Hayaki, Reina (2003). “Actualism and Higher-Order Worlds”. *Philosophical Studies* **115** (2), 149-178.

Actualism and Higher-Order Worlds

Actualism is the ontological thesis that everything that exists is actual. Although it seems so common-sensical as to be platitudinous, it has been attacked as being inadequate on various counts. In particular, opponents have alleged that actualism is incompatible with the standard semantics for quantified modal logic, because it cannot handle iterated modality. I shall argue that an honest actualist can accommodate both iterate modality and quantified modal logic generally by adopting a revised semantics.

First, a point of terminology: there are at least two distinct theses that sometimes go by the name of ‘actualism’. As I shall use the term, actualism is contrasted with possibilism, according to which there exist (in some broad sense) objects which are strictly non-actual. They do not inhabit this world, but they do inhabit others, so they exist outright by virtue of existing somewhere in modal space. The most famous proponent of possibilism is, of course, David Lewis, especially in Lewis (1986). This issue is distinct from the issue of whether every possible object is in fact (identical to) an actual object, i.e., whether all possible worlds contain exactly the same inhabitants. A positive view on this latter issue is occasionally called ‘actualism’ as well. Actualism_2 (fixed domains across worlds) implies actualism_1 (no non-actual objects), but not vice versa. In this paper I shall use the term ‘actualism’ to refer only to actualism_1. Interestingly, the problems that confront the actualist_1 are easily dissolved if one also adopts actualism_2. I shall not do so, however.

1. Actualism

It seems that there could have been objects other than those that actually exist. For example, I could have had an older brother in addition to the one younger sister I actually have. If we assume the popular doctrine of the necessity of origin, any such brother would be a non-actual object: an object not identical to any actual object.

Can we really talk about non-actual objects in this way?

For possibilists, there is no problem. Modal space is filled with non-actual objects like this older brother of mine. Non-actual objects exist, in some real sense, just as actual ones do. When we ordinarily speak of objects existing or not existing, we mean that they do or don't exist in *this* world; so we would regard as false a statement like "Blue horses exist". But according to possibilists, blue horses *do* exist, just not here; they exist in other worlds, or in other ways, whatever those might turn out to be. Apart from not inhabiting this world, non-actual but possible objects are really quite similar to actual objects in the kind of properties they have and the kind of relations they enter into.

However, my view is that all such talk is, at best, a mere *façon de parler*. We cannot truthfully talk of non-actual objects, because such objects do not exist; to say that they do is to commit oneself to an undesirable account of metaphysical possibility. I speak falsely if I say of an older brother of mine that "*he* is a non-actual object". There is no such object as *an older brother of mine*, so the word 'he' fails to refer.[1] All apparent references to non-actual objects are circumlocutions either for *de dicto* statements about ways the world might have been, or for *de re* statements about ways that actual objects might have been, or for some combination of both.

I could have had an older brother as well as the younger sister that I actually have. Using a little formal symbolism, we might express this as follows (I am RH and my sister JH):

(1) Possibly ∃*x*(*x* is RH's and JH's older brother).

The possibilist believes that (1) licenses the inference to

(1′) ∃*x* Possibly (*x* is RH's and JH's older brother),

if in (1′) '∃' is interpreted as ranging over all possibilia.[2] This is essentially because, for the possibilist, the possibility operator is a kind of existential quantification. David Lewis (1968) converts standard quantified modal logic to his counterpart theory by (among other steps) replacing the operator '◇' with existential quantification over possible worlds. In first-order logic '∃*w*∃*xFwx*' is equivalent to '∃*x*∃*wFwx*'; so '∃*w*∃*x*(object *x* is such-and-such at world *w*)' is equivalent to '∃*x*∃*w*(object *x* is such-and-such at world *w*)'.

On the other hand, the actualist does not interpret quantifiers or modal operators in this way. Quantifiers range not over possibilia, but over actualia only (because that's all there is), and the possibility operator means 'There is a possible world *according to which*...'. This is a context which does not permit the extraction of objects outside the scope of the operator. There are many such contexts, all well-known, in which reification of a so-called 'intensional object' is forbidden. I might *believe* that there are spies, without believing of any specific person that he or she is a spy (Quine 1956); a writer might tell a *fictional story* in which there is a violin-playing detective, without the story being about any actual person's crime-solving and violin-playing; in a non-factive sense of 'see', Macbeth *saw* that there was a dagger before him, the handle toward his hand (*Macbeth* II.i.33-34), but since he was hallucinating there was no object that he saw as a dagger before him (a particularly ghostly kind of dagger, "a dagger of the mind" (II.i.38)).

Reification is also illicit for modality. From the mere fact that I could have had an older brother, it does not automatically follow that there exists (i.e., exists in this world) an object which is possibly my older brother. This inference would be valid if we accepted the Barcan Formula (Marcus 1946):

(BF) $\Diamond\exists x\phi x \supset \exists x\Diamond\phi x$

However, most actualists, including myself, do not accept the Barcan Formula.[3] Without independent grounds for doing so, inferring the existence of possible objects from the existence of possibilities is illicit. There simply is no object *of which* it can be said that possibly *it* is RH's and JH's older brother. The sentence 'I could have had an older brother' means only that there is a possible world which represents me as having a certain property I don't actually have, namely that of sharing parents with an object that is male and older than me. The upshot is that we may not reify non-actual objects, any more than we may reify 'objects of thought'.

2. Iterated modality and nested worlds

All well and good — but the actualist faces a rather nasty stumbling-block with the problem of iterated modality. The following formulation of the problem is due to Alan McMichael (1983).

Suppose you are an actualist, but you want to help yourself to possible-worlds talk. Necessity is truth according to all possible worlds, and possibility is truth according to some possible world. (Insert 'accessible' if you don't like S5.) But as an actualist, you think that these 'possible worlds' are actual objects — something like maximal consistent sets of sentences in an idealized language, sets of states of affairs (Plantinga 1974), or some other kind of abstract object. They cannot have non-actual objects as constituents, in the way that the actual world has actual objects as constituents.

The charge is that you cannot make sense of iterated modal statements. The problem sentence (adapted from McMichael) is:

(2) I could have had a brother who was not a concert pianist but could have been one (if he had practised harder).

Clearly (2) is true. For convenience we will ignore the parenthetical comment about practising harder, and symbolize (2) as follows, where ‘*Bx*’ abbreviates ‘*x* is my brother’ and ‘*Px*’ abbreviates ‘*x* is a concert pianist’:

(2′) $\Diamond\exists x(Bx \wedge \neg Px \wedge \Diamond Px)$.

You are an actualist, so you paraphrase (2′) as something like the following:

(2″) There is a world-story, *w*, according to which I have a brother, *x*, who isn’t a concert pianist; and there is a world-story, *w*′, (accessible from *w*,) according to which *x* is a concert pianist.

You take *w* to be something like a very long (infinite) list of formulas, including ‘$\exists x(Bx\wedge\neg Px\wedge\Diamond Px)$’. *w*′ is a different list that includes ‘$\exists x(Bx\wedge Px)$’, or maybe just ‘$\exists xPx$’, since the pianist could fail to be my brother (perhaps because I don’t exist at that world). Our world contains both lists. If all that mattered for the truth of (2) were the inclusion of these formulas in their respective lists, you would be done. But what is required is that there be a special connection between the list *w* and the list *w*′: the formula ‘$\exists x(Bx\wedge\neg Px\wedge\Diamond Px)$’ in *w* and the formula ‘$\exists x(Bx\wedge Px)$’ (or ‘$\exists xPx$’) in *w*′ must be *about* the same object *x*. If, as before, we accept the necessity of origin so that no actual person will qualify as a truth-maker for (2), there is no such *x*! ‘*x*’ does not refer to a genuine object; it is merely a place-holder in a very long story, and then again, completely separately, in another. It is, of course, completely accidental that the same typographical character should have been used in both formulas to represent the bound variable; there is no connection between the two. There is no means by which a place-holder in world *w* can be identified with a place-holder in a different world, *w′*. A possibilist might use a name, but you are an actualist and may not. How can you name things that don't exist? You must either give up iterated modality and the traditional semantics, or give up actualism.

Clutching at straws, you might suggest trying to finesse the iterated modality problem by using 'dummy names' rather than variables. Whenever a non-actual world *w* the sentence '∃!*xFx*' is made true at a non-actual world *w* by an object distinct from every actual object, we can introduce a constant, or 'dummy name', say '*a*', by the addition of the formula 'ι*xFx* = *a*' to the world *w*. Then the dummy name would become available to other worlds, too, which could then include more sentences using '*a*', e.g., '¬*Fa*'. The worlds would then be 'talking about the same object'.

I am not unsympathetic to this manoeuvre, and indeed shall employ a somewhat similar one myself. But the mere introduction of dummy names is not an adequate solution to the iterated modality problem. A dummy name (also known as a dummy *variable*) is simply not the same thing as a name. Informally, the idea is that the range of a dummy variable cannot expand beyond the range for which it was originally stipulated: it cannot do more work than it was supposed to. For example, one might say at the start of a proof in geometry, "Consider Triangle *ABC*". 'Triangle *ABC*' is a dummy name. There is no triangle that we're considering; or rather, there is no triangle such that we are considering *it*. Even worse, take the proof that there is no largest prime: "Suppose that there were a largest prime; call it *n.* Consider the number *n*!+1…." Here '*n*' does not name anything at all. Dummy names simply are not names.

Suppose one is performing a truth tree derivation in first-order logic. One applies the rule for existential instantiation to a formula '∃*xFx*' by choosing a constant not previously used in the branch, and derives a formula such as '*Fa*'. One may not then export the name '*a*' for use in another branch or, worse yet, in another derivation altogether. I shall have more to say later about dummy variables and stipulations, but the illegitimacy of exporting a name from one branch of a derivation to another translates precisely to the illegitimacy of exporting a dummy name, introduced in one world by a simple definite description, to another world.[4]

Various suggestions have been made as to how the actualist can respond to the iterated modality problem. They usually involve importing non-actual objects into the actual world, either surreptitiously[5] or openly.[6] In contrast, my approach does not incur such an ontological cost.

The problem we are trying to solve is how the actualist can ascribe modal properties to non-actual objects — or, in actualistically acceptable terms, how to explain the truth of statements involving iterated modality. Let us first consider ascriptions of modal properties to actual objects. Standardly, a sentence such as

(3) Humphrey could have won

is analysed as

(3′) There is a possible world according to which Humphrey won.

The way this is sometimes put is that the ascription of a modal property to Humphrey — that of possibly having won — amounts to identifying the losing Humphrey of this world with the winning Humphrey of another world.[7] This other world, being abstract, doesn't include Humphrey as a physical constituent, but is *about* Humphrey; and that it is (partly) about Humphrey is a matter of brute stipulation: "'possible worlds' are *stipulated*, not *discovered* through powerful telescopes" (Kripke 1980, pp. 44). Thus the so-called 'problem of trans-world identity' is defused as having put the cart before the horse.

Brute stipulation guarantees that a world is *about* what one intends it to be about when making a *de re* modal statement about an actual object. My solution to McMichael's problem is: the same goes for iterated modality (apparent *de re* ascriptions of modal properties to non-actual objects).

This is not quite as straightforward as it sounds. It should not be surprising if the asymmetry between actual objects and non-actual ones (recall that the latter do not exist!) were reflected in an asymmetry between apparent ascriptions of modal properties to them. An apparent *de re* ascription of a simple (non-iterated) modal property to an actual object is exactly what it looks like. An apparent *de re* ascription of a simple (non-iterated) modal property to a non-actual object is really a description of a world within a world, and is hence a stipulation within a stipulation. We will need to quantify over worlds-within-worlds, rather than over all worlds outright, when we analyse apparent ascriptions of modal properties to non-actual objects. When I say that I could have had a brother who was not a concert pianist but could have been one, what I mean is that there is a world-story according to which I have a brother who is not a pianist, and it's true in that world-story that there is a nested world-story according to which *that very person* is a concert pianist. This second world-story has exactly the same relation to the first as the first does to the actual world: it is something stipulated to enable the ascription of a modal property to an object in the outer world.

To account for iterated modality, what we need are nested world-stories. In exactly the same way that the modal properties of actual objects are explicated by the existence of suitable first-order world-stories, the modal properties of non-actual objects in *n*th-order world-stories are explicated by the existence of suitable (*n*+1)-st order world-stories. What is crucial is that we abandon the traditional view that worlds are all first-order.

The much simpler case of fiction would probably be user-friendly here. I have, for this exposition, written a very short (and poor) fictional short story called ‘The No-Good Novelist’, the entirety of which is as follows:

> There lives somewhere on the island of Manhattan a tormented novelist called Barton. The president of the U.S. is Fillard Millmore, an ineffectual buffoon. Barton spends all his leisure time writing letters of protest to President Millmore and fan letters to the novelist J. D. Salinger, who is his hero. Neither Millmore nor Salinger ever writes back. Barton’s magnum opus is a novel called *Heroes and Villains*; it is set on a desolate Earth populated only by Fillard Millmore and a Martian (a fictional character created by Barton), each of whom writes a book.

First-order world-stories can be *about* the right objects by brute stipulation; so can first-order fictional stories. ‘The No-Good Novelist’ is, at least partly, *about* Manhattan and J. D. Salinger, by authorial intent and the appropriate causal links between those objects and my use of their names. (It is also, of course, ‘about’ Barton and President Millmore, in a secondary sense: the story implies the existence of such people.) Some actual objects survive through to *Heroes and Villains* (e.g., the planet Earth). Some (J. D. Salinger) make it to ‘The No-Good Novelist’, but not to *Heroes and Villains.* The interesting case is what happens to the first-order fictional characters, i.e., the ones that I (the author) created for ‘The No-Good Novelist’. According to ‘The No-Good Novelist’, Barton includes at least one of them, namely Fillard Millmore, in *his* novel, *Heroes and Villains*; he writes *about* Millmore. He has no more trouble doing this, according to my story, than I did in writing about Manhattan and J. D. Salinger. I used authorial intent to stipulate that my story was about Salinger; according to my story, Barton uses authorial intent to stipulate that his story is about various objects in his world (Millmore, the Earth). Furthermore, the nesting can go indefinitely wide and deep: according to my story, Barton writes a story according to which Millmore and a Martian both write stories.[8]

Nested world-stories are structurally comparable to nested fictions: there is an outermost world-story, containing many first-order world-stories, each of which contains second-order world-stories … and so on.[9]

Let's spell out the details of this nesting relation. The relation *S*, or **stipulation**[10], holds between an outer and an inner (nested) world. The nesting relation *S* is distinct from the accessibility relation *R* that appears in Kripke models; the connection between *S* and the traditional accessibility relation, and between my models and Kripke models, will soon become apparent.

The actual world is the one and only level 0 world; let us call it 'α'. In α (that is, according to α), it is true that I could have had a brother who was in fact a banker, but could have been a concert pianist. This involves stipulating a possible world, w_1, according to which I have a brother who is a banker but could have been a pianist. w_1 is a level 1 world, at which it is true that *this very brother*, previously stipulated by α, could have been a concert pianist. Hence we must stipulate a yet more deeply nested world, w_2, at which *this very brother* is a concert pianist.

It is time to come clean about what, or who, is involved in stipulation. We most emphatically do not want an account of modality on which the truth of a modal sentence requires some person or other to perform a mental act, however untaxing. Stipulation cannot be a purposive act (as story-telling is). Rather, stipulation is a function of the way in which language works. The truth of the sentence 'I could have had a banker brother' does not depend on *my* making a stipulation, or anyone else's making a stipulation; it is stipulated by the conventions of the world-making language, by the way in which the language is used.

Perhaps this is a little obfuscatory, but for an honest actualist it will have to do. It seems to be inevitable that an actualist must take modality as primitive.[11] A world counts as a *possible* world only if the situation it describes is consistent, and for this a purely syntactic notion of consistency will not suffice: the concept must be metaphysical. In other words, we cannot fully analyse modality away. ‘Reina Hayaki could have had a banker brother’ is true because there is a *consistent* world-story in which I have a banker brother. That this world-story counts as metaphysically consistent is a function of the way in which we *in fact* use the name ‘Reina Hayaki’ and the word ‘could’. The rules governing the use of the name ‘Reina Hayaki’ allow a certain latitude in the properties that its referent can have. Its referent can have a banker brother; its referent cannot be a fried egg. The world-story in which I have a banker brother is stipulated, not by a world-story-teller, but by the rules of our language.

The formal relation S is, of course, a relation between worlds, not between some set of linguistic rules and a world. This is not a problem. The rules of our language generate modal possibilities; and once possibilities — i.e., consistent world-stories — are stipulated, we can talk about possible worlds to which the actual world bears the relation S. It is possibly true that ϕ iff there is a world S-related to (‘stipulated by’) the actual world, at which it is true that ϕ.

In short: if a proposition ϕ is possibly true at a world w, this involves stipulating a world w' at which ϕ is true. w bears the relation S to w'. If, on the other hand, ϕ is necessarily true at w, then ϕ is true in all worlds w' to which w bears S.

So far, S is behaving very much like the Kripkean accessibility relation. However, it will shortly become apparently that the quasi-accessibility relation S is far more restrictive than the traditional accessibility relation R, being bound by conditions that do not apply to R. Some of these conditions must be in place for propositional modal logic, while others come into play only with quantification.

3. Properties of *S*, and AA models

The relation *S* for nested possible worlds must satisfy the following conditions, for both propositional and quantified modal logic.

(A) *S* forms a tree: it is (1) generated, (2) antisymmetric and (3) anticonvergent;

(B) *S* is irreflexive;

(C) *S* is intransitive.

(A1) S is generated. A generated relation has a point of origin which is related to all other points in a finite number of steps. *S* is generated because the actual world is the only world of level 0. All other possible worlds are *S*-related to the actual world in a finite number of steps; they are nested a finite number of levels deep. A world stipulated by the actual world is a level 1 world; a world stipulated by a level 1 world — a (sub–)stipulation within a stipulation — is a level 2 world; and so on.

(A2) S is antisymmetric. If a relation is antisymmetric, then for no two distinct elements does the relation hold in both directions. Suppose there are two distinct possible worlds, w_1 and w_2, such that w_1 stipulates w_2. w_2 is then a stipulation within a stipulation and cannot in turn stipulate w_1, since it is contained within this outer stipulation. w_2 could perhaps stipulate a world very, very similar to w_1, but it could not stipulate w_1 itself.

It may be objected that antisymmetry is an undesirable restriction on *S*, for it precludes symmetry if *S* holds between at least one pair of distinct worlds. Many people would be loath to give up the 'Brouwer axiom', a familiar and much-beloved axiom of modal logic that guarantees the symmetry of the accessibility relation *R*:

(**B**) $\phi \rightarrow \Box \Diamond \phi$

It would seem that requiring *S* to be antisymmetric involves rejecting this axiom. This, surely, is too high a price to pay for a supposedly superior semantics.

Indeed, I want my revised semantics to be as metaphysically neutral as possible, except of course for its commitment to actualism. Suppose it turned out that the restrictions on *S* were so intrusive as to settle by fiat all debates regarding the nature of possibility, e.g., whether **B** is an appropriate an axiom, or S5 the appropriate logic, for metaphysical modality; and that the chips fell in a way that most metaphysicians regard as counterintuitive. In that case there would be very little motivation for adopting the revised semantics.

Fortunately, the price is not nearly so high. We can adopt the revised semantics, despite all the apparent intractability of the stipulation relation *S*, without thereby committing ourselves to accepting or rejecting any particular modal logic, at least at the propositional level.

In particular, our revised semantics does not exclude the modal logic B. Suppose the axiom **B** holds at the actual world (α). For an arbitrary true sentence ϕ, $\Diamond\phi$ is true at all worlds *w* accessible from α (i.e., all level 1 worlds). Thus for every world *w* stipulated by α, there must be a world w' stipulated by *w* at which ϕ is true. Traditional semantics would guarantee this by the symmetry of the accessibility relation *R*. From every *w* such that $\alpha R w$, we can return to α ($wR\alpha$). If ϕ is true in the actual world, $\Diamond\phi$ is true in every world *R*-related to the actual world: $\Box\Diamond\phi$ is true in the actual world.

Clearly, we cannot do this with *S*. What we *can* do, however, is to have a *copy* of the actual world stipulated by each level 1 world. Suppose we decided that we approve of axiom **B**, and want to see it validated semantically. The motivation would presumably be that any situation which is actual should be in principle recoverable, even if lost in a possible world.[12] Rather than requiring, nonsensically, that all level 1 worlds stipulate the level 0 world, what we need to require is that every level 1 world stipulate a level 2 world at which exactly the same sentences are true as in the level 0 world.

To generalize: if we want axiom **B** to be validated, what we need is that for every world of level *n* which stipulates worlds of level *n*+1, each of those worlds of level *n*+1 should stipulate a world of level *n*+2 which duplicates the level *n* world. Each world has its (single) level and its place in the nesting structure essentially.

(A3) S is anticonvergent. An anticonvergent relation is one where branches, once separated, do not rejoin: if neither of two distinct elements is related to the other, there is no third element, distinct from both the two, to which both are related.

Suppose that for some worlds *a*, *b* and *c* (where $b \neq c$), we have *aSb* and *aSc*. We cannot have a world *d* such that *bSd* and *cSd*. If there were such a world *d*, it would be a stipulation within world *b* and, simultaneously, a stipulation within world *c*. The only way for these stipulations-within-stipulations to coincide would be for the outer stipulations (*b* and *c*) to coincide themselves: that is, for *b* and *c* to be the same world.

A comparison with other intentional objects may be in order here. In a famous discussion, P. T. Geach (1967) considers sentences like the following, which describes some (fictional) inhabitants of a village gripped by ‘witch mania’:

(4) Hob thinks a witch has blighted Bob’s mare, and Nob wonders whether she (the same witch) killed Cob’s sow.

Sentences like (4) are notoriously intractable if we take them at face value. We cannot analyse (4) as

(4*a*) $\exists x$(Hob thinks *x* is a witch and blighted Bob's mare $\wedge$ Nob wonders whether *x* killed Cob's sow),

because there may not be a specific person whom Hob suspects of witchery and mare-blighting and whose possible culpability for sow-killing is entertained by Nob. It is hard to see how Hob's belief can be about the same person that Nob's wondering is about if there are no external criteria by which we can judge (such as Hob's belief and Nob's wondering each being about a specific woman whom they can identify).[13] A small step forward would be to assimilate (4) to

(4*b*) Hob thinks a witch has blighted Bob's mare, and Nob wonders whether the witch who blighted Bob's mare killed Cob's sow.

The putative 'intentional identity' — between the witch who blighted Bob's mare and the witch who (may have) killed Cob's sow — is now at least confined to one mind, Nob's. But even if Hob is taken out of the picture, we still don't have a satisfactory analysis for the remainder of the sentence. The following is inadequate, for the same reason as (4*a*):

(4*c*) $\exists x$(Nob thinks that (*x* is a witch and blighted Bob's mare) $\wedge$ Nob wonders whether *x* killed Cob's sow).

There may be no actual person who is the object of Nob's thinking and wondering. Perhaps we could try:

(4*d*) Nob thinks that $\exists x$(*x* is a witch $\wedge$ *x* blighted Bob's mare $\wedge$ *x* may have killed Cob's sow).

We have lost the flavour of the original phrase 'wonders whether', but this seems as close as we can get. The most notable aspect of this attempt is that, in order to make sense of 'intentional identity', the 'intentional object' had to be shoehorned into one intentional context. The same would be the case if we adopted the approach of the previous footnote and concluded that (4) is true only if Hob's belief and Nob's wondering are both causally derived from the same witch-myth.

In the same spirit, if world *a* stipulates *b*, which in turn stipulates *d*; and *a* also stipulates a different *c*, which *seems* to stipulate the very same *d*; then we must insist that *c* does not in fact stipulate *d* but a doppelgänger of *d*.

As with antisymmetry, it might be objected that anticonvergence is an unwanted restriction on *S*, because it would invalidate the axiom **G**:

(**G**) $\Diamond\Box\phi \rightarrow \Box\Diamond\phi$

The logic S4.2, which is S4 with axiom **G** added, is complete for Kripke models with (transitive and) convergent accessibility relations: that is, whenever *aRb* and *aRc*, then there is a *d* such that *bRd* and *cRd*. Branches that separate must immediately rejoin. S4.2 is a proper subset of S5, considered by many to be the correct logic for metaphysical necessity. If we reject S4.2 by rejecting the axiom **G**, we must reject S5.

Happily, we are not bound to this conclusion. Suppose we decide that we like the axiom **G**, whether out of allegiance to S5 or on its own merits. Rather than requiring that separated branches literally rejoin, we require that separated branches continue in (qualitatively) identical ways.

(B) S is irreflexive. This is perhaps the most unpalatable of my restrictions on *S*, an unpalatability that corresponds to the strength of our intuition about reflexivity in Kripke models. Indeed, Kripke himself builds the reflexivity of 'relative possibility' (accessibility) into his semantics (Kripke 1963), although the requirement can easily be removed. If we have any intuitions about stipulatability, surely it is that, given a particular possible world, we can there stipulate that very same possibility again. We want the axiom **T** to be validated:

(**T**) $\Box\phi\rightarrow\phi$

If we want to retain **T** (as I think we do), what we need is not that every world stipulate itself, which simply cannot be done. Rather, we insist that every world stipulate a duplicate of itself: a world at which exactly the same sentences are true. The duplicate is not, of course, a perfect duplicate. It will be a world one level higher (one level more deeply nested) than the original.

Although **T** is as plausible as axioms of modal logic get, it is not built into the semantics, for the sake of maximum flexibility and neutrality.

(C) S is intransitive. Although the requirement of intransitivity should be obvious from what has gone before, I shall run through it briefly for completeness' sake.

Transitivity is a popular requirement for the accessibility relation *R*,[14] and perhaps it should also be a requirement for *S*. If *b* is stipulated by *a*, and *c* is stipulated by *b*, shouldn't *c* be directly stipulatable from *a*? Even more pressingly, isn't *c already* stipulated by *a*, by being stipulated by a world stipulated by *a*?

In traditional models, the transitivity of accessibility is guaranteed by the axiom **4**:

(**4**) $\Box\phi\rightarrow\Box\Box\phi$

I have no particular complaint with (**4**). If one wishes it to hold in a revised model, the right way to ensure this is not to have the stipulation relation *S* be transitive, but to require that whenever *aRb* and *bRc*, *a* directly stipulates a duplicate of *c* (rather than *c* itself). Again, the duplicate is not a perfect copy, as it will be one level lower (less nested) than the original. Nested worlds, like nested fictions, have their levels and their places in the nesting structure necessarily. Worlds are thus strictly regimented.

With the addition of quantification comes a final restriction on *S*. Actually, it is not so much a restriction on *S* but on the inhabitants of worlds. To express it, we need some new terms. Let S^* be the ancestral of the stipulation relation *S*: that is, $w_1S^*w_n$ iff $w_1S^nw_n$ for some $n>0$. In other words, there is a finite sequence of worlds $w_2, \ldots, w_{n-1}$, such that w_1Sw_2, w_2Sw_3, …, and $w_{n-1}Sw_n$. When wS^*w', I shall say that *w* is an **ancestor** of w', or w' is a **descendant** of *w*. The final restriction on *S* is as follows:

(D) If the same non-actual object is present in two distinct worlds, then either one is an ancestor of the other, or there is a third world which is an ancestor of both. (If the same non-actual object is present in each of two distinct worlds w_1 and w_2, then either $w_1S^*w_2$, or $w_2S^*w_1$, or $w_3S^*w_1$ and $w_3S^*w_2$ for some world w_3.)

Equivalently:

(D′) Suppose a non-actual object *z* is first introduced at a world *w*: that is, *z* is not present in any ancestor of *w*. Then *z* can appear only in *w* or its descendants.

Condition (D′) has the effect of confining non-actual objects to a single branch[15], just as the closest accurate analysis of sentence (4), namely (4d), confined the ‘object of thought’ to a single belief context.

The motivation for (D′) is this. I argued earlier that *S* is anticonvergent, i.e., that the same higher-order world (an ‘intentional object’) cannot be stipulated by two distinct lower-order worlds. For analogous reasons, we cannot have the same ‘non-actual object’ appearing in two distinct worlds at the same level. The only way that a ‘non-actual object’ can gain a foothold in two distinct worlds is by being part of one large stipulation that covers both worlds: that is, if one of the worlds is nested within the other, or they are both nested within a third, lower-level world.

The upshot of condition (D′) is that non-actual objects introduced for the first time in two different branches can never be identical. In fact, if questions should arise about the identity or distinctness of non-actual objects in different branches, it would be more appropriate to rule such questions illegitimate, than to answer them in the negative. However, this is an imperfection I shall ignore in order to streamline the semantics as much as possible. It can be motivated, perhaps, by considering whether Hob's witch and Nob's witch are identical. The question may be improper, but if an answer had to be given it should clearly be in the negative (barring the presence of common external factors that ground their mental acts).

A more basic problem for the committed actualist might be the legitimacy not of non-identity statements involving non-actual objects, but of *any* talk about such objects. They are not supposed to be there to be talked about, let alone to be named ('call it *z*').

Indeed, we cannot name non-actual objects. However, we *can* give them dummy names. As we saw earlier, a proof that there is no largest prime starts, 'Suppose there were a largest prime; call it *n*.' The proof then proceeds as though '*n*' were a name (i.e., named something), but it does *not* name anything, as the proof itself shows. What we have here is a dummy name, or what Geach calls a quasi-name.

> [W]e could have quasi-names corresponding to merely *intentional* identity. Suppose we hear of a man who dreams of the same girl night after night […]; for convenience of conversation, we may say he dreams of Petronella every night, without either committing ourselves to the view that there is a real live girl he is dreaming of, or meaning that the name 'Petronella' is the name he gives the girl in his dreams. For us, 'Petronella' is then functioning not as a name but as a quasi-name. Geach (1969), p. 161

> This possibility of introducing quasi-names may seem to give some sort of reality to the dream girl, Petronella […]: as though what I have called quasi-names were

> really names, only of second-grade entities. But no such thing need be assumed. These quasi-names were introduced only in object position after intentional verbs; the use of them as logical subjects of predication neither is explained, nor could be justified, by this introduction. (A word intended to be used as a proper name is simply a different word from a word intended to be used as a quasi-name, even if they look and sound the same […].) Geach (1969), p. 162

Our dummy names are confined to a single branch, so we do not run afoul of the injunction against cross-derivation importation.

Actualistically acceptable (AA) models for propositional and for quantified modal logic are provided in the appendix. They resemble Kripke models, in that an AA model for QML will contain a set of worlds, only one of which is actual but each of which contains objects both actual and non-actual. Objects that are actual will be designated by names in the object language; those that are non-actual will be designated by constants that are in fact quasi-names. However, the distinction will come into play only rarely.

4. The relationship between AA models and Kripke models

AA models are very similar to Kripke models: both contain sets of worlds, a relation on that set, and in the quantified case a domain and a function which assigns a subset of that domain to each world. Truth at a world is defined in similar (though not identical) ways. Doesn't this mean that an AA model is simply a type of Kripke model, teeming with extra worlds, where the accessibility relation *R* (and not this new-fangled 'stipulation relation *S*') is arboreal, irreflexive and intransitive, and non-actual objects are confined to a single branch?

The answer is that an AA model should not be thought of as a special type of Kripke model, but as an altogether different beast, for both philosophical and formal reasons. The philosophical reason is that traditional models cannot fully honour our commitment to actualism. The formal reason is apparent if we consider the relationship between Kripke models and certain modal logics. For example, the modal logic S4 (also known as KT4) is sound and complete for Kripke models whose accessibility relations are reflexive and transitive. However, we cannot say that S4 is sound and complete for AA models whose stipulation relations are reflexive and transitive, because the stipulation relation is irreflexive and intransitive by definition.

The loss of soundness and completeness results would be a crippling blow for AA semantics. Happily, we need not bear such a burden. S4 *is* sound and complete for a certain class of AA models: not the empty class of AA models with reflexive and transitive stipulation relations, but for a class of AA models where a close relative of the stipulation relation — quasi-stipulation — is reflexive and transitive. The same goes for all our favourite (propositional) modal logics.

Consider the logic T, axiomatized by adding **T** to the basic modal logic K:

(**T**) $\Box\phi\rightarrow\phi$

T is valid in all Kripke models with reflexive accessibility relations. What would it take for **T** to be valid in an AA model? It can't be reflexivity of the stipulation relation *S*. What we want is not for every world to stipulate itself, which is impossible, but for every world to stipulate a copy of itself, identical except in being one order higher up.

The notion of a 'copy' needs to be formalized. Suppose we have an AA model for propositional logic, $<\alpha, W, S, V>$, where *W* is a set of worlds, α is a member of *W* (the actual world), *S* is a relation on *W* (stipulation), and *V* is a valuation that assigns either **T** or **F** to each propositional variable at each world. *V* is extended to all compound sentences in the usual way. For any two worlds *w* and *w′* in *W*, I shall say that *w* and *w′* **agree** (**w≈w′**) iff there is a bijection *f* that maps *w* onto *w′* and every descendant of *w* onto a descendant of *w′*, so that for all worlds w_1 within the domain of *f*,

(i) $w_1 S w_2$ iff $f(w_1) S f(w_2)$, and

(ii) $V(\phi, w_1) = V(\phi, f(w_1))$ for every sentence ϕ.

f is one-to-one, so each of w and its descendants gets assigned its own unique world from among w' and its descendants. f is onto, so there are no descendants of w' left over. Exactly the same sentences must be true in a world and its image. If w and w' agree, they and their descendants must be structurally (1) and qualitatively (2) identical. Agreement is obviously an equivalence relation.

We can now return to the question of what type of AA model validates the logic T. The axiom **T** will be valid in all AA models where each world w stipulates a world w' such that $w \approx w'$. (For a proof, see §3 of the appendix.)

The logic B (KTB) includes all the axioms and rules of T, plus the following:

(**B**) $\phi \rightarrow \Box\Diamond\phi$

B is valid in all Kripke models with (reflexive and) symmetric accessibility relations. When we move to AA models, again we are faced with the constraint that the stipulation relation is antisymmetric. In order for **B** to hold in an AA model, each world must stipulate not its immediate ancestor, but a copy of its ancestor, i.e., a world which agrees with its ancestor.

The logic S4 (KT4) includes all the axioms and rules of T, plus the following:

(**4**) $\Box\phi \rightarrow \Box\Box\phi$

4 is valid in all Kripke models with (reflexive and) transitive accessibility relations. Again, in AA models we may not require that the stipulation relation be transitive. In order for **4** to hold in an AA model, each world must stipulate copies of each of its descendants. (If $w_1 S w_2$ and $w_2 S w_3$, then w_1 stipulates a world w_3' which agrees with w_3.)

There is a neater way to describe these various constraints on AA models. Let us define a relation *R* on *W*:

(R) wRw' iff there is a w'' such that wSw'' and $w' \approx w''$.

R lets us ignore differences between worlds that agree. Even if *w* and *w′* are not *S*-related, they can be *R*-related as long as *w* is *S*-related to a *w″* that is sufficiently similar to (agrees with) *w′*. I shall refer to this defined relation *R* as **quasi-stipulation**: *w* quasi-stipulates *w′* iff *w* stipulates a world that agrees with *w′*. Trivially, *w* quasi-stipulates *w′* (wRw') if *w* stipulates *w′* (wSw'), but the converse does not hold.

The quasi-stipulation relation behaves just like the traditional Kripkean accessibility relation (with which it shares the letter *R*): the **T** axiom is valid in an AA model iff *R* is reflexive; **B** is valid iff *R* is symmetric; **4** is valid iff *R* is transitive; and so on. (Soundness and completeness results are provided in the appendix.)

A final note about agreement: with quantification, the notion of agreement becomes more complicated. We need to add some bells and whistles to take account of the fact that Kripke models, unlike AA models, do not necessarily follow the constraint (D) that non-actual objects be confined to a single branch.

Suppose we have an AA model containing two worlds *w* and *w′* that are not *S*-related or S^*-related (where S^* is the ancestral of *S*). Each contains exactly three objects: two actual objects, a_1 and a_2, and one non-actual object (i.e., an object not present in $D(\alpha)$): *z* in *w* and *z′* in *w′*. Everything that is true of *z* in *w* is true of *z′* in *w′*, and the two worlds agree about a_1 and a_2. The same goes for all the descendants of *w* and *w′*: they are structurally identical, and qualitatively identical with the exception that where *z* appears in the descendants of *w*, *z′* appears in the descendants of w'. In this case, we want to say that w agrees with w', because their only disagreement is over a non-actual object, a dummy name.

Take an AA model for quantified modal logic, $<\alpha, W, S, D, Q, V>$, where W is a set of worlds, α is a member of W, S is a relation on W, D is the domain, Q is a function assigning a subset $D(w)\subseteq D$ to each world in W, and V is a valuation that assigns extensions to each name and each predicate (including E) at each world. Where w and w' are any worlds in W, w agrees with w' ($w\approx w'$) iff there is a bijection f that maps w onto w', every descendant of w onto a descendant of w', and every object in the domain of w or of a descendant of w onto an object in the domain of w' or of a descendant of w', such that

(i) $f(a)=a$ for every actual object a in the domain of w or of a descendant of w;

(ii) w_1Sw_2 iff $f(w_1)Sf(w_2)$; and

(iii) $<a_1, \ldots, a_n>\in V(P^n, w^*)$ iff $<f(a_1), \ldots, f(a_n)>\in V(P^n, f(w^*))$, for every n-place predicate P^n and every world w^* which is w or a descendant of w.

It is a simple matter to prove, by induction on formula complexity, that even where ϕ is non-atomic, it is true at w iff a suitably altered sentence is true at $f(w)$. (The alterations involve the identities of non-actual objects.) These conditions ensure that the only disagreements between w and its descendants on the one hand, and w' and its descendants on the other, are over the identity of non-actual objects. They cannot differ in the number and placing of descendants, in the properties ascribed to actual objects in any of the worlds, or in the number or qualitative properties of non-actual objects.

AA models are faithful to the spirit of Kripke models, while grounded in a more robust actualism. It is not to be expected that they will replace their admittedly simpler predecessors.

Rather, they illustrate that only relatively minor changes need be made to put our semantics on a more secure footing.[16]

Appendix: AA (actualistically acceptable) models

1. AA models for propositional modal logic

I shall now define an AA model for propositional modal logic, following as closely as possible the definition of a traditional Kripke model (Kripke 1963).

An **AA model structure**, or **AA frame**, is an ordered triple, $<\alpha, W, S>$, where W is a set, α is in W, and S is an arboreal[17], irreflexive and intransitive relation on W, with α being the point of origin. α, W and S are intuitively thought of as the actual world, a set of possible worlds, and the stipulation relation.

An **AA model** consists of an AA frame with a **valuation** $V(P, w)$, where P varies over propositional variables, w varies over elements of W, and the range of V is the set $\{\mathbf{T}, \mathbf{F}\}$. V specifies for each atomic formula (propositional variable) and each world whether the formula is true at that world. We then extend V recursively to non-atomic formulas in the usual way:

$V(\neg\phi, w) = \mathbf{T}$ iff $V(\phi, w) = \mathbf{F}$;

$V(\phi\rightarrow\psi, w) = \mathbf{T}$ iff $V(\phi, w) = \mathbf{F}$ or $V(\psi, w) = \mathbf{T}$; and

$V(\Box\phi, w) = \mathbf{T}$ iff for all $w' \in W$ such that wSw', $V(\phi, w') = \mathbf{T}$.

(The connectives '$\wedge$', '$\vee$' and '$\leftrightarrow$' are defined in terms of '$\neg$' and '$\rightarrow$'; the modal operator '$\Diamond$' is defined in terms of '$\neg$' and '$\Box$'.)

2. AA models for quantified modal logic

One's semantics for quantified modal logic will depend upon one's choice of logic. For a variety of reasons rehearsed by James W. Garson (1984), my choice of logic is what Garson calls Q1R.

Q1R is a free logic with world-relative domains and rigid terms, the axioms and rules of which can be found in Garson (1984, pp. 252, 255). Q1R is sound and complete for a class of Kripke models called Q1R-models. For comparative purposes, I shall describe these models, then tweak them to give a fully actualistically acceptable semantics.

A **Q1R-model** is an ordered quintuple $<W, R, D, Q, V>$, where:

- *W* is a set of worlds;
- *R* is a binary relation on *W*, the traditional accessibility relation;
- *D* is the domain;
- *Q* is a function that assigns to each world $w \in W$ a subset $D(w) \subseteq D$ (intuitively, the domain of quantification of *w*);
- *V* is a valuation that assigns, for each world $w \in W$,
  - an object in *D* to each term *t* (including variables),
  - a set of ordered *n*-tuples of elements of *D* to each *n*-ary predicate, and
  - the set $D(w)$ to the predicate constant *E*.

Additionally, *V* must satisfy the rigidity condition:

(VRT) $V(t, w) = V(t, w')$ for all w, w' in *W*.

V provides extensions for each term and each predicate (including *E*) at each world, so that the extension of each term is the same member of *D* at every world.

We then extend *V* to assign one of the values {**T**, **F**} to each atomic formula at each world:

$$V(P^n t_1 t_2 \ldots t_n, w) = \mathbf{T} \quad \text{iff} \quad <V(t_1, w), V(t_2, w), \ldots, V(t_n, w)> \in V(P^n, w);$$

$$V(Et, w) = \mathbf{T} \quad \text{iff} \quad V(t, w) \in D(w); \text{ and}$$

$V(t{=}t', w) = \mathbf{T}$ iff $V(t, w) = V(t', w)$.

V is defined for the compound formulas $\neg\phi$, $\phi\rightarrow\psi$ and $\Box\phi$ just as in propositional modal logic. For the universal quantifier, we have the following:

$V(\forall x\phi, w) = \mathbf{T}$ iff for every $d\in D(w)$, $V(d/x)(\phi, w) = \mathbf{T}$,

where $V(d/x)$ is the valuation just like V, except that $V(x)=d$.

We shall now proceed to modify these models to be fully acceptable to the actualist. An AA model for the logic Q1R is not an ordered quintuple $<W, R, D, Q, V>$ as above, but an ordered sextuple $<\alpha, W, S, D, Q, V>$. W, D and Q are exactly the same as in a Q1R-model. The two new elements are:

(i) α is an element of W (the actual world); and

(ii) S is the stipulation relation: an arboreal, irreflexive and intransitive relation on W; with α being the point of origin.

(Both α and S appeared in the specification of an AA model for propositional modal logic.)

V is slightly different in an AA model than it was in a Q1R-model. In a Q1R-model, V had to satisfy the rigidity condition:

(VRT) $V(t, w) = V(t, w')$ for all w, w' in W.

In an AA model, V must satisfy both the rigidity condition and the following, equivalent to both (D) and (D′) in the discussion of the restrictions on S, which I shall call the ‘single branch’ condition:

(SB) If a non-actual object z appears for the first time in a world w (i.e., is not present in any ancestors of w, including α), then z can appear only in w or its descendants.

V is then extended to assign truth-values to sentences in the usual way.

3. Soundness

AA models can be shown to be sound for a wide variety of popular modal logics, such as T, B, S4 and S5. In the interests of brevity, I shall not reproduce a proof for the soundness of the underlying logic Q1R, which proceeds exactly it would under traditional semantics (because none of the other rules or axioms involve iterated modality).

The logic T is sound for Kripke models whose accessibility relations are reflexive. We want now to show that T is sound for AA models whose quasi-stipulation relations are reflexive. It suffices to show that the axiom **T** ($\Box\phi\rightarrow\phi$) holds in all such AA models. Suppose **T** fails in an AA model: there is a world *w* at which **T** is false. At *w*, $\Box\phi$ is true but ϕ is false. Since $\Box\phi$ is true at *w*, ϕ must be true at any world which is stipulated by *w*. Thus there does not exist a world w' such that wSw' and $w\approx w'$. By the definition of *R*, *w* is not *R*-related to itself; *R* is not reflexive.

Similarly, we can show that the logic B, which adds the axiom **B** ($\phi\rightarrow\Box\Diamond\phi$) to T, is sound not just for Kripke models with symmetric accessibility relations, but AA models with symmetric quasi-stipulation relations. Suppose **B** fails in an AA model: there is a world w_1 at which it is false. At w_1, ϕ is true and $\Box\Diamond\phi$ is false; hence $\Diamond\Box\neg\phi$ is true. By the truth condition for $\Diamond$, w_1 stipulates a world w_2 at which $\Box\neg\phi$ is true. At any world stipulated by w_2, $\neg\phi$ must be true. By the definition of agreement, there does not exist a world stipulated by w_2 which agrees with w_1. Thus w_2 is not *R*-related to w_1. However, w_1 is *S*-related, hence *R*-related, to w_2. *R* is not symmetric.

Similar proofs can be given for S4 and S5. A modal logic is sound for the class of AA models with a particular kind of quasi-stipulation relation iff it is sound for the class of Kripke models with the same kind of accessibility relation.

4. Completeness: unravelling Kripke models into AA models

If a sentence ϕ is not a theorem of a particular modal logic (e.g., S4), then there is a Kripke model of a certain kind (for S4, reflexive and transitive) in which ϕ fails. Using a technique known as unravelling,[18] we can convert this Kripke model into an AA model whose quasi-stipulation relation will be have the same features. Thus S4 is complete for AA models with reflexive and transitive quasi-stipulation relations; B is complete for AA models with reflexive and symmetric quasi-stipulation relations; and so on.

For propositional modal logic, unravelling a Kripke model $<\alpha, W, R, V>$ into a model $<\alpha', W', R', V'>$ proceeds as follows.

(1) Deriving α' and W′. If α is the privileged element of W in the original Kripke model, the sequence $<\alpha>$ is the privileged element α' of W' in the unravelled model. If the sequence $<\alpha, w_1, \ldots, w_n>$ is an element of W', then the sequence $<\alpha, w_1, \ldots, w_n, w_{n+1}>$ is an element of W' iff $w_n R w_{n+1}$ in the original Kripke model. There are no other elements of W'. In other words, the elements of the unravelled model are all and only those finite sequences of elements of the original model where each member of the sequence is R-related to the next in the original model.

(2) Deriving R′. If $\underline{\vec{w}}_i$ and $\underline{\vec{w}}_j$ are elements of W' (and hence are sequences of worlds of the original model), then $\underline{\vec{w}}_i R' \underline{\vec{w}}_j$ iff the sequence $\vec{w}_j$ extends the sequence $\underline{\vec{w}}_i$ by one element.

(3) Deriving V′. Suppose α, w_1, …, w_n are elements of W, and <α, w_1, …, w_n> is an element of W': Where P is an atomic sentence, let $V'(P, <\alpha, w_1, \ldots, w_n>) = V(P, w_n)$. An atomic sentence is true at a world (sequence) in the new model iff it is true in the old model at the last world in the sequence.

The unravelled model, <α', W', R', V'>, will in fact serve as our AA model. We shall rename the stipulation relation S, to standardize nomenclature. We need to verify that <α', W', S, V'> is an AA model, by verifying the properties of S and V'.

Is S arboreal, irreflexive and intransitive? Clearly, it is generated; every sequence begins with α, so every element of W' is a descendant of <α>, as required. It is antisymmetric: if one sequence is an extension of another, the second cannot also be an extension of the first. It is anticonvergent: if $\underline{\vec{w}}_1 S \underline{\vec{w}}_2$, $\underline{\vec{w}}_1 S \underline{\vec{w}}_3$, and $\underline{\vec{w}}_2 \neq \underline{\vec{w}}_3$, then $\underline{\vec{w}}_2$ and $\underline{\vec{w}}_3$ both extend $\underline{\vec{w}}_1$ by one element, but their last elements will be different. Hence there cannot be a $\underline{\vec{w}}_4$ which is an extension of both $\underline{\vec{w}}_2$ and $\underline{\vec{w}}_3$. Thus S is arboreal. It is also irreflexive: a finite sequence cannot be an extension of itself. Finally, S is intransitive: if $\underline{\vec{w}}_1 S \underline{\vec{w}}_2$ and $\underline{\vec{w}}_2 S \underline{\vec{w}}_3$, then $\underline{\vec{w}}_2$ extends $\underline{\vec{w}}_1$ by one element and $\underline{\vec{w}}_3$ extends $\underline{\vec{w}}_2$ by one element. We cannot have $\underline{\vec{w}}_1 S \underline{\vec{w}}_3$, because $\underline{\vec{w}}_3$ extends $\underline{\vec{w}}_1$ by two elements rather than one. S meets all the requirements for a stipulation relation in an AA model.

Does V' provide adequate truth-conditions? V' was defined so that an original world w_n and a new world <α, w_1, …, w_n> (henceforth abbreviated $\underline{\vec{w}}$) must agree on atomic sentences; but we want them to agree on all sentences, not just atomic ones. We can show this by induction on formula complexity. Suppose w_n and $\underline{\vec{w}}$ agree on a formula ϕ: $V(\phi, w_n) = V'(\phi, \underline{\vec{w}})$. Since the value at a world of the formula $\neg\phi$ is the reverse of the value for ϕ, $V(\neg\phi, w_n) = V'(\neg\phi, \vec{w})$.

Clearly, w_n and $\underline{\vec{w}}$ will also agree on formulas of the form $\phi\rightarrow\psi$. Suppose that a sentence of the form $\Box\phi$ is true at w_n. This means that for all $w'\in W$ such that w_nRw', ϕ is true at w'. By the construction of the new model, for every one of these worlds w' in the original model, there is a sequence $<\alpha, w_1, \ldots, w_n, w'>\in W'$ such that $<\alpha, w_1, \ldots, w_n>S<\alpha, w_1, \ldots, w_n, w'>$. Each of these new worlds $<\underline{\vec{w}}, w'>$ will agree with the old worlds w' about ϕ: for every old w' such that w_nRw', and every new $<\underline{\vec{w}}, w'>$ such that $\underline{\vec{w}}S<\underline{\vec{w}}, w'>$, $V(\phi, w') = V'(\phi, <\underline{\vec{w}}, w'>) = \mathbf{T}$. Since ϕ is true at every new world S-related to $\underline{\vec{w}}$, $\Box\phi$ must be true at $\underline{\vec{w}}$, as at w_n. Hence $\underline{\vec{w}}$ and w_n agree on all sentences.

Thus if a sentence fails in a Kripke model, it fails in the unravelled model, and hence in an AA model. This shows that the basic propositional modal logic (K) is complete for AA models.[19]

Suppose the propositional modal logic is not K, but T. If a formula ϕ is a non-theorem of T, then ϕ fails in a Kripke model whose accessibility relation R is reflexive. What happens when we unravel this Kripke model? We do of course get an AA model, with an irreflexive stipulation relation S. The crucial desideratum is not that S be reflexive, but that *quasi*-stipulation be reflexive.

We can secure this. Every world w in the original model is R-related to itself. Therefore, for every sequence $<\alpha, \ldots, w>$ in W', there is a sequence $<\alpha, \ldots, w, w>\in W$ such that the former is S-related to the latter. Call these sequences $\underline{\vec{w}}_1$ and $\underline{\vec{w}}_2$. Since both sequences end with w, both agree with w on all sentences. Thus they agree with each other: $\underline{\vec{w}}_1\approx\underline{\vec{w}}_2$. Recall the definition of quasi-stipulation: w_1 quasi-stipulates w_2 iff w_1 stipulates a world that agrees with w_2. $\underline{\vec{w}}_1$

stipulates $\underline{\vec{w}}_2$, and $\underline{\vec{w}}_2$ agrees with $\underline{\vec{w}}_1$, so $\underline{\vec{w}}_1$ quasi-stipulates itself. This holds for every sequence in W'; quasi-stipulation is reflexive.

The completeness proofs for propositional B, S4 and S5 are similar.

For the quantified modal logic Q1R, converting a Q1R-model into a quantified AA model is a little more involved than the conversion for propositional modal logic. The reason, of course, is that with quantification there are additional conditions that a model must meet in order to count as an AA-model: in particular, the condition that non-actual objects must be confined to a single branch. After a Q1R-model is unravelled, it must be tweaked a little more to get it into the right shape for an AA model.

First, we shall unravel a Q1R-model to produce a model $\langle\alpha', W', R', D, Q', V'\rangle$. As before, the elements of W' are finite sequences of elements of W, and R' is the one-element extension relation. The domain, D, remains the same after unravelling. The domain of each new world should reflect the domain of the relevant old world: $Q'(\langle\alpha, \ldots, w_n\rangle) = Q(w_n)$. The new valuation, V', should generally agree with V. If V assigns an object $d \in D$ to a term t at a world w_n, V' should assign d to t at the new world $\langle\alpha, \ldots, w_n\rangle$. If V assigns a set of ordered n-tuples of elements of D to the n-ary predicate P at w_n, then V' should assign that same set of ordered n-tuples to P at $\langle\alpha, \ldots, w_n\rangle$. V assigns the set $D(w_n)$ to the predicate constant E at w_n; V' should assign the set $D(\langle\alpha, \ldots, w_n\rangle) = D(w_n)$ to the constant E at $\langle\alpha, \ldots, w_n\rangle$. As in the propositional case, it can be shown by induction on formula complexity that V' and V will then agree on all sentences.

We now have the unravelled model <α', W', R', D, Q', V'>. Of course, there is absolutely no guarantee that this model is actualistically acceptable, because it may violate the single branch condition:

(SB) If a non-actual object z appears for the first time in a world w (i.e., is not present in any ancestors of w, including α), then z can appear only in w or its descendants.

Suppose that the original model, M, is such that an object d is not in the domain of α, but is in the domain of two distinct worlds (w_1 and w_2) accessible from α. The unravelled model M' will violate (SB), as a non-actual object d will be present in both <α, w_1> and <α, w_2> cannot be allowed to stand, as it violates the single branch condition on AA-models. We must revise M' in order to accommodate (SB). First, we shall remove the offending non-actual object d and replace it with two non-actual objects, d_1 and d_2, the former to inhabit <α, w_1> and the latter to inhabit <α, w_2>. Next, we shall define a new valuation V'' to cover these new objects: <…, d_1, …>$\in V''(P^n, \vec{w})$ iff <…, d, …>$\in V'(P^n, \vec{w})$. In other words, exactly the same things are to be true of d_1 as of b. The same goes for d_2. V'' is otherwise identical to V'. We now have an AA-model, M''. If the original Q1R-model M is a model for the sentence ϕ, then the model M', derived from M by unravelling, is still a model for ϕ; and *as long as ϕ contains no constants*[20], the model M'', derived from M' by renaming the non-actual objects so as to be confined to a single branch, is an AA-model for ϕ.

It can be shown by induction on formula complexity that as long as ϕ contains no constants, M is a Q1R-model for ϕ iff M'' is an AA model for ϕ. Where R is the accessibility relation in M and S is the stipulation relation in M'', the quasi-stipulation relation S' will be reflexive if *R* is reflexive, symmetric if *R* is symmetric, and so on. Renaming non-actual objects will not matter, since agreement (≈) between worlds within AA models ignores differences in non-actual objects.

Notes

[11] The possibilist might object that of course the word ‘he’ fails to refer, because I have not given enough information to specify which of my many possible brothers I am talking about. This would have to be done by giving additional details about the brother (‘the older brother of mine that resulted from the union of ovum *O* and sperm *S*’, for example). But the multiplicity of brothers is not why the word ‘he’ fails to refer. On the actualist view, it is that *no* object, rather than too many objects, is possibly an older brother of mine. Hence the failure of reference cannot be fixed by providing more details about such a possible brother.

[2] The existential quantifier in (1), however, does *not* range over all possibilia, but over the domain of a single world. In fact, if the quantifier were interpreted as ranging over all possibilia, it would be true necessarily (rather than merely possibly) that there is an object *x* which is RH’s and JH’s older brother, as the contents of the totality of modal space are presumably necessary.

[3] My reasons are given in another paper, “Contingent objects and the Barcan Formula”. However, some actualists do accept the Barcan Formula (Linsky and Zalta 1994, 1996; Williamson 1998, 2000), while David Lewis, a possibilist, rejects it.

[4] There is also a technical wrinkle with this particular ‘dummy name’ response to McMichael: some objects may not get named. We can conceive of a world inhabited only by two qualitatively indistinguishable objects (Black 1952). We would then have no definite description with which to pick out either object. One could say that they have different modal properties (e.g., one could exist without the other), so that one and not the other might ‘pick up’ a dummy name in some other world. But then we are still faced with the illegal importation of such a world into the original two-object world.

[5] Plantinga’s unactualized essences (Plantinga 1976) are a good example of surreptitious importation. Although my older brother who is a banker but could have been a concert pianist is non-actual, the essence of this object is not. It is a perfectly respectable (although of course abstract) actual object, which may be quantified over, ascribed properties, and so on. In brief, my worry with this proposal is that it’s not clear whether *all* non-actual objects — rather than just the interesting ones, like non-actual people — have essences that can be specified in a non-circular way. Some interesting consequences of this worry are discussed in another paper, “The transience of possibility”.

[6] Open importation amounts to holding domains fixed across worlds, i.e., to accepting the Barcan Formula; see note 3 above.

[7] Considerations of space preclude a discussion here of the merits of a quantified modal logic involving trans-world identity over those of counterpart theory; my own views are given in my doctoral dissertation, *Actualism and Quantified Modal Logic* (Princeton, 2002), the main ideas of which this paper is an abstract.

[8] I discuss the logic of nested fictions in another paper, “Fictions within fictions”.

[9] A. P. Hazen (1996) provides a semantics that gives somewhat similar results. He starts by stripping the modal language of all *de re* modalities, leaving only *de dicto* modalities, which are to be represented by ‘purely qualitative worlds’. He constructs models called ‘world-fans’ which, like my nested worlds, form a tree structure (hence the name). The referent of an individual constant is determined by a growing cluster of counterpart functions that are defined separately along each branch. The most important result (for actualists) is that iterated modalities can be accounted for but individual constants cannot represent non-actual objects.

Although Hazen's and my accounts are superficially similar, there are three advantages to my programme. First, *de re* modalities can be preserved in (indeed, are crucial to) my account; this, I think, provides a sturdier philosophical foundation than Hazen offers. Second, Hazen's account is counterpart-theoretic. It is true that constraints could be placed on the myriad counterpart functions that would make them behave more or less as trans-world identity would. However, such a move would, again, be metaphysically unmotivated. Third, Hazen explicitly rejects the view that possible worlds can differ in kind (i.e., by being of different orders), endorsing instead what he calls 'semantic uniformity' (the modal operators must range over the same type of object). Although semantic uniformity has great initial attraction, I believe the parallel construction of worlds-within-worlds and fictions-within-fictions is enough to suggest that uniformity is not the way to go.

[10] From the dictum already quoted, that "'possible worlds' are *stipulated*, not *discovered* through powerful telescopes" (Kripke 1980, p. 44).

[11] David Lewis (1986) regards this as a strong point in favour of his version of possibilism, which does not trade one modality for another but takes spatio-temporal systems not spatio-temporally related to ours as primitive. However, it has been suggested that this very feature demonstrates that full-blown Lewisian 'possible worlds' talk is not really modal talk at all.

[12] I argue in "The transience of possibility" that this is in fact not the case, and hence that B should be rejected.

[13] An example of external criteria being applied successfully to a case of 'intentional identity' would be in analysing the sentence, "Hannibal and Hasdrubal worshipped the same god [Baal]." (This example comes from Geach 1969.) We don't want to say that there exists an object *x* such

that *x* is a god and Hannibal worshipped *x* and Hasdrubal worshipped *x*. What Hannibal and Hasdrubal *do* share is a religion. The religion of Baal-worship, unlike Baal, is ontologically respectable. We can purge the original sentence, "Hannibal and Hasdrubal both worshipped the same god", of any apparent ontological commitment to nonexistent entities by rephrasing it as "Hannibal and Hasdrubal were co-religionists." In order to incorporate specific mention of Baal (as opposed to some other non-existent god), we can tell a story about the historical origins of Baal-worship and say that Hannibal and Hasdrubal were participating in this tradition rather than another, say Jupiter-worship.

Along these lines, we could say that Hob's belief and Nob's wondering are 'about the same witch' if they are both derived from the same witch-myth. Certain historical facts (about the origin and perpetuation of the myth) would have to hold for this analysis to work.

[14] Although not universally liked. Nathan Salmon (1989) gives a sorites argument against the transitivity of accessibility.

[15] By 'branch' I do not mean an (infinitely long) path just one world wide, but a world together with all its descendants. A non-actual object's branch starts at the first world where it appears (the unique world which contains the object and whose ancestors do not). Of course, objects need not appear in *every* world in a branch.

[16] I am indebted to John Burgess, Gideon Rosen and Tim Williamson for comments on earlier drafts of this work.

[17] 'Arboreal' is shorthand for 'generated, antisymmetric and anticonvergent' (i.e., forming a tree).

[18] The pedigree of this technique is briefly described in section 13 of Bull and Segerberg (1984). Unravelling first appeared in Dummett and Lemmon (1959). Henrik Sahlqvist uses the technique, without citing Dummett and Lemmon, in Sahlqvist (1975), to “turn generated models into equivalent generated models” (Bull and Segerberg 1984, p. 125). The chief result is that, since we can turn a model which isn’t irreflexive into an equivalent irreflexive model, there is no axiom that corresponds to irreflexivity. The same goes for asymmetry, as well as intransitivity.

[19] It may be objected that there could be worlds in the original Kripke model which don’t get ‘reproduced’ in the unravelled model: namely, those worlds not accessible from α in a finite number of steps. If a non-theorem ϕ fails at one of these worlds, then the unravelled model may no longer be a countermodel for ϕ: there may be no world in the new model at which ϕ is false. This is an important objection, but it can be overcome. If S is any complete normal modal system, then S is complete for any class of frames for S which contains all the generated frames for S (see Hughes and Cresswell 1984). What this means is that if ϕ is not a theorem of a logic S, there is a *generated* countermodel for ϕ (which may also be reflexive, transitive, etc., depending on *S*). Unravelling a generated Kripke model preserves all the worlds in the original model, thus ensuring that the new (AA) model is also a countermodel for ϕ.

[20] A discussion of the complications that can arise when constants are included is provided in my dissertation.

Bibliography

Black, M. (1952): "The identity of indiscernibles", *Mind* 61(242), pp. 153-164.

Bull, R. and K. Segerberg (1984): "Basic modal logic", in D. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, Dordrecht: Kluwer Academic Publishers, ii, pp. 1-88.

Dummett, M. A. E. and E. J. Lemmon (1959): "Modal logics between S4 and S5", *Zeitschrift für mathematische Logik und Grundlagen der Mathematik* 5, pp. 250-264.

Garson, J. W. (1984): "Quantification in modal logic", in D. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, Dordrecht: Kluwer Academic Publishers, ii, pp. 249-308.

Geach, P. T. (1967): "Intentional identity", *Journal of Philosophy* 64(20), pp. 627-32.

Geach, P. T. (1969): "The perils of Pauline", *Review of Metaphysics* 23, pp. 287-300.

Hazen, A. P. (1996): "Actualism again", *Philosophical Studies* 84, pp. 155-81.

Hughes, G. E. and M. J. Cresswell (1984): *A Companion to Modal Logic*, London: Methuen.

Kripke, S. (1963): "Semantical considerations on modal logic", *Acta Philosophica Fennica* 16, pp. 83-94.

Kripke, S. (1980): *Naming and Necessity*, Cambridge, Mass.: Harvard University Press.

Lewis, D. (1968): "Counterpart theory and quantified modal logic", *Journal of Philosophy* 65, pp. 113-126.

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

Linsky, B. and E. Zalta (1994): "In defense of the simplest quantified modal logic", in *Philosophical Perspectives, vol. 8: Logic and Language*, Atascadero: Ridgeview, pp. 431-58.

Linsky, B. and E. Zalta (1996): "In defense of the contingently nonconcrete", *Philosophical Studies* 84, pp. 283-94.

Marcus, R. B. (1946): "A functional calculus of first order based on strict implication", *Journal of Symbolic Logic* 11, pp. 1-16.

McMichael, A. (1983): "A problem for actualism about possible worlds", *Philosophical Review* 92, pp. 49-66.

Plantinga, A. (1974): *The Nature of Necessity*, Oxford: Oxford University Press.

Plantinga, A. (1976): “Actualism and possible worlds”, *Theoria* 42, pp. 139-60.

Quine, W. V. (1956): “Quantifiers and propositional attitudes”, *Journal of Philosophy* 53, pp. 177-87.

Sahlqvist, H. (1975): “Completeness and correspondence in the first and second order semantics for modal logic”, in S. Kanger (ed.), *Proceedings of the Third Scandinavian Logic Symposium*, Amsterdam: North-Holland, pp. 110-43.

Salmon, N. (1989): “The logic of what might have been”, *Philosophical Review* 98, pp. 3-34.

Williamson, T. (1998): “Bare possibilia”, *Erkenntnis* 48(2 and 3), pp. 257-273.

Williamson, T. (2000): “Existence and contingency”, *Proceedings of the Aristotelian Society* 100(1), pp. 117-39.