First-Order Intensional Logic

Melvin Fitting
Dept. Mathematics and Computer Science
Lehman College (CUNY), 250 Bedford Park Boulevard West
Bronx, NY 10468-1589
e-mail: fitting@lehman.cuny.edu
web page: comet.lehman.cuny.edu/fitting

January 27, 2003

Abstract

First-order modal logic is very much under current development, with many different semantics proposed. The use of rigid objects goes back to Saul Kripke. More recently several semantics based on counterparts have been examined, in a development that goes back to David Lewis. There is yet another line of research, using intensional objects, that traces back to Richard Montague. I have been involved with this line of development for some time. In the present paper I briefly sketch several of the approaches to first-order modal logic. Then I present one that I call **FOIL** (for first-order intensional logic) in the Montague tradition that, I believe, is both expressive and natural. I briefly discuss in what sense it can be made to encompass the other approaches. Finally I provide tableau rules to go with the **FOIL** semantics.

1 Introduction

What is first-order modal logic for? Since this is obviously not a simple question; perhaps we should begin by asking, what is propositional modal logic for? Here we are on well-explored ground. With propositional modal logic, and its relational semantics, we want to explicate various constructs from natural language, and explore nuances of certain concepts arising in philosophical investigations. We want to model knowledge, at least in an ideal sense. We want to reason about action. And there is another purpose as well, one that has become clearer over the years. In studying propositional modal logics—primarily those characterized by classes of frames—we are also studying fragments of classical first-order (and higher-order) logic. This is known as correspondence theory. For this purpose axiomatizability (or not) is a central issue. In addition, axiom systems allow the construction of canonical models, which provides a metamathematical methodology that is uniform across many logics. Details matter a great deal, of course, but the broad outlines of propositional modal logics have been standardized for some time.

But the original question above was, what is *first-order* modal logic for? What do quantifiers add to the mix? Motivations based on natural language and philosophy are still central, though we have a much richer variety of things we can potentially formalize and investigate. Of course we want a semantics that agrees with our intuitive understanding, but now intuitions can, and do, differ substantially from person to person. Are designators rigid? Can objects exist in more than one possible world? Should there be a distinction between identity and necessary identity? And for that matter, is the whole subject a mistake from the beginning, as Quine would have it? Rather than a semantics on which we all generally agree, quite a disparate range has been proposed. We are still exploring what first-order modal semantics should be; the propositional case was settled long ago.

One motivation has disappeared, however. It does not seem to be useful to think of first-order modal logic as a way of studying fragments of classical logic. This means the role of frames, while still interesting, is not as central. And issues of axiomatizability, while also interesting, are not as critical either. In guiding a robot, or understanding natural language, automatibility is fundamental. This suggests a shift in emphasis to tableaus or resolution. But we still have the problem of just what a first-order modal logic should look like.

Current research in the semantics of first-order modal logic seems to have divided into two broad camps, one based on *counterpart relations*, the other based on *intensional objects*. The counterpart relation approach originated in work of David Lewis and was originally meant to address philosophical problems. More recently, beginning with Ghilardi, counterpart semantics has had an additional layer of complexity added to it, resulting in a semantics that behaves very well metamathematically, though its philosophical motivation is more obscure. Intensional objects trace back to Montague (and Carnap), and are the basis of a good deal of work of my own. Personally I'm more comfortable with this approach, both informally and technically, but that may be just my prejudice.

In this paper I'll sketch the background of first-order modal semantics, leading up to why elaborations like counterpart semantics or intensional objects were introduced in the first place. But primarily I'll present a rich system based on intensional objects—I call it *first-order intensional logic*, or **FOIL** for short. **FOIL** is not really new—versions of it have been around for some time, but this is the first presentation of it as a fully elaborated first-order modal system. I also present a sketch of counterpart semantics, and discuss how such an approach can be embedded into **FOIL**. This suggests an extension of the notion of frame which has a certain naturalness, and which provides satisfying completeness results. Finally I'll also give a tableau system for **FOIL**. Basically I hope to find buyers among those logicians who are still shopping around.

2 General Considerations

Some comments on terminology, before I really begin. Throughout this paper I'll be working with first-order languages in which atomic formulas have relation symbols and variables, and formulas are built up from these using logical connectives, quantifiers, and $\Box$ and $\Diamond$ in the usual way. It is simplest in this survey paper to omit constant and function symbols, though it is straightforward to add them. Some additional syntactic machinery will be introduced in Section 5.2. I'll use x, y, ..., with and without subscripts, as variables (later a second kind of variable will be introduced, but these will still be present). I'll use P, Q, ... as relation symbols. I'll call the resulting language the *basic first-order modal language*.

Throughout this paper I'll use $\mathcal{G}$ as a set of possible worlds and $\mathcal{R}$ as an accessibility relation. If $\mathcal{M}$ is a model (whatever that may be), I'll write $\mathcal{M}, \Gamma \Vdash_v \Phi$ to symbolize that formula Φ is true at possible world Γ of model $\mathcal{M}$, with respect to valuation v, which assigns values to free variables. This terminology and notation is far from standard, so the literature must be read with care, though there will always be some version of valuation and truth at a world with respect to it, whatever it is called and however it is notated.

There are only two general conditions that apply throughout—I state them here once and for all.

Negation $\mathcal{M}, \Gamma \Vdash_v \neg X \Leftrightarrow \mathcal{M}, \Gamma \nVdash_v X$.

Propositional Connectives $\mathcal{M}, \Gamma \Vdash_v X \wedge Y \Leftrightarrow \mathcal{M}, \Gamma \Vdash_v X$ and $\mathcal{M}, \Gamma \Vdash_v Y$, and similarly for other propositional connectives.

Conditions for quantifiers and modalities can vary—I'll discuss them on a case by case basis. I'll assume an equality relation symbol is part of the language, and if it is interpreted by the equality relation I'll say a model is *normal*.

3 Rigidity

Variables may be given world-independent meanings in models. Such meanings are said to be *rigid*. Historically this was the first quantified modal semantics to be introduced [17, 13], and technically it is the simplest approach. There are several variations in detail, but essentially the two discussed below are the main versions.

3.1 Varying Domain Models

A *varying domain* model is a structure, $\mathcal{M} = \langle \mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$, meeting the conditions that follow. $\mathcal{D}$ is a *domain function* from $\mathcal{G}$ to non-empty sets. For $\Gamma \in \mathcal{G}$, $\mathcal{D}(\Gamma)$ is the *domain of world* Γ. The *domain of the model* is $\cup_{\Gamma \in \mathcal{G}} \mathcal{D}(\Gamma)$. $\mathcal{I}$ is an *interpretation function*, mapping n-place relation symbols to functions from $\mathcal{G}$ to n-ary relations on the domain of the model. Note that the values $\mathcal{I}$ assigns are not rigid—they are functions on worlds. This is the case with every semantics proposed for quantified modal logic; after all, if relation symbols behave the same way in all worlds, the semantics essentially collapses since worlds can't be distinguished.

A *valuation* is a mapping from variables to the domain of a model. Since the value assigned to a variable does not depend on worlds, it is *rigid*—the same in every world. Now, the conditions for atomic formulas, modalities, and quantifiers are these familiar ones.

Atomic $\mathcal{M}, \Gamma \Vdash_v P(x_1, \ldots, x_n) \Leftrightarrow \langle v(x_1), \ldots, v(x_n) \rangle \in \mathcal{I}(R)(\Gamma)$.

Necessity $\mathcal{M}, \Gamma \Vdash_v \Box X \Leftrightarrow \mathcal{M}, \Delta \Vdash_v X$ for all $\Delta \in \mathcal{G}$ such that $\Gamma \mathcal{R} \Delta$.

Possibility $\mathcal{M}, \Gamma \Vdash_v \Diamond X \Leftrightarrow \mathcal{M}, \Delta \Vdash_v X$ for some $\Delta \in \mathcal{G}$ such that $\Gamma \mathcal{R} \Delta$.

Universal Quantifier $\mathcal{M}, \Gamma \Vdash_v (\forall x)\Phi \Leftrightarrow \mathcal{M}, \Gamma \Vdash_w \Phi$ for every valuation w that is like v except possibly on x, but $w(x)$ must be some member of $\mathcal{D}(\Gamma)$.

Existential Quantifier $\mathcal{M}, \Gamma \Vdash_v (\exists x)\Phi \Leftrightarrow \mathcal{M}, \Gamma \Vdash_w \Phi$ for some valuation w that is like v except possibly on x, but $w(x)$ must be some member of $\mathcal{D}(\Gamma)$.

Think of $\mathcal{D}(\Gamma)$ as the set of existent objects at possible world Γ—this set is allowed to vary from world to world, which is why the semantics is called varying domain. At each world quantifiers range over the existents of that world; this is ensured by the condition placed on w in the quantifier cases above. There is no requirement that an existent at one world should also be an existent at another; neither is it forbidden. It can happen that in evaluating the truth of a formula at a world we may encounter a variable designating a non-existent at that world. There is no good reason why variables must always designate existents—we can talk about a flying horse, we just don't think one exists.

The quantification used in this semantics is *actualist*. One thinks of a quantifier as ranging over what actually exists. Something of the flavor of free logic is apparent—we do not have the validity of $(\forall x)P(x) \supset P(y)$ for instance, because at a particular world the designation of the variable y may be a non-existent, and so not be in the range of the quantifier. There is an obvious shortcoming here: one can never simply refer to "everything," existents and non-existents alike. We can say "y exists" provided we have equality available—it is expressed by the formula $(\exists x)(x = y)$. Hence we can say y does not exist, $\neg(\exists x)(x = y)$—but we can't say "there are non-existents," since quantifiers only range over existents. For some purposes this is a real limitation.

Very commonly special conditions are imposed on the domain function $\mathcal{D}$: *monotonicity* ($\Gamma\mathcal{R}\Delta$ implies $\mathcal{D}(\Gamma) \subseteq \mathcal{D}(\Delta)$), or *anti-monotonicity* ($\Gamma\mathcal{R}\Delta$ implies $\mathcal{D}(\Delta) \subseteq \mathcal{D}(\Gamma)$). Anti-monotonicity corresponds to the validity of the well-known Barcan schema. But thinking of most intended applications, monotonicity and anti-monotonicity alone are not particularly natural, and I will not discuss them further here. On the other hand, if both conditions are imposed, we essentially get constant-domain models, and these are worth a section to themselves.

3.2 Constant Domain Models

A *constant domain model* is a structure, $\mathcal{M} = \langle \mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{I} \rangle$, where $\mathcal{D}$ is a non-empty *set*, instead of a function on worlds—take it to be analogous to the domain of the model in the varying domain case. The interpretation $\mathcal{I}$ is as before. Truth conditions are the same as in the varying domain case except for the quantifier conditions, which become the following.

Universal Quantifier $\mathcal{M}, \Gamma \Vdash_v (\forall x)\Phi \Leftrightarrow \mathcal{M}, \Gamma \Vdash_w \Phi$ for every valuation w in $\mathcal{D}$ such that w is like v except possibly on x.

Existential Quantifier $\mathcal{M}, \Gamma \Vdash_v (\exists x)\Phi \Leftrightarrow \mathcal{M}, \Gamma \Vdash_w \Phi$ for some valuation w in $\mathcal{D}$ such that w is like v except possibly on x.

One can, of course, think of a constant domain model as a varying domain model in which the domain function is constant. Near the beginning of the development of first-order modal semantics it was discovered that this could be expressed proof-theoretically rather elegantly: a formula is valid in constant domain models just in case it is valid in all varying domain models in which the Barcan and Converse Barcan formulas are valid. Together these say $\Diamond$ and $(\exists x)$ commute, as do $\Box$ and $(\forall x)$.

Once again, variables are interpreted rigidly. The significant feature is that the domain of quantification is the same from world to world. Now quantifiers are called *possibilist*—we can think of them as ranging over what does and what could exist, where existence in an alternative possible world is taken to be possible existence in this one. It is as if we started with a varying domain model, added quantification over the domain of the model, and suppressed the actualist quantifiers. This suggests that a semantics in which both actualist and possibilist quantifiers were available would be natural, and useful. Fortunately, one can get the effect of this rather easily.

Suppose we introduce into the language a special one-place predicate symbol, E, which we can think of as a primitive existence predicate. Relativized quantifiers can then be introduced, $(\forall^{\mathsf{E}} x)\Phi$ for $(\forall x)(\mathsf{E}(x) \supset \Phi)$, and $(\exists^{\mathsf{E}} x)\Phi$ for $(\exists x)(\mathsf{E}(x) \wedge \Phi)$, [13, 8]. In a constant domain model we can think of relativized quantifiers as actualist, while the unrelativized quantifiers are possibilist; in this way we easily get the best of both semantics. Technically, to make this work we need a few assumptions on the behavior of E in models. First, at each world its interpretation should be non-empty; and second, each member of $\mathcal{D}$ should be in the interpretation of E at some world. And even these conditions can be relaxed without extreme discomfort.

Since varying domain semantics can be simulated using constant domain semantics and relativized quantifiers, from a semantic point of view there is really little point in studying the varying domain version in much detail. These are not the only considerations, however. Axiomatic systems intended for constant domain systems have more complex completeness proofs. On the other hand, prefixed tableau systems for constant domain systems are considerably simpler than the varying domain versions. In the present work I'll stick to constant domain semantics from now on, allowing a primitive existence predicate and relativized quantifiers as appropriate.

3.3 But the Problems Are

In the semantical systems sketched above, variables are interpreted *rigidly*—not changing from world to world. While this is natural for many purposes—natural numbers are certainly rigid—it leads to difficulties in expressing certain things that we may want to say. For instance, in normal models $(x = y) \supset \Box(x = y)$ and $(x \neq y) \supset \Box(x \neq y)$ are both valid. Now, the morning star and the evening star are in fact identical objects, so if x and y refer to the *object* that the phrases 'morning star' and 'evening star' designate (in the real world), certainly $\Box(x = y)$ is the case, since there is only one object involved, and one cannot become two. But then, how do we express the very natural thought that the morning and evening stars might have been distinct, as the ancient Babylonians believed to be the case? There is simply no way we can do it. Rigidity does not allow the possible distinctness of objects that are equal in fact to ever be true.

4 Counterparts

David Lewis [21, 22] argues against rigid semantics because it requires that objects always have being and be identifiable across possible worlds. I will not go into the details of his objections. The key point for us is that he proposes a looser notion: allow an object in a possible world to have *counterparts* in other worlds, rather than itself being in other worlds as well. An object of this world could have multiple counterparts in another, or multiple objects here could have a single counterpart in another world. This provides much greater flexibility in the semantics. In particular, it becomes possible to provide counter-models to $(x = y) \supset \Box(x = y)$ and to $(x \neq y) \supset \Box(x \neq y)$. The counterpart idea has been further generalized, starting with the functor semantics of Ghilardi, [10, 11], to allow for multiple counterpart relations as well. I'll sketch the ideas briefly here, beginning with a variation on the Lewis version, then moving to the more complex, multiple counterpart version.

4.1 Counterpart Models—Lewis Version

What is presented here is not actually the semantics proposed by Lewis, but is a close relative in the Lewis style. The present version is sufficient to get across the basic ideas, and can be made closer to the real Lewis version via straightforward modifications. For instance, Lewis requires that different worlds contain distinct sets of objects, while we do not, but such a condition can be imposed on a model by using the E predicate and placing semantic restrictions on it.

The *domain* of a binary relation Q is the set of all x such that $\langle x, y\rangle \in Q$ for some y; the *codomain* is the set of all y such that $\langle x, y\rangle \in Q$ for some x. A *counterpart relation* on a set $\mathcal{D}$ is a binary relation C whose domain and codomain is $\mathcal{D}$. (The conditions on domain and codomain can be relaxed to give a more general semantics. It's more than I want to consider here.) If v and w are two valuations in $\mathcal{D}$ and C is a counterpart relation on $\mathcal{D}$, I'll say w is a *C-counterpart* to v provided, for each variable x, $\langle v(x), w(x)\rangle \in C$.

A *Lewis counterpart model* is a structure $\mathcal{M} = \langle \mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{C}, \mathcal{I} \rangle$, with everything just as in Section 3.2 except that $\mathcal{C}$ has been added to the setup, where $\mathcal{C}$ is a function mapping each member of $\mathcal{G} \times \mathcal{G}$ to a counterpart relation on $\mathcal{D}$. The idea is, if $\langle x, y \rangle \in \mathcal{C}(\Gamma, \Delta)$ then y is a counterpart, in world Δ, of the object x in world Γ. If Δ is not accessible from Γ the value of $\mathcal{C}(\Gamma, \Delta)$ does not really matter—I'll take it to be the empty relation by convention.

Truth conditions are the same as for constant domain models except for the modal conditions, which are replaced by the following.

Necessity $\mathcal{M}, \Gamma \Vdash_v \Box X \Leftrightarrow \mathcal{M}, \Delta \Vdash_w X$ for all $\Delta \in \mathcal{G}$ such that $\Gamma \mathcal{R} \Delta$ and for every valuation w that is a $\mathcal{C}(\Gamma, \Delta)$ counterpart of v.

Possibility $\mathcal{M}, \Gamma \Vdash_v \Diamond X \Leftrightarrow \mathcal{M}, \Delta \Vdash_w X$ for some $\Delta \in \mathcal{G}$ such that $\Gamma \mathcal{R} \Delta$ and for some valuation w that is a $\mathcal{C}(\Gamma, \Delta)$ counterpart of v.

In effect, when moving from one world to another we replace talk of an object by talk of its counterparts. Necessary truth requires truth at all alternative worlds, as usual, but now we require this truth to be for all counterparts of the objects from the original world. Possible truth is the dual.

Example 4.1 To illustrate how this works, here is a very simple Lewis counterpart model, $\mathcal{M} = \langle \mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{C}, \mathcal{I} \rangle$. The set of possible worlds is $\mathcal{G} = \{\Gamma, \Delta\}$. Accessibility is just $\Gamma \mathcal{R} \Delta$, with no other case holding. The domain is $\mathcal{D} = \{a, b\}$. The counterpart function is $\mathcal{C}(\Gamma, \Delta) = \{\langle a, a \rangle, \langle b, b \rangle, \langle a, b \rangle\}$, with $\mathcal{C}(X, Y)$ being the empty relation in all other cases. Finally, say P is a two-place relation symbol, and $\mathcal{I}(P)$ is the function assigning the empty set to Γ and $\{\langle a, b \rangle, \langle b, b \rangle\}$ to Δ. What $\mathcal{I}$ assigns to other relation symbols won't matter for this example.

To make the discussion more perspicuous, let us say x and y are the only variables (they are the only ones that appear in this example). Let w_1 be the valuation such that $w_1(x) = a$ and $w_1(y) = b$; since $\langle a, b \rangle \in \mathcal{I}(P)(\Delta)$ we have $\mathcal{M}, \Delta \Vdash_{w_1} P(x, y)$. Likewise let w_2 be the valuation such that $w_2(x) = b$ and $w_2(y) = b$; since $\langle b, b \rangle \in \mathcal{I}(P)(\Delta)$ we have $\mathcal{M}, \Delta \Vdash_{w_2} P(x, y)$.

Let v be the valuation such that $v(x) = a$ and $v(y) = b$ (the same as w_1 as it happens). Valuations w_1 and w_2 are the only $\mathcal{C}(\Gamma, \Delta)$ counterparts of v, so we have $\mathcal{M}, \Gamma \Vdash_v \Box P(x, y)$.

Since it is allowed for an object to have more than one counterpart at an alternative world, it is easy to construct Lewis counterpart models that invalidate $(x = y) \supset \Box(x = y)$, though we retain the validity of $(x = x) \supset \Box(x = x)$. Likewise, since different objects might have the same counterpart at an alternative world, we also lose the validity of $(x \neq y) \supset \Box(x \neq y)$. This was part of the motivation for introducing counterpart models in the first place. On the other hand, if the counterpart relation between any two worlds is the identity function, a Lewis counterpart model is essentially a rigid model in the sense of Section 3, so this semantics extends the earlier one.

4.2 Counterpart Models—Multiple-Counterpart Version

In order to provide a mathematically smooth approach to completeness for first-order modal logics, counterpart semantics has been generalized beyond the Lewis version presented above. Ghilardi, [10, 11], uses category theory in a fundamental way; the notion of a frame is modified so that it becomes a category with worlds as objects and morphisms supplying the counterpart notion. This functor-semantic approach has been quite successful for logics above **S4**. A different semantics, also based on the use of category theory, is the metaframe semantics of [28, 27, 30]. Functor semantics without category theory, loosely speaking, has been developed in [14, 15, 18, 20, 16]. Most people will find this more appealing than the direct use of category theory and, in addition, restrictions to logics above **S4** are no longer imposed by the mathematics. This section presents a mild variation of that semantics.

A *multiple counterpart model* is a structure $\mathcal{M} = \langle \mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{C}, \mathcal{I} \rangle$, where everything is just as in Lewis counterpart semantics, except that now $\mathcal{C}$ is a function mapping each member of $\mathcal{G} \times \mathcal{G}$ to a *set* of counterpart relations on $\mathcal{D}$. Rather like before, I assume that if $\Gamma \mathcal{R} \Delta$ does not hold then $\mathcal{C}(\Gamma, \Delta)$ is the empty set.

Lewis counterpart truth conditions are replaced by the following more complicated version.

Necessity $\mathcal{M}, \Gamma \Vdash_v \Box X \Leftrightarrow \mathcal{M}, \Delta \Vdash_w X$ for all $\Delta \in \mathcal{G}$ such that $\Gamma \mathcal{R} \Delta$ and for every valuation w that is a C-counterpart of v, for every $C \in \mathcal{C}(\Gamma, \Delta)$.

Possibility $\mathcal{M}, \Gamma \Vdash_v \Diamond X \Leftrightarrow \mathcal{M}, \Delta \Vdash_w X$ for some $\Delta \in \mathcal{G}$ such that $\Gamma \mathcal{R} \Delta$ and for some valuation w that is a C-counterpart of v, for some $C \in \mathcal{C}(\Gamma, \Delta)$.

This carries the Lewis counterpart ideas one step further. Now necessary truth at a world requires truth at all alternative worlds, with objects replaced by any of their counterparts, as in the Lewis version, but also *with respect to every counterpart relation.*

Example 4.2 Let $\mathcal{M}$ be the multiple counterpart model whose structure is the same as the Lewis counterpart model of Example 4.1, except that $\mathcal{C}$ now is the function such that $\mathcal{C}(\Gamma, \Gamma) = \mathcal{C}(\Delta, \Gamma) = \mathcal{C}(\Delta, \Delta) = \emptyset$, and $\mathcal{C}(\Gamma, \Delta) = \{C_1, C_2\}$ where $C_1 = \{\langle a, a \rangle, \langle a, b \rangle, \langle b, b \rangle\}$ and $C_2 = \{\langle a, b \rangle, \langle b, a \rangle\}$. (Note: in the Lewis counterpart model of Example 4.1, *the* counterpart relation between Γ and Δ was C_1 of the present example.) With valuations w_1 and w_2 as in Example 4.1, we of course still have $\mathcal{M}, \Delta \Vdash_{w_1} P(x, y)$ and $\mathcal{M}, \Delta \Vdash_{w_2} P(x, y)$. Then if valuation v is as in Example 4.1, we have $P(x, y)$ true at Δ with respect to every valuation that is a C_1 counterpart of v. (Essentially, this is all as before.) But, if w_3 is the valuation given by $w_3(x) = b$ and $w_3(y) = a$, then $\mathcal{M}, \Delta \nVdash_{w_3} P(x, y)$ since $\langle b, a \rangle \notin \mathcal{I}(P)(\Delta)$. Since w_3 is a C_2 counterpart of v we do not have $\mathcal{M}, \Gamma \Vdash_v \Box P(x, y)$.

Of course the multiple counterpart semantics just presented includes the Lewis version as a special case—just take $\mathcal{C}(\Gamma, \Delta)$ to be a set consisting of a single counterpart relation whenever $\Gamma \mathcal{R} \Delta$. Multiple counterpart *models* can always be converted to Lewis counterpart models, but this is not the case for *frames* [19].

4.3 Problems Still

Lewis counterpart semantics provides us with a version of intensional logic. We tend to believe there are things, objects, independent of us, and there are concepts, which are our creations. Concepts are what we use when we talk; objects are what we believe we are talking about. The morning star and the evening star are concepts, both of which designate a particular object, the planet Venus.

In counterpart semantics, objects are present since they are what counterpart relations connect, but the counterpart network is fundamental, and an object, at a world, is actually something like a slice across that network. The morning star/evening star object in this world has, in an alternative Babylonian world, two counterparts, one playing the morning star role, the other the evening star role. This raises problems of ontology. In counterpart semantics what, exactly, is the morning star? For that matter, what is the evening star? They can't be the object Venus, because Venus is Venus, no matter what, yet we don't have necessary identity between the morning star and the evening star. The morning star is something more like a web of relationships, connecting Venus in our world with some (actually non-existent) objects in the world of the Babylonians, and those with still other objects in other worlds, and so on—relationships that sometimes split and sometimes merge. In counterpart semantics the morning star is there in the network of relationships, somehow, but I find myself unable to point at it, figuratively speaking. In short, I have a problem identifying the subject matter of this semantics. Indeed, while the notion of counterpart is fundamental, there is no way of saying this object and that one are counterparts in the formal modal language.

All of the above applies to both the Lewis version of counterpart semantics and to the multiple counterpart version. Multiple counterpart semantics is mathematically better behaved than Lewis counterpart semantics—[15] proves a completeness theorem that applies to all modal predicate logics, in their terminology, making use of what they call canonical models, which are particular multiple counterpart models. Nonetheless, the multiple counterpart semantics has even greater problems of informal interpretation than the Lewis version. Intuitively, what is the role of the many counterpart relations? Perhaps eventually this will be found to be natural, but at the moment it is rather puzzling.

5 First-Order Intensional Logic

I mentioned earlier that the two major approaches to first-order modal semantics currently under development either involve some notion of counterpart relations, or involve intensional objects. I have been working with intensional objects for some time, and now want to present what I think is the fullest version of a first-order semantics that uses them. I will discuss, in Section 6, in what sense counterpart relation semantics can be encompassed within an intensional object approach.

What I propose is a semantics (and tableau system) which I call **FOIL**, for *first-order intensional logic*. Actually **FOIL** is a family, depending on a choice of the underlying propositional modal logic. When necessary to be more specific, I will write **FOIL-S4**, **FOIL-S5**, and the like. For the most part such specificity is not needed here, and I will simply refer to **FOIL**, meaning any one of the general family of logics. **FOIL** is a variation on the system developed in [8], and amounts to the first-order part of the system of [7], which itself is an extension of the intensional logic of Montague and Gallin, [24, 25, 26, 9].

In **FOIL** semantics models have a domain of objects, just as they did in Section 3. This is the same for all worlds—if a varying domain version is desired we can use an existence predicate E, as described in Section 3.2. There is quantification over objects, identity entails necessary identity; in short, we have a version of the earlier rigid semantics.

In addition to objects there will be what we call *intensions* or *intensional objects* or *concepts*. Typical informal intensions are *the morning star*, *the oldest person in the world*, or simply *that*. Intensions designate different objects under different circumstances—they are *non-rigid* designators. As such, they will be modeled by functions from possible worlds to objects. There will be quantification over intensions, as well as quantification over objects.

When working with intensions two different aspects naturally stand out. We can consider an intension as a whole, or we can consider what the intension happens to designate. Equality of intensions considered as wholes is more often called "synonymy." Thus *the morning star* and *the evening star* designate the same object in the present world, but need not do so under other circumstances—one might say they are equal but not synonymous. On the other hand, *the evening star* and *the first heavenly body seen in the evening* are presumably synonymous and so, as a consequence, designate the same object under all circumstances. These are distinctions we wish to capture in a formal semantics. But this leads to a problem that deserves a section to itself.

5.1 *De Re/De Dicto* Difficulties

Suppose we have a possible-world semantics with a domain of objects, as in earlier sections. And suppose we have an *intension*, f, that picks out an object in each world. Say, for instance, that the domain consists of numbers, worlds are time instances, and at each time instance f picks out the size of the world's population. Suppose P is a one-place relation symbol. We could take $P(f)$ to mean that the *intension* f has the property P, or that the *object designated by* f has the property P. **FOIL** will provide both alternatives, since both are useful. But there is a difficulty with the second version which requires some adjustments to the language.

Suppose $P(f)$ is to mean the object designated by f (at a world) has property P. Then how should $\Diamond P(f)$ be interpreted, at world Γ say? One alternative is to understand it as saying the thing designated by f at Γ (call it f_Γ) has the 'possible-P' property, and so at some alternative world Δ we have that f_Γ has property P. This is traditionally called the *de re* reading—a possible property is ascribed to a thing. But there is another alternative: taking the possibility operator as primary, we could understand $\Diamond P(f)$ at Γ to mean that at some alternative world, Ω, we have $P(f)$, and so at Ω the thing designated by f (call it f_Ω) has property P. This is traditionally called the *de dicto* reading—possibility applies to a sentence. Now, even if Δ and Ω turn out to be the same, there is no reason why f_Γ and f_Ω should be identical. (Unless designation is rigid, which essentially puts us back in Section 3.)

Clearly the *de re* and *de dicto* readings are different, and both are plausible. An abstraction mechanism has become standard to distinguish between them. Here the *de re* reading will be symbolized $\langle\lambda x.\Diamond P(x)\rangle(f)$ and the *de dicto* will be symbolized $\Diamond\langle\lambda x.P(x)\rangle(f)$. Disambiguation involves a move to a somewhat more complex language.

5.2 The FOIL Language

So far we have been using the basic first-order modal language, sketched in Section 2. Now it must be extended. For **FOIL** there are two sorts of variables, *object variables*, x, y, ... , as in previous sections, and *intension variables*, f, g, ... I'll assume each relation symbol has a *type* associated with it, where a type is an n-tuple whose entries are in $\{O, I\}$. An *atomic formula* is an expression of the form $P(\alpha_1, \ldots, \alpha_n)$ where P is a relation symbol whose type is $\langle t_1, \ldots, t_n\rangle$ and, for each i, if $t_i = O$ then α_i is an object variable, and if $t_i = I$ then α_i is an intension variable. I will assume there is a two place relation symbol $=$ of type $\langle O, O\rangle$. Equality of type $\langle I, I\rangle$ could also be added if desired—I will not do so. It would play the role of a synonymy symbol.

FOIL formulas are built up from atomic formulas exactly as usual, with the following piece of machinery added. If Φ is a formula, x is an object variable, and f is an intension variable, then $\langle\lambda x.\Phi\rangle(f)$ is a formula, in which the free variable occurrences are those of Φ except for x, together with the displayed occurrence of f. I will sometimes abbreviate $\langle\lambda x.\langle\lambda y.\Phi\rangle(g)\rangle(f)$ by $\langle\lambda x, y.\Phi\rangle(f, g)$, and so on. The idea is that $\langle\lambda x.\Phi\rangle(f)$ should say, at a world, the object designated by f at that world has the Φ property. This will be given an exact meaning in the semantics below, but perhaps some informal examples will be useful.

Suppose P is a relation symbol of type $\langle I\rangle$. The atomic formula $P(f)$ is intended to assert, at a world, that the *concept* f has the property P. For instance, say possible worlds are people, and f is the *favorite-book* concept picking out, for each person, that person's favorite book. And suppose P is intended to be the *is-an-important-concept* predicate, which different persons will apply in different ways. For a person who considers reading important, $P(f)$ will most likely be true—the concept of a favorite book would be important for that person. But $P(f)$ would most likely be false for a person who does not value reading or literature. Again, suppose Q is of type

$\langle O\rangle$, so that $\langle\lambda x.Q(x)\rangle(f)$ is a formula. Let us say Q is intended to be the *is-an-important-book* predicate. I certainly think $\langle\lambda x.Q(x)\rangle(f)$ is true—for me it says my favorite book is an important book (for me). I would not think $\langle\lambda x.\Box Q(x)\rangle(f)$ to be true—for me it says that my favorite book is an important book for everybody. On the other hand I probably would think $\Box\langle\lambda x.Q(x)\rangle(f)$ to be true—for me it says that everybody thinks their favorite book is important.

I have presented versions of **FOIL** before, but with some variations. In [8] essentially this system appears, with constant symbols of intension type but no intension variables. It was sufficient for the purposes we had in mind in that book, and adding intension variables and quantifiers is straightforward. The present system appeared in [5], but there was an additional piece of machinery: an explicit operator mapping an intension to an object at a world. Strictly speaking it was not necessary, but was included to make the relationship with the higher-order system of [4, 7] clear (such an operator is essential when higher orders are considered). I have dropped it here, in the interests of simplicity.

5.3 FOIL Models

A **FOIL** model is a structure $\mathcal{M} = \langle\mathcal{G}, \mathcal{R}, \mathcal{D}_O, \mathcal{D}_I, \mathcal{I}\rangle$ meeting the following conditions. $\mathcal{G}$ and $\mathcal{R}$ are as usual. $\mathcal{D}_O$ is a non-empty set, the *object domain.* $\mathcal{D}_I$ is a non-empty set of functions from $\mathcal{G}$ to $\mathcal{D}_O$; this is the *intension domain.* Finally, $\mathcal{I}$ is an *interpretation* such that, if P is a relation symbol of type $\langle t_1, \ldots, t_n\rangle$ then $\mathcal{I}(P)$ is a mapping from $\mathcal{G}$ to subsets of $\mathcal{D}_{t_1} \times \cdots \times \mathcal{D}_{t_n}$. I'll assume that $\mathcal{I}(=)$ is the constant function mapping each world to the identity relation on $\mathcal{D}_O$.

A *valuation* in **FOIL** model $\mathcal{M}$ is a mapping that assigns to each object variable a member of $\mathcal{D}_O$ and to each intension variable a member of $\mathcal{D}_I$.

Most of the various clauses of a truth definition are as usual. We now have two kinds of quantifiers, object and intension, since we have two kinds of variables, but semantic conditions for quantifiers are straightforward—if object variables are involved, quantification is over $\mathcal{D}_O$, and if intension variables are involved, quantification is over $\mathcal{D}_I$. The one substantially new clause is the following.

Abstraction $\mathcal{M}, \Gamma \Vdash_v \langle\lambda x.\Phi\rangle(f)$ if $\mathcal{M}, \Gamma \Vdash_w \Phi$ where w is like v except that $w(x) = v(f)(\Gamma)$.

The idea is, $\mathcal{M}, \Gamma \Vdash_v \langle\lambda x.\Phi\rangle(f)$ says the object designated by f at Γ has the property specified by Φ at Γ.

It is easy to see that at a world of a **FOIL** model, $\langle\lambda x, y.(x = y)\rangle(f, g)$ asserts intension variables f and g denote the same object. We do *not* have validity of the following.

$$(\forall f)(\forall g)[\langle\lambda x, y.(x = y)\rangle(f, g) \supset \Box\langle\lambda x, y.(x = y)\rangle(f, g)] \tag{1}$$

Formula (1) says that if f and g designate the same object in the present world, they will designate the same object in all accessible worlds (they will be synonymous), and this is certainly not the case. On the other hand, we *do* have validity of the following.

$$(\forall f)(\forall g)[\langle\lambda x, y.(x = y)\rangle(f, g) \supset \langle\lambda x, y.\Box(x = y)\rangle(f, g)] \tag{2}$$

In words, (2) says that if the objects denoted by f and g are identical, these objects are necessarily identical. This is so because identity between *objects* is necessary identity. Indeed, since objects behave as they did in Section 3, we have the validity of the following as well.

$$(\forall x)(\forall y)[(x = y) \supset \Box(x = y)] \tag{3}$$

5.4 Other Interesting Formulas

I'll begin with a simple formula that will be of use from time to time.

$$\mathbf{D}(f, x) \text{ abbreviates } \langle \lambda y.y = x\rangle(f) \text{ (where } x \text{ and } y \text{ are distinct object variables).} \tag{4}$$

If $\mathbf{D}(f, x)$ were atomic, rather than being an abbreviation for a more complex formula, it would be of type $\langle I, O\rangle$. Working through the **FOIL** semantics, $\mathcal{M}, \Gamma \Vdash_v \mathbf{D}(f, x)$ is true just in case $v(f)(\Gamma) = v(x)$. In words, it essentially says the intension f designates the object x at Γ.

The domain $\mathcal{D}_I$ of intensions in a model is required to be non-empty, and that is all. It is not required to be the set of *all* functions from $\mathcal{G}$ to $\mathcal{D}_O$. Such a requirement is undesirable for two reasons. First, not everything reasonably should be considered an intension. There isn't much plausibility to an intention that is a wrench in this world, a baby robin in another, and the number 7 in a third. Intensions should have some coherence to them, and though I don't know how to characterize that, clearly not everything mathematically possible will meet a reasonable coherence condition. The second reason for not taking the entire set of functions from $\mathcal{G}$ to $\mathcal{D}_O$ as $\mathcal{D}_I$ is more practical: if we do, a complete proof procedure is almost certainly beyond reach.

All this notwithstanding, there are a few extra requirements we might want to impose on $\mathcal{D}_I$, saying there are 'enough' members for various purposes. I now consider two such requirements.

There is no *a priori* reason to believe that every object is designated by some intension. After all, "There are more things in heaven and earth, Horatio, than are dreamt of in your philosophy." But under special circumstances we might want a requirement to this effect. Using the abbreviation of (4), we can do this by restricting ourselves to models in which we have the validity of

$$(\forall x)(\exists f)\mathbf{D}(f, x). \tag{5}$$

If we require (5), object quantification is reducible to intensional quantification:

$$(\forall x)\Phi \equiv (\forall f)\langle \lambda x.\Phi\rangle(f). \tag{6}$$

That is to say, the implication (5) $\supset$ (6) is valid in **FOIL** semantics.

The other condition we might sometimes want to impose on models is the existence of *choice functions*. Suppose that to each possible world Γ of a model we, somehow, associate an object d_Γ in $\mathcal{D}_O$. If our way of choosing d_Γ can be specified by a formula of the language, it is not unreasonable to insist that it amounts to an intension. Such a requirement can't quite be captured in a postulate (except for **FOIL-S5**), but here is something that comes close. For each formula Φ:

$$\Box(\exists x)\Phi \supset (\exists f)\Box\langle \lambda x.\Phi\rangle(x) \tag{7}$$

In a series of papers Hartley Slater has argued that modal logic should be built on the epsilon calculus; [29] makes the fundamental case for this. If one does so, he argues, then $\Box(\exists x)\Phi \supset (\exists x)\Box\Phi$ will be valid. While this principle is not valid, or desirable, in the present system, I note that (7) is actually a variation on it, split across the two types of **FOIL**.

5.5 Decidability Issues

Ordinarily when quantifiers are not present in a formula, whatever the logic, the problem of its validity is essentially a propositional one. But once the abstraction operator and intension variables are present, things are not so simple. The following is from [6].

For a propositional modal logic **L**, determined by a class of frames in the usual sense, **FOIL-L** is the intensional logic built on that class of frames in the obvious way. Now, let **FOIL-L-**λ be the restriction of **FOIL-L** to the sublanguage without quantifiers.

1. If **L** is one of **K**, **T**, or **D**, **FOIL-L**-λ is decidable. This can be shown using tableaus (see Section 7).

2. **FOIL-S5**-λ is undecidable, with or without equality.

3. If $=$ is interpreted by equality on $\mathcal{D}_O$ (as we have required), **FOIL-L**-λ is undecidable for any **L** between **K4** and **S5**.

4. The two preceding items remain true even if formulas are restricted to contain no object variables and only a single intension variable.

Clearly, adding intensional variables and abstraction is a major step, and not as elementary as may first appear.

5.6 Partiality

In **FOIL** models members of $\mathcal{D}_I$, the domain of intensions, are functions defined on the entire set of possible worlds. A plausible next step is to allow partial functions. If this is done, one can also extend the language to allow definite descriptions, where $\imath x.\Phi$ is of intension type, and is semantically treated as a partial function on worlds. I do not develop the machinery here, but refer to [8] where a version of it is presented at considerable length.

6 Relationships With Other Systems

It is obvious that the semantics of Section 3.2 embeds in **FOIL**. If one ignores the intensional aspects of **FOIL**, rigid constant domain semantics is what remains. Relationships between the counterpart semantics of Sections 4.1 and 4.2 and **FOIL** are not so simple. A connection involving models would be nice. Better yet would be a connection involving frames. Of course this requires us to specify what a **FOIL** frame is—I propose a definition, in 6.2 below, of *Riemann **FOIL** frames* that seems new to the literature, but that works quite well for these purposes. But let us begin with cases where everything works well and there are no complications. In particular, I begin with Lewis counterpart semantics, and postpone the more complex multiple counterpart semantics.

6.1 Simply Connected Frames

Under reasonable circumstances a relation can be looked at as a collection of functions. For instance, the *parent-of* relation is a union of the *father-of* function and the *mother-of* function. This carries over to Lewis counterpart models, provided the accessibility structure is not too complicated. The ideas in this section were independently developed in [16], with somewhat different terminology.

Definition 6.1 If $\langle \mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{C}, \mathcal{I} \rangle$ is a Lewis counterpart model (multiple counterpart model), I'll call $\langle \mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{C} \rangle$ a *Lewis counterpart frame* (*multiple counterpart frame*). Likewise if $\langle \mathcal{G}, \mathcal{R}, \mathcal{D}_O, \mathcal{D}_I, \mathcal{I} \rangle$ is a **FOIL** model, I'll call $\langle \mathcal{G}, \mathcal{R}, \mathcal{D}_O, \mathcal{D}_I \rangle$ a ***FOIL** frame.* And finally, I'll call $\langle \mathcal{G}, \mathcal{R} \rangle$ simply a *frame*, as usual.

Many Lewis counterpart models can be based on a Lewis counterpart frame, by varying the choice of an interpretation. Similarly in the multiple counterpart case. Likewise choices of interpretation allows us to build different **FOIL** models on a **FOIL** frame.

Definition 6.2 Let $\langle \mathcal{G}, \mathcal{R} \rangle$ be a frame. Call two members Γ and Δ in $\mathcal{G}$ *related* if $\Gamma \mathcal{R} \Delta$ or $\Delta \mathcal{R} \Gamma$. By a *path* in the frame I mean a sequence of members of $\mathcal{G}$, $\Gamma_1, \Gamma_2, \ldots, \Gamma_n$, such that each term of the sequence is related to the next. A frame is *simply connected* if there is just one path between any two members of $\mathcal{G}$. If $\mathcal{G}_0 \subseteq \mathcal{G}$, I'll say $\mathcal{G}_0$ is simply connected if the frame $\langle \mathcal{G}_0, \mathcal{R}_0 \rangle$ is simply connected, where $\mathcal{R}_0$ is $\mathcal{R}$ restricted to $\mathcal{G}_0$. I'll say a Lewis counterpart frame $\langle \mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{C} \rangle$ is simply connected if $\langle \mathcal{G}, \mathcal{R} \rangle$ is simply connected, and similarly for **FOIL** frames.

Being simply connected really consists of two requirements: between any two worlds there should be at least one path, and there should not be more than one. The first condition is not as significant. A frame in which there are no paths between some of the worlds can be thought of as multiple frames 'stuck together,' and these can be broken into separate frames. Details are straightforward, and well-known. The key item is the requirement that there should not be multiple paths between any pair of worlds—no loops, in other words.

Definition 6.3 Suppose $\mathcal{F} = \langle \mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{C} \rangle$ is a Lewis counterpart frame. Let f be a function defined on a simply connected subset of $\mathcal{G}$, mapping worlds to members of $\mathcal{D}$. I'll say f is $\mathcal{F}$ *compatible* provided, for each Γ, Δ in the domain of f such that $\Gamma \mathcal{R} \Delta$, $f(\Delta)$ is a counterpart of $f(\Gamma)$, that is, $\mathcal{C}(\Gamma, \Delta)(f(\Gamma), f(\Delta))$.

I omit the proof, but it is easy to show that if $\mathcal{F} = \langle \mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{C} \rangle$ is a simply connected Lewis counterpart frame, then any $\mathcal{F}$ compatible function can be extended to an $\mathcal{F}$ compatible function defined on the entire of $\mathcal{G}$.

Definition 6.4 If $\mathcal{F} = \langle \mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{C} \rangle$ is a simply connected Lewis counterpart frame, by its *companion* **FOIL** frame I mean $\langle \mathcal{G}, \mathcal{R}, \mathcal{D}_O, \mathcal{D}_I \rangle$ in which $\mathcal{D}_O = \mathcal{D}$ and $\mathcal{D}_I$ is the set of all $\mathcal{F}$ compatible functions with domain the entire of $\mathcal{G}$.

Figures 1 and 2 provide an elementary example of this. Figure 1 pictures a simply connected Lewis counterpart frame. The domain $\mathcal{D}$ consists of two objects, 'dot' and 'square'; the counterpart relation is shown by the arrows; the accessibility relation holds between the displayed world in the middle and the two outer worlds. Figure 2 shows the **FOIL** frame that is the companion to this. Instead of the counterpart relation, three functions are displayed using connected lines.

In a companion frame the original counterpart relations have been replaced by functions and, from them and the accessibility relation, we can reconstruct the original counterpart notion. Simple connectedness makes all this possible. Formulas that were evaluated in models using counterpart relations must be translated into formulas involving functions, intensions. (A translation that is closely to the one that follows can be found in [16].)

Definition 6.5 Let Φ be a formula in the basic first-order modal language of Section 2 (hence not containing intension variables or quantifiers, and not containing abstracts). A formula Φ^* in the **FOIL** language is defined as follows.

1. If A is atomic, $A^* = A$.

2. $[A \wedge B]^* = [A^* \wedge B^*]$, and similarly for the other propositional connectives.

3. $[(\forall x)A]^* = (\forall x)A^*$, and similarly for the existential quantifier.

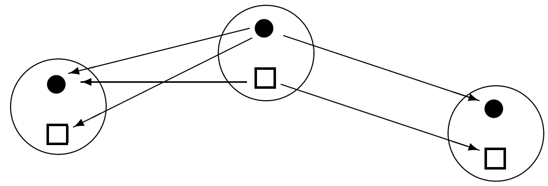

Figure 1: Counterpart Version

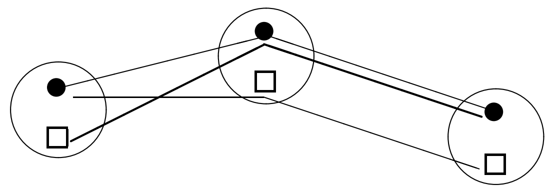

Figure 2: Function Version

4. Suppose the free variables of A are $x_1, \ldots, x_n$ (all object variables), and $g_1, \ldots, g_n$ are intension variables that do not occur in A^*. Recall the abbreviation of (4).

$$\begin{aligned}[\Box A]^* = (\forall g_1)\ldots(\forall g_n)\{[\mathbf{D}(g_1, x_1) \wedge \ldots \wedge \mathbf{D}(g_n, x_n)] \\ \supset \Box\langle\lambda x_1, \ldots, x_n.A^*\rangle(g_1, \ldots, g_n)\} \\ [\Diamond A]^* = (\exists g_1)\ldots(\exists g_n)\{[\mathbf{D}(g_1, x_1) \wedge \ldots \wedge \mathbf{D}(g_n, x_n)] \\ \wedge \Diamond\langle\lambda x_1, \ldots, x_n.A^*\rangle(g_1, \ldots, g_n)\}\end{aligned}$$

If $\mathcal{I}$ is an interpretation in the sense of a Lewis counterpart model, it is also an interpretation in the sense of a **FOIL** model (except that it assigns no values to relation symbols involving intension variables, something we can ignore for now). Similarly in the other direction. Now, the following is not hard to show.

Proposition 6.6 *Let $\langle\mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{C}\rangle$ be a simply connected Lewis counterpart frame, and let $\langle\mathcal{G}, \mathcal{R}, \mathcal{D}_O, \mathcal{D}_I\rangle$ be the companion **FOIL** frame. For any interpretation $\mathcal{I}$ on either frame, and for any formula Φ of the basic first-order modal language, Φ is valid in $\langle\mathcal{G}, \mathcal{R}, \mathcal{D}, \mathcal{C}, \mathcal{I}\rangle$ if and only if Φ^* is valid in $\langle\mathcal{G}, \mathcal{R}, \mathcal{D}_O, \mathcal{D}_I, \mathcal{I}\rangle$.*

What this means is that any logic characterized by a class of simply connected Lewis counterpart frames is also characterized (under translation) by a class of simply connected **FOIL** frames. When simply connected frames suffice, **FOIL** includes, under translation, Lewis counterpart semantics.

6.2 Riemannization

An embedding between Lewis counterpart frames and **FOIL** frames works well in the simply connected case, but what about frames that are not simply connected? The problem is that if a frame is not simply connected, it may not be possible to replace a counterpart relation by a collection of functions. Figure 3 shows the difficulties that can arise. Following the upper accessibility chain, the 'dot' object in world Γ corresponds to 'dot' which corresponds to 'dot' in world Δ, while following the lower chain it corresponds to 'box' which corresponds to 'box' in world Δ. A similar thing happens starting with the object 'box' of world Γ. Such counterpart relations cannot be treated as the union of a collection of functions, as we did when simply connected frames were considered.

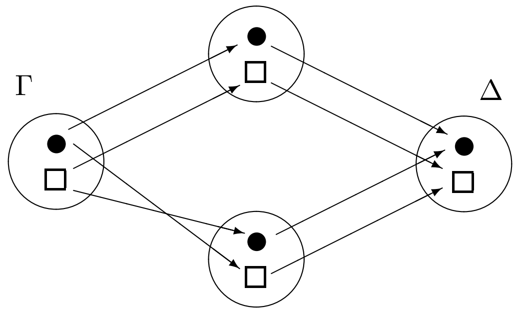

Figure 3: Not Simply Connected

One way of putting the difficulty is that replacing counterpart relations by functions can lead to many-valued functions, as it would in the example shown in Figure 3. A similar problem once arose in complex analysis with the square root function, the logarithm function, and many others. The solution was to break the complex plane into multiple copies of itself—into separate sheets—producing a Riemann surface. A similar solution applies here. To continue with the example of Figure 3, suppose we introduce two copies of world Δ. What makes them copies is that we will require any interpretation to assign the same relations to relation symbols at both copies. With Δ split in two in this way, the counterpart relationship can be 'functionalized,' as shown in Figure 4. Perhaps one could plausibly think of the split version of Δ as being a situation seen in two different ways (with the possibility of not recognizing we really have two aspects of the same thing). This leads us to the following generalization of the notion of **FOIL** frame.

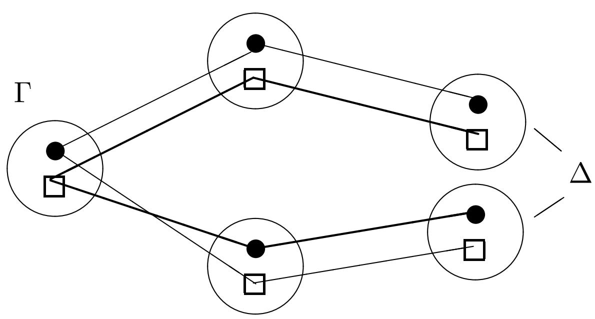

Figure 4: "Riemannized" Version

Definition 6.7 A *Riemann **FOIL** frame* is a pair $\langle \mathcal{F}, \mathcal{E} \rangle$, where $\mathcal{F} = \langle \mathcal{G}, \mathcal{R}, \mathcal{D}_O, \mathcal{D}_I \rangle$ is a **FOIL** frame as in Definition 6.1, and $\mathcal{E}$ is an equivalence relation on $\mathcal{G}$. If possible worlds $\Gamma, \Delta \in \mathcal{G}$ are in the relation $\mathcal{E}$ to each other, I'll say each is a *copy* of the other. A **FOIL** model $\langle \mathcal{G}, \mathcal{R}, \mathcal{D}_O, \mathcal{D}_I, \mathcal{I} \rangle$

is *based on* the Riemann **FOIL** frame $\langle\mathcal{F},\mathcal{E}\rangle$ provided the interpretation $\mathcal{I}$ *respects the equivalence relation* $\mathcal{E}$ in the following sense: for each relation symbol P and for each $\Gamma,\Delta\in\mathcal{G}$ that are in the relation $\mathcal{E}$, $\mathcal{I}(P)(\Gamma)=\mathcal{I}(P)(\Delta)$. A formula is *valid in a Riemann frame* if it is valid in every **FOIL** model based on the frame.

In effect, what makes two possible worlds copies is that in any model based on a Riemann **FOIL** frame atomic formulas behave the same at the two worlds. Of course beyond the atomic level this need not be the case. The notion of Riemann **FOIL** frame extends that of **FOIL** frame as used earlier—simply take as equivalence relation the one that relates each world to itself and to nothing else. Now, using the notions of frame for counterpart semantics from Definition 6.1 we have the following fundamental result.

Proposition 6.8 *For any Lewis or multiple counterpart frame $\mathcal{S}$ there is a Riemann **FOIL** frame $\mathcal{T}$ such that, for every formula Φ in the basic first-order modal language, Φ is valid in $\mathcal{S}$ if and only if Φ^* is valid in $\mathcal{T}$ (using the translation of Definition 6.5).*

An example illustrating Proposition 6.8 when Lewis counterpart frames are involved is given by Figure 3 and Figure 4. I'll leave details to you. The situation with multiple counterpart frames is, perhaps, of greater interest. Example 6.9 is based on an interesting such frame due to Oliver Kutz.

Example 6.9 Let $\mathcal{S}$ be the multiple counterpart frame $\langle\mathcal{G},\mathcal{R},\mathcal{D},\mathcal{C}\rangle$ specified as follows. $\mathcal{G}=\{\Gamma\}$, $\Gamma\mathcal{R}\Gamma$, $\mathcal{D}=\{a,b\}$, and $\mathcal{C}=\{C_1,C_2\}$, where $C_1=\{\langle a,a\rangle,\langle b,b\rangle\}$ and $C_2=\{\langle a,b\rangle,\langle b,a\rangle\}$. Oliver Kutz [19] provides an axiomatization of the logic of this frame (in the basic first-order modal language, involving equality), and shows it cannot be characterized by a class of Lewis counterpart frames. In short, multiple counterpart relations are essential.

Now, let $\mathcal{T}$ be the **FOIL** frame $\langle\mathcal{G}',\mathcal{R}',\mathcal{D}_O,\mathcal{D}_I\rangle$ sketched in the following diagram. In this, $\mathcal{G}'=\{\Gamma_1,\Gamma_2\}$, $\mathcal{R}'$ holds universally, $\mathcal{D}_O=\mathcal{D}=\{a,b\}$, and $\mathcal{D}_I$ consists of the two functions indicated in the diagram by the light and dark lines.

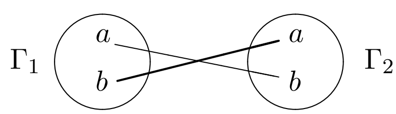

Finally, consider the Riemann **FOIL** frame $\langle\mathcal{T},\mathcal{E}\rangle$ where $\mathcal{E}$ is the universal equivalence relation on $\mathcal{G}'$ (thus being the same as $\mathcal{R}'$, though this is simply coincidence). I leave it to you to verify for this example that a basic modal formula Φ is valid in $\mathcal{S}$ if and only if Φ^* is valid in $\langle\mathcal{T},\mathcal{E}\rangle$.

As a particularly interesting special case, note that even though $\mathcal{S}$ is a one-world model, it does not follow that every formula of the form $X\supset\Box X$ is valid. For instance, $A(x,y)\supset\Box A(x,y)$ fails using an interpretation $\mathcal{I}$ that assigns to A at Γ the set $\{\langle a,b\rangle\}$, and using a valuation v such that $v(x)=a$ and $v(y)=b$. Now, $[A(x,y)\supset\Box A(x,y)]^*$ is the formula $A(x,y)\supset(\forall f)(\forall g)\{[\mathbf{D}(f,x)\wedge\mathbf{D}(g,y)]\supset\Box\langle\lambda x,y.A(x,y)\rangle(f,g)\}$. This fails at Γ_1 of the Riemann **FOIL** frame using an interpretation that agrees with $\mathcal{I}$ at both Γ_1 and Γ_2, and the same valuation v.

The status of the schema $\Box X\supset X$ is also worth thinking about.

I do not include a proof of Proposition 6.8 here. One can be found, using different terminology, in [16]. The methodology involved amounts to unwinding a counterpart frame.

7 Tableaus

Tableau systems using prefixes go back to [2] and [3]. A smoother, more uniform approach was developed in [23, 12]. In [8, 7] variations that allowed for intensional constant symbols and variables were studied at some length. For **FOIL**, prefixed tableaus give us a natural and simple proof procedure, at least for standard logics like **K**, **S4**, **S5**, and so on. I sketch the ideas here.

Prefixes are intended to be syntactic names for possible worlds, with a structure that reflects the accessibility relation in a convenient way. I'll present a system for **FOIL-K**; for several other **FOIL**-based logics modifications are straightforward, and for **FOIL-S5** things can be made considerably simpler, but **FOIL-K** will serve as a paradigm case.

A *prefix* is a finite sequence of positive integers, and a *prefixed formula* is an expression of the form $\sigma\,\Phi$, where σ is a prefix and Φ is a formula. I write prefixes as 1.2.3.2.1, for example. If σ is a prefix and n is a positive integer, $\sigma.n$ is σ with n adjoined. Intuitively, a prefix σ names a possible world in some model, and $\sigma\,\Phi$ says that Φ is true at the world σ names. The prefix $\sigma.n$ is intended to name a world accessible from the one that σ names. A tableau will contain prefixed formulas as node labels.

A tableau proof of a *closed formula* Φ is a tree with $1\,\neg\Phi$ at its root, that is 'grown' according to certain *Branch Extension Rules*, and that is *closed*. A tableau is closed if each branch is closed, and a branch is closed if it contains $\sigma\,X$ and $\sigma\,\neg X$, for some σ and some X. Here are the Branch Extension Rules for **FOIL-K**. In the following, σ is an arbitrary prefix. Also, if $\Phi(x)$ has been displayed, then $\Phi(y)$ denotes the formula $\Phi(x)$ with all free occurrences of x replaced by occurrences of y.

First the propositional rules, taking $\wedge$ as a representative binary connective.

$$\frac{\sigma\, X \wedge Y}{\begin{array}{c}\sigma\, X\\ \sigma\, Y\end{array}} \qquad \frac{\sigma\, \neg(X \wedge Y)}{\sigma\, \neg X \mid \sigma\, \neg Y} \qquad \frac{\sigma\, \neg\neg X}{\sigma\, X}$$

Next the modal rules.

$$\frac{\sigma\, \Diamond X}{\sigma.n\, X} \qquad \frac{\sigma\, \neg\Box X}{\sigma.n\, \neg X} \text{ where } \sigma.n \text{ is new to the branch}$$

$$\frac{\sigma\, \Box X}{\sigma.n\, X} \qquad \frac{\sigma\, \neg\Diamond X}{\sigma.n\, \neg X} \text{ where } \sigma.n \text{ already occurs on the branch}$$

For existential quantifiers the language is extended in the usual way. That is, for each of the types, object and intension, new alphabets of variables are added to the language—these are called *parameters*. Parameters, though variables, are never quantified, and only appear in proofs. In the following, y is an object parameter that is *new to the branch* and g is an intension parameter that is *new to the branch.*

$$\frac{\sigma\, (\exists x)\Phi(x)}{\sigma\, \Phi(y)} \qquad \frac{\sigma\, \neg(\forall x)\Phi(x)}{\sigma\, \neg\Phi(y)} \qquad \frac{\sigma\, (\exists f)\Psi(f)}{\sigma\, \Psi(g)} \qquad \frac{\sigma\, \neg(\forall f)\Psi(f)}{\sigma\, \neg\Psi(g)}$$

To take care of abstraction still more variables are introduced which will only appear in proofs. If f is a parameter of intension type and σ is a prefix, f_σ is a new variable of object type. Think of it as what the non-rigid symbol f designates at the world named by σ.

$$\frac{\sigma\, \langle\lambda x.\Phi(x)\rangle(f)}{\sigma\, \Phi(f_\sigma)} \qquad \frac{\sigma\, \neg\langle\lambda x.\Phi(x)\rangle(f)}{\sigma\, \neg\Phi(f_\sigma)}$$

For universal quantifiers we have more-or-less the expected rules. In the following g is any parameter of intension type, and y is either a parameter of object type, or a subscripted intension parameter whose subscript already occurs on the branch.

$$\frac{\sigma\ (\forall x)\Phi(x)}{\sigma\ \Phi(y)} \quad \frac{\sigma\ \neg(\exists x)\Phi(x)}{\sigma\ \neg\Phi(y)} \quad \frac{\sigma\ (\forall f)\Psi(f)}{\sigma\ \Psi(g)} \quad \frac{\sigma\ \neg(\exists f)\Psi(f)}{\sigma\ \neg\Psi(g)}$$

Example 7.1 Here is a sample proof in **FOIL-K**, of $(\forall f)[(\forall x)\Box P(x) \supset \Box\langle\lambda x.P(x)\rangle(f)]$.

$$\begin{array}{ll}
1 & \neg(\forall f)[(\forall x)\Box P(x) \supset \Box\langle\lambda x.P(x)\rangle(f)] \quad 1. \\
1 & \neg[(\forall x)\Box P(x) \supset \Box\langle\lambda x.P(x)\rangle(g)] \quad 2. \\
1 & \quad (\forall x)\Box P(x) \quad 3. \\
1 & \neg\Box\langle\lambda x.P(x)\rangle(g)] \quad 4. \\
1.1 & \neg\langle\lambda x.P(x)\rangle(g) \quad 5. \\
1.1 & \neg P(g_{1.1}) \quad 6. \\
1 & \quad \Box P(g_{1.1}) \quad 7. \\
1.1 & \quad P(g_{1.1}) \quad 8.
\end{array}$$

Line 2 is from 1 by an existential rule, with g as a new intension parameter; 3 and 4 are from 2 by a propositional rule; 5 is from 4 by a modal rule (1.1 is new to the branch); 6 is from 5 by an abstraction rule; 7 is from 3 by a universal rule ($g_{1.1}$ is of object type); 8 is from 7 by a modal rule (1.1 already occurs on the branch). Closure is by 6 and 8.

Soundness of the tableau system is straightforward. Here is a sketch of completeness. If Φ is not provable, systematically construct a tableau beginning with $1\,\neg\Phi$, so that every applicable rule is eventually applied. Choose an open branch of the resulting tableau. (Construction may run infinitely long, and König's Lemma may be needed here.) Use that open branch to construct a counter-model to Φ as follows. Take as worlds the prefixes occurring on the branch, with $\sigma.n$ accessible from σ. Objects are object parameters and subscripted intension parameters. Intensions are intension parameters. At world σ the intension parameter f designates the object f_σ. A relation symbol P is interpreted to be the mapping assigning to possible world σ the set $\{\langle t_1, \ldots, t_n\rangle \mid \sigma\, P(t_1, \ldots, t_n) \text{ is on } \sigma\}$. It is not hard to show this gives a model in which Φ is false (at world 1).

To extend the tableau system to allow for equality, two more rules must be added.

First, reflexivity: If t is a parameter of object type, or a subscripted parameter of intension type whose subscript already occurs on the branch, and σ is a prefix that occurs on the tableau branch, then $\sigma\,(t = t)$ can be added to the end of the branch.

Second, substitutivity: Let $\Phi(x)$ be a formula in which at most x occurs free, and let t and u be terms of object type. If $\sigma_1(t = u)$ and $\sigma_2\Phi(t)$ both occur on a tableau branch, $\sigma_2\Phi(u)$ can be added to the end.

Example 7.2 Here is a sample proof in **FOIL-K** with equality, of $(\forall f)\{\langle\lambda x.\Box\langle\lambda y.x = y\rangle(f)\rangle(f) \supset$

$[\Box\langle\lambda x.P(x)\rangle(f) \supset \langle\lambda x.\Box P(x)\rangle(f)]\}$.

$$
\begin{array}{ll}
1 & \neg(\forall f)\{\langle\lambda x.\Box\langle\lambda y.x = y\rangle(f)\rangle(f) \supset [\Box\langle\lambda x.P(x)\rangle(f) \supset \langle\lambda x.\Box P(x)\rangle(f)]\} \quad 1.\\
1 & \neg\{\langle\lambda x.\Box\langle\lambda y.x = y\rangle(g)\rangle(g) \supset [\Box\langle\lambda x.P(x)\rangle(g) \supset \langle\lambda x.\Box P(x)\rangle(g)]\} \quad 2.\\
1 & \quad \langle\lambda x.\Box\langle\lambda y.x = y\rangle(g)\rangle(g) \quad 3.\\
1 & \neg[\Box\langle\lambda x.P(x)\rangle(g) \supset \langle\lambda x.\Box P(x)\rangle(g)] \quad 4.\\
1 & \quad \Box\langle\lambda x.P(x)\rangle(g) \quad 5.\\
1 & \neg\langle\lambda x.\Box P(x)\rangle(g) \quad 6.\\
1 & \neg\Box P(g_1) \quad 7.\\
1.1 & \neg P(g_1) \quad 8.\\
1.1 & \quad \langle\lambda x.P(x)\rangle(g) \quad 9.\\
1.1 & \quad P(g_{1.1}) \quad 10.\\
1 & \quad \Box\langle\lambda y.g_1 = y\rangle(g) \quad 11.\\
1.1 & \quad \langle\lambda y.g_1 = y\rangle(g) \quad 12.\\
1.1 & \quad g_1 = g_{1.1} \quad 13.\\
1.1 & \neg P(g_{1.1}) \quad 14.
\end{array}
$$

Line 2 is from 1 by an existential rule; 3 and 4 are from 2, and 5 and 6 are from 4 by propositional rules; 7 is from 6 by an abstraction rule; 8 is from 7 by a modal rule; 9 is from 5 by a modal rule; 10 is from 9 by an abstraction rule; 11 is from 3 by an abstraction rule; 12 is from 11 by a modal rule; 13 is from 12 by an abstraction rule; 14 is from 8 and 13 by the substitutivity rule. Closure is by 10 and 14.

For other logics besides **K**, if there is a propositional prefixed tableau system it readily adapts to a **FOIL** version, as above. For **S5** things are particularly simple, since worlds can be taken to be mutually accessible. In this case the prefix structure can be simplified to just integers, instead of sequences of them.

8 Conclusion

What first-order modal logics are and ought to be is still a question very much under discussion. Mathematically, various versions of counterpart semantics have been intensively examined in recent years. I put **FOIL**-based logics forward as a counter-foil, so to speak. I suggest **FOIL** provides a very rich setting with a satisfying intuition. If Riemann **FOIL** frames are allowed, the expressive power is as great as with multiple counterpart semantics. Of course one might argue that having multiple copies of worlds has no corresponding intuition, but I think this is equally the case with having multiple counterpart relations. The question basically is, without such machinery added, how intuitively plausible and how easy to use is the resulting semantics. I believe a very good case can be made for **FOIL**. This paper amounts to a partial presentation of that case.

References

[1] M. D'Agostino, D. Gabbay, R. Hähnle, and J. Posegga, editors. *Handbook of Tableau Methods*. Kluwer, Dordrecht, 1999.

[2] M. C. Fitting. Tableau methods of proof for modal logics. *Notre Dame Journal of Formal Logic*, 13:237–247, 1972.

[3] M. C. Fitting. *Proof Methods for Modal and Intuitionistic Logics*. D. Reidel Publishing Co., Dordrecht, 1983.

[4] M. C. Fitting. Databases and higher types. In J. Lloyd, V. Dahl, U. Furbach, M. Kerber, K.-K. Lau, C. Palamidessi, L. M. Pereira, Y. Sagiv, and P. Stuckey, editors, *Computational Logic—CL 2000*, pages 41–52. Springer Lecture Notes in Artificial Intelligence, 1861, 2000.

[5] M. C. Fitting. Modality and databases. In R. Dyckhoff, editor, *Automated Reasoning with Analytic Tableaux and Related Methods*, pages 19–39. Springer Lecture Notes in Artificial Intelligence, 1847, 2000.

[6] M. C. Fitting. Modal logics between propositional and first-order. *Journal of Logic and Computation*, 12:1017–1026, 2002.

[7] M. C. Fitting. *Types, Tableaus, and Gödel's God*. Kluwer, 2002.

[8] M. C. Fitting and R. Mendelsohn. *First-Order Modal Logic*. Kluwer, 1998. Paperback, 1999.

[9] D. Gallin. *Intensional and Higher-Order Modal Logic*. North-Holland, 1975.

[10] S. Ghilardi. Incompleteness results in Kripke semantics. *Journal of Symbolic Logic*, 56:516–538, 1991.

[11] S. Ghilardi. Quantified extensions of canonical propositional intermediate logics. *Studia Logica*, 51:195–214, 1992.

[12] R. Goré. Tableau methods for modal and temporal logics. 1996. In [1].

[13] G. E. Hughes and M. J. Cresswell. *An Introduction to Modal Logic*. Methuen, London, 1968.

[14] M. Kracht. How completeness and correspondence theory got married. In M. de Rijke, editor, *Diamonds and Defaults*, Dordrecht, 1993. Kluwer.

[15] M. Kracht and O. Kutz. The semantics of modal predicate logic I, counterpart frames. In F. Wolter, H. Wansing, M. de Rijke, and M. Zakharyaschev, editors, *Advances in Modal Logic*, volume 3, Stanford, 2001. CSLI Publications.

[16] M. Kracht and O. Kutz. The semantics of modal predicate logic II, modal individuals revisited. In R. Kahle, editor, *Intensionality. Proceedings of a conference held at Munich, 27th–29th October 2000*, Lecture Notes in Logic. A. K. Peters, 2003. To appear.

[17] S. Kripke. Semantical considerations on modal logics. In *Acta Philosophica Fennica, Modal and Many-valued Logics*, pages 83–94, 1963.

[18] O. Kutz. Kripke-type Semantiken für die modale Prädikatenlogik. Master's thesis, Humboldt-Universität zu Berlin, 2000.

[19] O. Kutz. A logic that needs multiple counterpart relations. Private communication, 2001.

[20] O. Kutz. New semantics for modal predicate logics. In *Foundations of the Formal Sciences II*, Trends in Logic. Kluwer, 2002.

[21] D. Lewis. Counterpart theory and quantified modal logic. *Journal of Philosophy*, 65:113–126, 1968.

[22] D. Lewis. Counterparts of persons and their bodies. *Journal of Philosophy*, 68:203–211, 1971.

[23] F. Massacci. Strongly analytic tableaux for normal modal logics. In A. Bundy, editor, *Proceedings of CADE 12*, volume 814 of *Lecture Notes in Artificial Intelligence*, pages 723–737, Berlin, 1994. Springer-Verlag.

[24] R. Montague. On the nature of certain philosophical entities. *The Monist*, 53:159–194, 1960. Reprinted in [31], 148–187.

[25] R. Montague. Pragmatics. pages 102–122. 1968. In *Contemporary Philosophy: A Survey*, R. Klibansky editor, Florence, La Nuova Italia Editrice, 1968. Reprinted in [31], 95–118.

[26] R. Montague. Pragmatics and intensional logic. *Synthèse*, 22:68–94, 1970. Reprinted in [31], 119–147.

[27] H. Shirasu. Duality in superintuitionistic and modal predicate logics. In M. Kracht, editor, *Advances in Modal Logic*, pages 223–236, Stanford, 1998. CSLI Publications.

[28] D. P. Skvortsov and V. B. Shehtman. Maximal Kripke-type semantics for modal and superintuitionistic predicate logics. *Annals of Pure and Applied Logic*, 63:69–101, 1993.

[29] B. H. Slater. Modal semantics. *Logique & Analyse*, 127-128:195–209, 1989.

[30] N.-Y. Suzuki. Algebraic Kripke sheaf semantics for non-classical logics. *Studia Logica*, 63:387–416, 1999.

[31] R. H. Thomason, editor. *Formal Philosophy, Selected Papers of Richard Montague*. Yale University Press, New Haven and London, 1974.