach appears to be more of a problem for the "conceptual covers" part of that theory rather than for the "dynamic" part of that theory. Note, for example, that moving to a *static* conceptual covers theory wouldn't help matters much. For presumably such a theory would assign to (4) the following truth conditions:

$$\llbracket \forall x \diamond x = w \rrbracket^{v.g.\mathcal{R}.\mathcal{C}} = 1 \text{ iff for every } c \in \mathcal{C}, \text{ there is a } v' \in \mathcal{W} \text{ such that } v\mathcal{R}v' \text{ and } c(v') \text{ is the winning ticket in } v'.$$

We want (4) to come out (determinately) true in the one hundred ticket scenario when we are thinking of the tickets via their shapes. But it is unclear how any theory of this sort could yield that prediction. Again, the problem is that many conceptual covers are shape-preserving, and (4) will be false relative to some of those covers.

Note also that the problem here appears to be independent of the features that make the conceptual covers approach distinct from other systems of contingent identity. The source of the problem is not the conceptual covers theory's commitment to the *existence* and *uniqueness* conditions (§4.2). For any contingent identity theory that hopes to model the context-sensitivity of a sentence like (4) will likely have to say that the context determines the set of individual concepts over which the quantifiers quantify. But in a contexts like the one discussed above, the only clear constraint that the context issues is that the relevant set of individual concepts be shape-preserving. But that clearly isn't enough to rule out 'bad' sets of concepts relative to which (4) is false/unassertable.

So the problem facing the conceptual covers approach here cannot simply be solved by moving to a static version of the theory, nor by dropping some of the constraints on what counts as a conceptual cover.⁴⁶

⁴⁶The one hundred ticket case, like the two ticket case, has a certain special feature: we know how many elements are in the domain of quantification. But obviously we often use quantified sentences even without knowing this ('Everyone on this airplane can hear that baby crying'). In such cases, the size of the domain will vary across epistemically possible worlds, which makes it hard see how Aloni's *uniqueness condition* on covers could be retained. In any case, the conceptual covers response to our challenge should apply to cases like this as well, and to cases in which we know, for example, that every ticket is either red or blue, but we don't know how many red and how many blue.

6 Static vs. dynamic semantics

Let us take stock. We argued in §2, that the two ticket lottery puzzle posed various problems for both standard static and standard dynamic approaches to quantifiers and epistemic modals. We extended this argument in §3, arguing that neither of those approaches could accommodate the Quinean insight, which we took to be the key to solving the puzzle. In §4, we argued that accommodating the Quinean insight and resolving the puzzle requires adopting a non-standard theory of transworld representation. We looked at two ways of doing this: static counterpart theory and Aloni’s dynamic conceptual covers theory. And we have just been comparing those two approaches, arguing that the one hundred ticket lottery case poses a challenge to the conceptual covers approach, a challenge that the static counterpart approach appears to avoid.

Now the static counterpart theory and the dynamic conceptual covers theory differ along two dimensions: counterpart relations vs. conceptual covers, and static vs. dynamic semantics. As noted above, the problem we raised for the dynamic conceptual covers theory appears to be more of a problem for the conceptual covers part of that theory, rather than for its “dynamism.” So even if one is persuaded by the above argument that we should use counterpart relations rather than conceptual covers, one might still wonder whether to prefer a static or dynamic implementation of counterpart theory. An important question here is whether there is a viable dynamic version of counterpart theory, one that replicates the static counterpart treatment of the foregoing phenomena. Although I have not come across such a theory in the extant literature, I see no in-principle barrier to constructing such a theory. In any case, our discussion certainly leaves open the possibility of such a theory.⁴⁷

Nevertheless, I think our discussion does bear on the dispute between static and dynamic semantics (though it does not, of course, resolve it). In particular, the possibility of a static counterpart theoretic approach complicates a certain line of argument that might be taken to favor dynamic approaches over static ones. The argument centers on the observation that the following sentences are infelicitous (in most normal contexts):

(17) The red ticket is such that it might not be the red ticket. $(\lambda x.\Diamond x \neq r)(r)$

(18) Some red ticket is such that it might not be red. $\exists x(Rx \wedge \Diamond \neg Rx)$

(19) Some ticket that is red and might not be red lost. $\exists x((Rx \wedge \Diamond \neg Rx) \wedge Lx)$

(20) Every ticket that is red but might not be red lost. $\forall x((Rx \wedge \Diamond \neg Rx) \rightarrow Lx)$

(21) Every red ticket is such that it might not be red. $\forall x(Rx \rightarrow \Diamond \neg Rx)$

As Yalcin (2015, 485) observes, these sentences appear to be defective in some way, at least in most ordinary contexts.⁴⁸ Yalcin also notes that this is not

⁴⁷Thanks to the editors of the *Philosophical Review* for emphasizing this possibility.

⁴⁸Groenendijk et al. (1996) and Aloni (2001, Ch. 3) discuss examples similar to (18). For detailed discussion of these and related sentences, see Yalcin 2015, Klinedinst and Rothschild 2016, and Mandelkern 2017, Ch. 1.

something that we should expect given standard static theories, such as the ones discussed in §2.1, since those theories predict that these sentences have coherent truth-conditions. Furthermore, standard dynamic theories – such as the one we discussed in §2.2 – predict that these sentences are indeed defective, as Yalcin points out. So it would be quite reasonable to take these facts to constitute an argument for dynamic semantics over static semantics.⁴⁹

Now if the arguments of §§4–5 are correct, then a theory of quantified epistemic modality needs counterpart relations (or something similar). For as we saw in §2.2 and §3, standard dynamic semantics by itself is not enough to handle either the lottery puzzle or the context-sensitivity of (4). That then raises a question: if we already have counterpart relations in our theory, do we *also* need dynamic semantics in order to handle sentences (17)–(21)?

After all, it seems like the counterpart theorist is in a position to offer an alternative *static* story about why these sentences are defective. Take (18), for example. A sentence like this is most naturally evaluated in a context that delivers a counterpart relation that preserves the color of the relevant tickets across worlds, since that sentence makes salient the color of the tickets. And relative to a context like that, (18) will be false, because no red ticket has a color counterpart that is not red. So (18) is defective because it is false in the sort of context in which it is most naturally evaluated. Similar explanations of the other sentences listed above can likewise be given.

How might we decide between these two approaches? One relevant consideration is that the dynamic theorist’s explanation of (17)–(21) is *semantic*, while the counterpart explanation is *pragmatic*. Sentence (18), for example, is *inconsistent* according to dynamic semantics, in the technical sense that updating any state of information with it yields the empty set. In contrast, counterpart theory predicts that sentence (18) is *true* relative to some contexts, though these will be contexts in which the counterpart relation fails to preserve a ticket’s color across worlds. Sentence (18) is said to be defective by the counterpart theorist because it is *most naturally evaluated* relative to contexts that deliver a color preserving counterpart relation. This difference between the theories may help us to decide between them. If, for example, it is possible to describe contexts in which (17)–(21) can be truly uttered, that would appear to be a point in favor of the static approach. But if, on the other hand, this is not possible, that would seem to tell in favor of the dynamic approach.

Another relevant difference between the theories concerns the semantic status of sentences (11) and (12), though here the “semantic-pragmatic” tables are turned.

(11) Although any ticket might be the winning ticket, its false that the red ticket might be the winning ticket. $(\forall x \diamond x = w) \wedge \neg((\lambda x. \diamond x = w)(r))$

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

⁴⁹Yalcin himself doesn’t take this last step. He is careful to leave open the possibility that a (non-standard) static theory might explain the above data (Yalcin 2015, 518).

Static counterpart theory tells us that these sentences are contradictory, as they indeed seem to be. Neither of the dynamic theories discussed in this essay makes a similar prediction, though, as we saw, the dynamic conceptual covers theorist can offer a pragmatic explanation of why these sentences might seem contradictory. This difference between the two approaches may give us yet another consideration that bears on the choice between them. But how this dispute ultimately plays out is a matter I must leave for future inquiry.

References

- Aloni, Maria. 1997. “Quantification in Dynamic Semantics.” In Paul Dekker, Martin Stokhof, and Yde Venema (eds.), *Proceedings of the 11th Amsterdam Colloquium*. Amsterdam: ILLC, University of Amsterdam.
- . 2001. *Quantification under Conceptual Covers*. Ph.D. thesis, University of Amsterdam.
- Bacon, Andrew. 2014. “Representing Counterparts.” *The Australasian Journal of Logic* 11.
- Beaver, David I. 2001. *Presupposition and Assertion in Dynamic Semantics*. Stanford: CSLI Publications.
- Bressan, Aldo. 1972. *A General Interpreted Modal Calculus*. New Haven: Yale University Press.
- Burgess, John P. 1997. “Quinus Ab Omni Naevo Cindicatus.” *Canadian Journal of Philosophy* 27:25–65.
- Carnap, Rudolf. 1947. *Meaning and Necessity*. Chicago: University of Chicago Press.
- Crimmins, Mark and Perry, John. 1989. “The Prince and the Phone Booth: Reporting Puzzling Beliefs.” *Journal of Philosophy* 86:685–711. doi:10.2307/2027014.
- DeRose, Keith. 1991. “Epistemic Possibilities.” *Philosophical Review* 100:581–605.
- Dorr, Cian. 2010. “How to be a Modal Realist.” Unpublished manuscript, Oxford University.
- Dowell, Janice. 2011. “A Flexible Contextualist Account of Epistemic Modals.” *Philosophers’ Imprint* 14:1–25.
- Egan, Andy. 2007. “Epistemic Modals, Relativism, and Assertion.” *Philosophical Studies* 133:1–22.
- Egan, Andy, Hawthorne, John, and Weatherson, Brian. 2005. “Epistemic Modals in Context.” In Gerhard Preyer and Georg Peter (eds.), *Contextualism in Philosophy*, 131–168. Oxford: Oxford University Press.
- Fara, Michael and Williamson, Timothy. 2005. “Counterparts and Actuality.” *Mind* 114:1–30.
- von Fintel, Kai. 1999. “NPI Licensing, Strawson Entailment, and Context Dependency.” *Journal of semantics* 16:97–148.
- von Fintel, Kai and Gillies, Anthony. 2008. “CIA Leaks.” *Philosophical Review* 117:77–98.
- . 2011. ““Might” Made Right.” In Andy Egan and Brian Weatherson (eds.), *Epistemic Modality*, 108–130. Oxford: Oxford University Press.
- Fitting, Melvin. 2004. “First-order Intensional Logic.” *Annals of Pure and Applied Logic* 127:171–193.
- Fitting, Melvin and Mendelsohn, Richard L. 1998. *First-Order Modal Logic*. Dordrecht: Kluwer Academic Publishers.
- Frege, Gottlob. 1892. “Über Sinn und Bedeutung.” *Zeitschrift für Philosophie und philosophische Kritik*, C 25–50. English Translation: “On Sense and Meaning,” in McGuinness, B. (ed), Frege: Collected Works, Oxford: Basil Blackwell, 157–177.
- Garson, James W. 2006. *Modal Logic for Philosophers*. Cambridge: Cambridge University Press.
- Gauker, Christopher. 1997. “Universal Instantiation: A Study of the Role of Context in Logic.” *Erkenntnis* 46:185–214.

- Gerbrandy, Jelle. 1997. "Questions of Identity." In Paul Dekker, Martin Stokhof, and Yde Venema (eds.), *The Proceedings of the 11th Amsterdam Colloquium*, Amsterdam.
- . 1998. "Identity in Epistemic Semantics." In *Proceedings of the Third Conference on Information Theoretic Approaches to Logic, Language and Computation*, 147–159.
- Gibbard, Allan. 1975. "Contingent Identity." *Journal of Philosophical Logic* 4:187–221.
- Groenendijk, Jeroen, Stokhof, Martin, and Veltman, Frank. 1996. "Coreference and Modality." In Shalom Lappin (ed.), *Handbook of Contemporary Semantic Theory*, 179–216. Oxford: Blackwell.
- Hacking, Ian. 1967. "Possibility." *Philosophical Review* 76:143–168.
- Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst.
- . 1992. "Presupposition Projection and the Semantics of Attitude Verbs." *Journal of Semantics* 9:183–221.
- Holliday, Wesley H and Perry, John. 2014. "Roles, Rigidity, and Quantification in Epistemic Logic." In *Johan van Benthem on Logic and Information Dynamics*, 591–629. Springer.
- Hughes, G. E. and Cresswell, M. J. 1968. *A New Introduction to Modal Logic*. London and New York: Routledge.
- Kamp, Hans. 1981. "A Theory of Truth and Semantic Representation." In J.A.G. Groenendijk, T.M.V. Janssen, and M.B.J. Stokhof (eds.), *Formal Methods in the Study of Language*, volume 135, 277–322. Amsterdam: Mathematisch Centrum.
- Kaplan, David. 1989. "Demonstratives." In Joseph Almog, John Perry, and Howard Wettstein (eds.), *Themes from Kaplan*, 481–563. New York: Oxford University Press.
- Karttunen, Lauri. 1969. "Discourse Referents." In *Proceedings of the 1969 Conference on Computational Linguistics*, 1–38. Association for Computational Linguistics.
- Klinedinst, Nathan and Rothschild, Daniel. 2016. "Quantifiers, Pronouns, and Modals." Handout from a talk at the Pacific APA.
- Kment, Boris. 2012. "Individuals, Supervenience, and Fundamentality." *Philosophical Review* 121:573–609.
- Kratzer, Angelika. 1991. "Modality." In A. von Stechow and D. Wunderlich (eds.), *Semantics: An International Handbook of Contemporary Research*, 639–650. Berlin: Walter de Gruyter.
- Ladusaw, William. 1979. *Negative Polarity Items as Inherent Scope Relations*. Ph.D. thesis, University of Texas at Austin.
- Lennertz, Benjamin. 2015. "Quantificational Credences." *Philosophers' Imprint* 15:1–24.
- Lewis, David K. 1968. "Counterpart Theory and Quantified Modal Logic." *Journal of Philosophy* 65:113–126.
- . 1971. "Counterparts of Persons and Their Bodies." *Journal of Philosophy* 68:203–211.
- . 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- MacFarlane, John. 2011. "Epistemic Modals Are Assessment-Sensitive." In Andy Egan and Brian Weatherson (eds.), *Epistemic Modals*, 144–178. Oxford: Oxford University Press.
- . 2014. *Assessment-Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.
- Mandelkern, Matthew. 2017. *Coordination in Conversation*. Ph.D. thesis, MIT.
- Moss, Sarah. 2018. *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Ninan, Dilip. 2016. "Relational Semantics and Domain Semantics for Epistemic Modals." *Journal of Philosophical Logic* 1–16. Advance online publication.
- . 2018. "Naming and Epistemic Necessity." Unpublished manuscript, Tufts University.
- Quine, W.V.O. 1953. "Reference and Modality." In *From a Logical Point of View*. Cambridge: Harvard University Press.
- Russell, Bertrand. 1905. "On Denoting." *Mind* 14:479–493.

- Russell, Jeffrey Sanford. 2013. “Actuality for Counterpart Theorists.” *Mind* 122:85–134.
- Santorio, Paolo. 2012. “Reference and Monstrosity.” *Philosophical Review* 121:359–406.
- Schaffer, Jonathan. 2011. “Perspective in Taste Predicates and Epistemic Modals.” In Andy Egan and Brian Weatherson (eds.), *Epistemic Modality*, 179–226. Oxford: Oxford University Press.
- Stalnaker, Robert. 1970. “Pragmatics.” *Synthese* 22:272–289.
- . 1974. “Pragmatic Presuppositions.” In M. K. Munitz and P. Unger (eds.), *Semantics and Philosophy*. New York: New York University Press.
- . 1978. “Assertion.” In P. Cole (ed.), *Syntax and Semantics 9: Pragmatics*, 315–332. New York: Academic Press. Reprinted in Stalnaker 1999, 78–95. Page references are to the 1999 reprint.
- . 1999. *Context and Content*. New York: Oxford University Press.
- Stephenson, Tamina. 2007. “Judge Dependence, Epistemic Modals, and Predicates of Personal Taste.” *Linguistics and Philosophy* 30:487–525.
- Szabó, Zoltán Gendler. 1998. “Review: Christopher Gauker, Universal Instantiation: A Study of the Role of Context in Logic.” *The Journal of Symbolic Logic* 63:1610–1611.
- Veltman, Frank. 1996. “Defaults in Update Semantics.” *Journal of Philosophical Logic* 25:221–261.
- Willer, Malte. 2013. “Dynamics of Epistemic Modality.” *Philosophical Review* 122:45–92.
- Yalcin, Seth. 2007. “Epistemic Modals.” *Mind* 116:983–1026.
- . 2011. “Nonfactualism About Epistemic Modality.” In Andy Egan and Brian Weatherson (eds.), *Epistemic Modality*, 295–332. Oxford: Oxford University Press.
- . 2012. “A Counterexample to Modus Tollens.” *Journal of Philosophical Logic* 41:1001–1024.
- . 2015. “Epistemic Modality *De Re*.” *Ergo* 2:475–527.

Appendix

Assume a language of quantified modal logic with identity and individual constants whose primitive logical symbols are: $=$, \neg , \wedge , λ , \exists , and \diamond . The formation clause for the abstraction operator, λ , is as follows: for any formula ϕ , variable x , and individual constant a , $(\lambda x.\phi)(a)$ is a formula. The other logical symbols are defined in the usual way: for example, $\forall x\phi$ is $\neg\exists x\neg\phi$ and $\phi\vee\psi$ is $\neg(\neg\phi\wedge\neg\psi)$. The *terms* of the language are the variables and the individual constants.

Static semantics and counterpart semantics

A *static frame* is an n -tuple $\mathcal{F} = \langle \mathcal{W}, \mathcal{D}, \dots \rangle$ where \mathcal{W} is a non-empty set (of *worlds*) and \mathcal{D} is a non-empty set (of individuals), the domain of the frame. The ellipses indicate that a static frame may contain other elements as well, such as a set of contexts, a set of times, an accessibility relation, etc..

A *variable assignment* on a frame \mathcal{F} is a function from variables to individuals in the domain of \mathcal{F} . An *index* e on a frame \mathcal{F} is an n -tuple where the elements in e are drawn from \mathcal{F} (so e may, for example, be $\langle c, v, t \rangle$, where c is a context, v a world, and t a time). A *point of evaluation* e, g on a frame \mathcal{F} is a pair of an index e and a variable assignment g .

A *static model* is an n -tuple $\langle \mathcal{W}, \mathcal{D}, \dots, \mathcal{I} \rangle$ consisting of a static frame and an interpretation function \mathcal{I} that: (i) maps each individual constant a to a function

$\mathcal{I}(a)$ that maps each index e to an element of \mathcal{D} ; and (ii) maps each n -ary predicate P to a function that maps each index e to a subset of \mathcal{D}^n (the set of n -ary sequences of elements of \mathcal{D}).

Where \mathcal{M} is any static model, e, g any point of evaluation, and t any term, we have:

$$\llbracket t \rrbracket^{\mathcal{M}, e, g} = g(t), \text{ if } t \text{ is a variable;}$$

$$\llbracket t \rrbracket^{\mathcal{M}, e, g} = \mathcal{I}(t)(e) \text{ if } t \text{ is an individual constant.}$$

The ‘standard assumptions’ about the non-modal vocabulary of \mathcal{L} discussed in §2.1 are as follows:

$$\llbracket P(t_1, \dots, t_n) \rrbracket^{\mathcal{M}, e, g} = 1 \text{ iff } \langle \llbracket t_1 \rrbracket^{\mathcal{M}, e, g}, \dots, \llbracket t_n \rrbracket^{\mathcal{M}, e, g} \rangle \in \mathcal{I}(P)(e)$$

$$\llbracket t = t' \rrbracket^{\mathcal{M}, e, g} = 1 \text{ iff } \llbracket t \rrbracket^{\mathcal{M}, e, g} = \llbracket t' \rrbracket^{\mathcal{M}, e, g}$$

$$\llbracket \neg\phi \rrbracket^{\mathcal{M}, e, g} = 1 \text{ iff } \llbracket \phi \rrbracket^{\mathcal{M}, e, g} = 0$$

$$\llbracket \phi \wedge \psi \rrbracket^{\mathcal{M}, e, g} = 1 \text{ iff } \llbracket \phi \rrbracket^{\mathcal{M}, e, g} = 1 \text{ and } \llbracket \psi \rrbracket^{\mathcal{M}, e, g} = 1$$

$$\llbracket (\lambda x.\phi)(a) \rrbracket^{\mathcal{M}, e, g} = 1 \text{ iff } \llbracket \phi \rrbracket^{v, g[x/o]} = 1, \text{ where } o = \llbracket a \rrbracket^{\mathcal{M}, e, g} \text{ }^{50}$$

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

To get a complete recursive definition of truth (at a model and point of evaluation), we need to add a clause for the modal operator \diamond . But whatever clause we adopt, the inferences from (1)–(3) to (4) and from (4) to (5) will preserve truth at a model and point of evaluation, as discussed in §2.1.

The counterpart semantics of §4.1 takes a model to be a triple $\mathcal{M} = \langle \mathcal{W}, \mathcal{D}, \mathcal{I} \rangle$, with \mathcal{W} and \mathcal{D} defined as above, and with indices identified with elements of \mathcal{W} . An accessibility relation on such a model is a binary relation on \mathcal{W} ; an information state is a subset of \mathcal{W} ; and a counterpart relation is a reflexive binary relation on $\mathcal{W} \times \mathcal{D}$. A point of evaluation consists of a world, an accessibility relation (or an information state), a counterpart relation, and a variable assignment. The counterpart semantics retains all of the above assumptions about the semantics of the non-modal vocabulary, and then completes the recursive semantics by adding a distinctive clause for the modal operator, in which that operator shifts both the world and the variable assignment of the point of evaluation (as discussed in §4.1).

Dynamic semantics, with and without conceptual covers

In both versions of dynamic semantics discussed in this essay, a model is again a $\langle \mathcal{W}, \mathcal{D}, \mathcal{I} \rangle$. In standard dynamic semantics (that is, dynamic semantics without conceptual covers), the definition of a variable assignment remains unchanged: a variable assignment is a function from variables to individuals in \mathcal{D} . A *possibility*

⁵⁰For an extensive discussion of terms and predicate abstraction in the context of quantified modal logic, see Fitting and Mendelsohn (1998).

$i = \langle v, g \rangle$ is a pair of a world v and a variable assignment g , and a *state of information* is a set of possibilities.

The semantics takes the form of a recursive definition of *the update of a state s with a sentence ϕ relative to a model \mathcal{M}* , written $s[\phi]^{\mathcal{M}}$ (though we suppress reference to the model).⁵¹ To state the semantics, we first need some notation: Where $i = \langle v, g \rangle$ is any possibility:

$$\begin{aligned}
i(x) &= g(x) && \text{for any variable } x \\
i(a) &= \mathcal{I}(a)(v) && \text{for any individual constant } a \\
i(P) &= \mathcal{I}(P)(v) && \text{for any predicate } P \\
i[x/o] &= \langle v, g[x/o] \rangle && \text{for any object } o \\
s[x/o] &= \{i[x/o] : i \in s\} \\
i[x/a] &= \langle v, g[x/i(a)] \rangle \\
s[x/a] &= \{i[x/a] : i \in s\}
\end{aligned}$$

Where s any state of information on any model \mathcal{M} , we have:

$$\begin{aligned}
s[P(t_1, \dots, t_n)] &= \{i \in s : \langle i(t_1), \dots, i(t_n) \rangle \in i(P)\} \\
s[t = t'] &= \{i \in s : i(t) = i(t')\} \\
s[\neg\phi] &= s - s[\phi] \\
s[\phi \wedge \psi] &= s[\phi][\psi] \\
s[(\lambda x.\phi)(a)] &= \{i \in s : i[x/a] \in s[x/a][\phi]\} \\
s[\exists x\phi] &= \{i \in s : \exists o \in \mathcal{D} : i[x/o] \in s[x/o][\phi]\} \\
s[\diamond\phi] &= \begin{cases} s & \text{if } s[\phi] \neq \emptyset; \\ \emptyset & \text{otherwise.} \end{cases}
\end{aligned}$$

Given that $\forall x\phi$ is defined as $\neg\exists x\neg\phi$, we obtain the following result:

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

An important fact about this system that we employed in §2.2 is **Fact 2**:

$$\text{For any state } s \text{ on any model } \mathcal{M}, s[(\lambda x.\diamond\phi)(a)] = s[\diamond\phi(a/x)].$$

⁵¹What follows is a slightly modified version of the theory of Groenendijk et al. 1996. There are two main differences between the theory discussed here and the theory presented in Groenendijk et al. 1996. First, whereas Groenendijk et. al. are interested in providing an account of inter-sentential anaphora, that issue is not within our remit. This allows us to simplify the treatment of quantifiers slightly. Second, Groenendijk et. al. do not define their semantics over a language that contains an abstraction operator, and so we will need to extend their approach to such a language. (An alternative approach to the abstraction operator is mentioned in footnote 19.)

This follows from a more general fact about this system:

Fact 6. *For any state s on any model \mathcal{M} , $s[(\lambda x.\phi)(a)] = s[\phi(a/x)]$.*

This can be proved by induction on the complexity of formulas (see Ninan 2018 for relevant discussion).

In dynamic semantics *with* conceptual covers, a variable assignment is a function from variables to individual concepts. A conceptual cover over a model is a set of individual concepts that satisfies the existence and uniqueness conditions (see §4.2). Possibilities and states are defined as before, *modulo* the change in the definition of a variable assignment. The conceptual covers semantics recursively defines the notion of *the update of a state s with a sentence ϕ relative to a conceptual cover \mathcal{C} and model \mathcal{M}* , written $s[\phi]^{\mathcal{M},\mathcal{C}}$. The clauses of the recursive definition all remain the same (save for the addition of the superscript “ \mathcal{C} ”) except the one for the existential quantifier, which now runs as follows:

$$s[\exists x\phi]^{\mathcal{C}} = \{i \in s : \exists c \in \mathcal{C} : i[x/c] \in s[x/c][\phi]^{\mathcal{C}}\}$$

Given the definition of the universal quantifier, we obtain the following result:

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

The new notation used here is explained as follows. Where $i = \langle v, g \rangle$ is any possibility:

$$i[x/c] = \langle v, g[x/c] \rangle \quad \text{for any individual concept } c$$

$$s[x/c] = \{i[x/c] : i \in s\}$$

And some old notation is re-defined slightly. Where $i = \langle v, g \rangle$ is any possibility:

$$i[x/a] = \langle v, g[x/\mathcal{I}(a)] \rangle \quad \text{for any individual constant } a$$

In §4.2, we noted that the analogue of **Fact 2** still holds, and this too holds in virtue of the following more general fact

Fact 7. *For any state s and conceptual cover \mathcal{C} on any model \mathcal{M} ,*

$$s[(\lambda x.\phi)(a)]^{\mathcal{C}} = s[\phi(a/x)]^{\mathcal{C}}.$$

The proof of this fact is analogous to the proof of **Fact 6** (see Ninan 2018 again).