

An Actualist's Guide to Quantifying In

Agustín Rayo
Massachusetts Institute of Technology

July 30, 2012

Abstract: I develop a device for simulating quantification over merely possible objects from the perspective of a *modal actualist*—someone who thinks that everything that exists actually exists.

Keywords: Intentional Identity, Empty Names, Possible Worlds, Merely Possible Objects, Modal Actualism.

1 Introduction

Philosophers call on possible worlds to perform different kinds of jobs. One of these jobs is that of explaining what the truth of modal truths and the falsity of modal falsities *consists in*. “What does the possibility that p consist in?”, you ask. “It consists in the existence of a possible world at which p ”, someone might reply.

Possible worlds are also used as *semantic machinery*. The semanticist needs entities for her quantifiers to range over, and possible worlds—or, more generally, *possibilia*—can be used to construct them. One might take a proposition to be a set of possible worlds (structured or not); one might take the intension of a predicate to be a function that assigns each world a set of objects in that world (or some complication thereof); one might use possible worlds to characterize the semantic values of modal operators; and so forth.

(It is sometimes assumed that these two jobs are best performed by a single notion of possible world, as in Lewis (1986). But it is worth keeping in mind that such an assumption is not mandated. It would be in principle possible to develop two distinct notions of possible world, one for the foundationalist project and another for the semantic project.)

The main objective of this paper is to defend the claim that a specialized modal ontology is not needed *as semantic machinery*. (The further claim that it is not needed to explain what modal truth consists in is defended in chapters 2 and 5 of Rayo (forthcoming).) Some philosophers might be interested in this sort of project because of ontological scruples. Not me: my own view is that ontology is cheap. I am interested in the project because I think it helps clarify the role of possible worlds in our philosophical theorizing.

My proposal is an instance of what David Lewis called ‘ersatzism’. I will argue that the needs of the semanticist can be fulfilled by using *representatives* for possibilities in place of possibilities. Although there are other ersatzist proposals in the literature,¹ I hope the machinery developed here will earn its keep by delivering an attractive combination of frugality and strength.

The proposal is frugal in two different respects. First, it is metaphysically frugal: it is designed to be acceptable to modal actualists, and presupposes very little by way of ontology. (I help myself to set-theory, but do not assume a specialized modal ontology, or an ontology of properties.) Second, the proposal is ideologically frugal: it does not presuppose potentially controversial expressive resources such as infinitary languages or non-standard modal operators. The point of developing machinery that presupposes so little is that philosophers of different persuasions can put it to work without having to take a stance on difficult philosophical issues.

As far as strength is concerned, one gets a qualified version of the following claim: anything that can be said by quantifying over Lewisian possibilities can also be said by using the machinery developed here. The result is that the proposal can be used quite

freely in the context of semantic theorizing, without having to worry too much about running into expressive limitations. An especially useful feature of the proposal is that it allows one to enjoy the benefits of quantification over *sets* of possibilia, which are often appealed to in the course of semantic theorizing.²

Possible-words-theorists sometimes claim that the same individual exists according to distinct possible worlds. (There is a world according to which I have a sister who is a philosopher, and a world at which that very individual is a cellist rather than a philosopher.) Ersatzist representatives for such worlds might be said to be *linked*. Much of the paper will be devoted to the phenomenon of linking. An effective treatment of linking is crucial to the success of the ersazist program.

2 Possibility

Sadly, I don't have a sister. But I might have had a sister. In fact, I might have had a sister who was a philosopher. And, of course, had I had a sister who was a philosopher, she wouldn't have been a philosopher essentially: she might have been a cellist instead. The following is therefore true:³

SISTER

$\diamond(\exists x(\text{Sister}(x, \text{AR}) \wedge \text{Phil}(x) \wedge \diamond(\text{Cellist}(x) \wedge \neg\text{Phil}(x))))$

(*Read:* I might have had a sister who was a philosopher and might have been a cellist rather than a philosopher.)

On the most straightforward possible-worlds semantics for first-order modal languages, SISTER will only be counted as true if there are worlds w_1 and w_2 with the following properties: according to w_1 , there is an individual who is my sister and a philosopher; according to w_2 , that very individual—as one is inclined to put it—is a cellist rather than

a philosopher. It is therefore tempting to say the following:

$$\exists x([\text{Sister}(x, \text{AR}) \wedge \text{Phil}(x)]_{w_1} \wedge [\text{Cellist}(x) \wedge \neg \text{Phil}(x)]_{w_2})$$

(*Read:* There is an x such that: (i) according to w_1 , x is my sister and a philosopher, and (ii) according to w_2 , x is a cellist rather than a philosopher.)

But is there anything to make this existential quantification true? A merely possible sister could do the job. But modal actualists believe there is no such thing. The job could also be done by an actually existing object who is not my sister but might have been my sister. (See Williamson (2010).) But the claim that such objects exist would be a substantial metaphysical assumption—an assumption that it would best to avoid, if at all possible.

2.1 A Kripke-semantics for Actualists

There is a certain sense in which it is straightforward for an actualist to give an adequate Kripke-semantics for modal sentences. The trick is to have one's semantics quantify over *representations* of possibilities, rather than over the possibilities themselves. In this section I will describe one such semantics. It is an elaboration of an idea introduced in Roy (1995) and further developed in Melia (2001).

Let L be a first-order language, and let L^\diamond be the result of enriching L with the sentential operator ' \diamond '. An a -world (short for 'actualist-world') for L^\diamond is an ordered pair $\langle D, I \rangle$, consisting of a domain and an interpretation function, such that:

- The domain D is a set containing of two kinds of entities. First, it contains every every ordered pair of the form $\langle x, \text{'actual'} \rangle$, for x an individual in the domain of L . (These pairs will be used to represent actually existing objects.) Second, it may contain ordered pairs of the form $\langle x, \text{'nonactual'} \rangle$, where x is some *actually existing* individual, set-theoretic or otherwise. (These pairs will be used to represent merely

possible objects.)

- If C is an individual constant of L and x is its intended interpretation, the interpretation-function I assigns the pair $\langle x, \text{'actual'} \rangle$ to C . (Accordingly, if ‘SOCRATES’ is an individual constant referring to Socrates, I assigns the pair $\langle \text{Socrates}, \text{'actual'} \rangle$ to ‘SOCRATES’.)
- I assigns a subset of D^n to each n -place predicate-letter of L , and a function from D^n to D to each n -place function-letter of L .

The notions of truth and satisfaction at an a -world are characterized along standard lines, with the proviso that ‘ $x = x$ ’ is only satisfied at an a -world by objects in the domain of a -world, with the result that ‘ $x = x$ ’ can be used as an existence predicate. See Appendix A for details.

The best way to see how a -worlds are supposed to work is by considering a simple example. Take the a -world $\langle D_1, I_1 \rangle$ from figure 1. The domain of $\langle D_1, I_1 \rangle$ consists of the pairs $\langle \text{Agustín}, \text{'actual'} \rangle$ (which represents me) and $\langle \text{Socrates}, \text{'nonactual'} \rangle$ (which represents *not* Socrates, but a merely possible object), and I_1 assigns ‘Sister’ a set containing the pair $\langle \langle \text{Socrates}, \text{'nonactual'} \rangle, \langle \text{Agustín}, \text{'actual'} \rangle \rangle$. The result is that $\langle D_1, I_1 \rangle$ makes ‘ $\exists x(\text{Sister}(x, \text{AR}))$ ’ true, and therefore represents a possibility whereby I have a sister. (Why use the pair $\langle \text{Socrates}, \text{'nonactual'} \rangle$ to represent my merely possible sister rather than some other pair— $\langle \text{Plato}, \text{'nonactual'} \rangle$, say? No good reason. Just like one can write a story according to which I have a sister using any font one likes, so one can define an a -world according to which I have a sister using any pair $\langle z, \text{'nonactual'} \rangle$ one likes.)

It is useful to compare a -worlds to Lewisian worlds. Like a -worlds, Lewisian worlds can be thought of as representing possibilities. Here is Lewis:

How does a world, [Lewisian] or ersatz, represent, concerning Humphrey, that he exists?... A [Lewisian] world might do it by having Humphrey himself as

a part. That is how our own world represents, concerning Humphrey, that he exists. But for other worlds to represent in the same way that Humphrey exists, Humphrey would have to be a common part of many overlapping worlds... I reject such overlap... There is a better way for a [Lewisian] world to represent, concerning Humphrey, that he exists... it can have a Humphrey of its own, a flesh-and-blood counterpart of our Humphrey, a man very much like Humphrey in his origins, in his intrinsic character, or in his historical role. By having such a part, a world represents *de re*, concerning Humphrey—that is, the Humphrey of our world, whom we as his worldmates may call simply Humphrey—that he exists and does thus-and-so.⁴

(It is easy to lose track of Lewis’s representationalism because he also subscribed to a different claim: what it is for it to be possible that p is for there to be a Lewisian world at which (i.e. representing that) p . For further discussion of these issues, see section Rayo (forthcoming).)

Whereas Lewisian worlds represent by analogy, *a*-worlds represent by satisfaction. A Lewisian world represents the possibility that I have a sister by containing a person who is similar to me in certain respects, and has a sister. An *a*-world, on the other hand, represents the possibility that I have a sister by being such as to satisfy the formula ‘ $\exists x(\text{Sister}(x, \text{AR}))$ ’.

From the perspective of the Lewisian, an individual with a counterpart in the actual world represents its actual-world counterpart, and an individual with no counterpart in the actual world represents a merely possible object. From the present perspective, a pair of the form ‘ $\langle x, \text{‘actual’} \rangle$ ’ represents its first component, and a pair of the form ‘ $\langle x, \text{‘nonactual’} \rangle$ ’ represents a merely possible object (even though the pair itself, and both of its components, are actually existing objects).

Two representations are *linked* if—as one is inclined to put it—they concern the same individual, even if the individual in question doesn’t exist. In order for a Kripke

$\langle D_1, I_1 \rangle$

$$\begin{aligned} D_1 &= \{ \langle \text{Agustín, 'actual'} \rangle, \langle \text{Socrates, 'nonactual'} \rangle \} \\ I_1(\text{'Philosopher'}) &= \{ \langle \text{Socrates, 'nonactual'} \rangle \} \\ I_1(\text{'Cellist'}) &= \{ \} \\ I_1(\text{'Sister'}) &= \{ \langle \langle \text{Socrates, 'nonactual'} \rangle, \langle \text{Agustín, 'actual'} \rangle \rangle \} \\ I_1(\text{'AR'}) &= \langle \text{Agustín, 'actual'} \rangle \end{aligned}$$

$\langle D_2, I_2 \rangle$

$$\begin{aligned} D_2 &= \{ \langle \text{Socrates, 'nonactual'} \rangle \} \\ I_2(\text{'Philosopher'}) &= \{ \} \\ I_2(\text{'Cellist'}) &= \{ \langle \text{Socrates, 'nonactual'} \rangle \} \\ I_2(\text{'Sister'}) &= \{ \} \\ I_2(\text{'AR'}) &= \langle \text{Agustín, 'actual'} \rangle \end{aligned}$$

$\langle D_3, I_3 \rangle$

$$\begin{aligned} D_3 &= \{ \langle \text{Plato, 'nonactual'} \rangle \} \\ I_3(\text{'Philosopher'}) &= \{ \} \\ I_3(\text{'Cellist'}) &= \{ \langle \text{Plato, 'nonactual'} \rangle \} \\ I_3(\text{'Sister'}) &= \{ \} \\ I_3(\text{'AR'}) &= \langle \text{Agustín, 'actual'} \rangle \end{aligned}$$

$\langle D_4, I_4 \rangle$

$$\begin{aligned} D_4 &= \{ \langle \text{Agustín, 'actual'} \rangle, \langle \text{Plato, 'nonactual'} \rangle \} \\ I_4(\text{'Philosopher'}) &= \{ \langle \text{Plato, 'nonactual'} \rangle \} \\ I_4(\text{'Cellist'}) &= \{ \} \\ I_4(\text{'Sister'}) &= \{ \langle \langle \text{Plato, 'nonactual'} \rangle, \langle \text{Agustín, 'actual'} \rangle \rangle \} \\ I_4(\text{'AR'}) &= \langle \text{Agustín, 'actual'} \rangle \end{aligned}$$

These examples assume that the only non-logical expressions in L are 'Philosopher', 'Cellist', 'Sister' and 'AR', and that the domain of L is $\{ \text{Agustín} \}$.

Figure 1: Examples of a -worlds.

semantics based on *a*-worlds to verify SISTER, there must be linking amongst *a*-worlds. In particular, some *a*-world must represent a possibility whereby I have a sister who is a philosopher and another must represent a possibility whereby—as one is inclined to put it—*that very same individual* is a cellist rather than a philosopher.

Let us first see how linking gets addressed from a Lewisian perspective. l_1 and l_2 are Lewisian worlds: l_1 contains an individual a_1 who bears the right sort of similarity to me and an individual s_1 who is a_1 's sister and a philosopher; l_2 contains an individual s_2 who is a cellist. Accordingly, l_1 represents a possibility whereby my sister is a philosopher, and l_2 represents a possibility whereby someone is a cellist. But nothing so far guarantees linking. Nothing so far guarantees that—as one is inclined to put it—the individual l_1 represents as my sister is *the very individual* that l_2 represents as a cellist. What is needed for linking is that s_1 and s_2 be counterparts: that they be similar in the right sorts of respects.

The same maneuver can be used when it comes to *a*-worlds. Like the Lewisian, we shall use counterparthood amongst representations to capture linking. For Lewis, representations are counterparts just in case they are similar in the right sorts of respects. From the present perspective, we shall say that representations are counterparts just in case they are *identical* (though other ways of defining the counterpart relation could be used as well). Here is an example. The *a*-world $\langle D_1, I_1 \rangle$ represents a possibility whereby I have a sister who is a philosopher. Now consider *a*-worlds $\langle D_2, I_2 \rangle$ and $\langle D_3, I_3 \rangle$ from figure 1. Each of them represents a possibility whereby someone is a cellist rather than a philosopher. But only $\langle D_2, I_2 \rangle$ is linked to $\langle D_1, I_1 \rangle$. For $\langle D_1, I_1 \rangle$ and $\langle D_2, I_2 \rangle$ both employ $\langle \text{Socrates}, \text{'nonactual'} \rangle$ as a representation, and it is this that guarantees that—as one is inclined to put it—the individual who $\langle D_1, I_1 \rangle$ represents as my sister is *the very individual* that $\langle D_2, I_2 \rangle$ represents as a cellist. On the other hand, since $\langle D_3, I_3 \rangle$ represents a possibility whereby someone is a cellist by using $\langle \text{Plato}, \text{'nonactual'} \rangle$ rather than $\langle \text{Socrates}, \text{'nonactual'} \rangle$, what one gets is that—as one is inclined to put it—

the individual who $\langle D_1, I_1 \rangle$ represents as my sister is *distinct* from the individual that $\langle D_2, I_2 \rangle$ represents as a cellist.

It is important to keep in mind that an *a*-worlds-semantics is *not* a way of improving on the informal characterization of linking that I supplied a few paragraphs back. (Representations are linked if—as one is inclined to put it—they concern the same individual, even if the individual in question doesn’t exist.) In particular, it is *not* a way of dispensing with the qualifying phrase ‘as one is inclined to put it’.⁵ What an *a*-worlds-semantics delivers is an (actualistically acceptable) device for *representing* possibilities, which enjoys the following feature: it is clear when two representations are to be counted as linked. The reason this is helpful is that, as we shall see below, much of the theoretical work that can be carried out by quantifying over possibilities can be carried out by quantifying over representations of possibilities instead. So an *a*-worlds-semantics puts the actualist in a position to get on with certain kinds of theoretical work without having to worry about giving a proper characterization of linking.

2.2 Admissibility

There are *a*-worlds according to which someone is a married bachelor, and *a*-worlds according to which there might have been a human who wasn’t essentially human. Such representations need to be excluded from our semantics, on pain of getting the result that ‘ $\diamond(\exists x(\text{Married}(x) \wedge \text{Bachelor}(x)))$ ’ or ‘ $\diamond(\exists x(\text{Human}(x) \wedge \diamond(\exists y(y = x \wedge \neg \text{Human}(x))))$ ’ are true. What we need is a notion of *admissibility*. Armed with such a notion, one can say that ‘ $\ulcorner \diamond \phi \urcorner$ ’ is true just in case ϕ is true at some admissible *a*-world, and that ‘ $\ulcorner \Box \phi \urcorner$ ’ is true just in case ϕ is true at every admissible *a*-world.⁶

It is important to keep in mind the distinction between the *semantic* project of developing a Kripke-semantics for modal languages, on the one hand, and the project of accounting for the limits of metaphysical possibility, on the other. A semantics based on *a*-worlds is meant to address the former, but not the latter of these projects. Accord-

ingly, the notion of admissibility that is presupposed by the semantic project should be thought of as a placeholder for whatever limits on the metaphysically possible turn out to be uncovered by a separate investigation. (On reasonable assumptions, one can show that a suitable notion of admissibility is guaranteed to exist wherever the limits turn out to lie.⁷) A semantics based on *a*-words is compatible with different views about how to fix the limits of metaphysical possibility. One's account might appeal to a Principle of Recombination, for example, or to a set of 'basic modal truths'.⁸ (My own account is spelled out in chapters 2 and 5 of Rayo (forthcoming).)

Even if one appeals to an ontology of possible worlds in fixing the limits of metaphysical possibility, one might have good reasons for using *a*-worlds rather than possible worlds for the purposes of semantic theorizing. The easiest way to see this is by distinguishing between *sparse* and *abundant* conceptions of possible worlds.⁹ A sparse conception countenances worlds according to which there are objects that don't actually exist, but not worlds according to which it is true of *specific* non-existent objects that they exist. There is, for example, a possible world w_1 according to which I have a sister who is a philosopher and might have been a cellist rather than a philosopher, but no possible world according to which it is true of the specific individual who would have been my sister had w_1 obtained that she exists. (Not even w_1 is such a world, for even though w_1 is a world according to which I have a sister, it is not a world according to which it is true of some specific individual that she is my sister.) On an abundant conception of possible worlds, on the other hand, there are possible worlds according to which it is true of specific non-existent objects that they exist. There is, for instance, a possible world w_2 according to which it is true of the very individual who would have been my sister had w_1 been actualized that she is a cellist rather than a philosopher.

On the sparse conception of possible worlds, the existence of a world is conditional on the existence of the objects the world represents as existing, in the same sort of way that the existence of a set is conditional on the existence of its members. Had w_1

been actualized, I would have had a sister, and all manner of sets containing that very individual would have existed. But as things stand, my sister doesn't exist, and neither do sets having her as a member. Similarly—the story would go—had w_1 been actualized, I would have had a sister, and a world according to which that very individual is a cellist would have existed. But as things stand, my sister doesn't exist, and neither do possible worlds according to which she herself exists.

The absence of w_2 does not prevent a defender of the sparse conception from using possible worlds to determine a truth-value for SISTER. For, on the assumption that possible worlds track metaphysical possibility, the existence of w_1 is enough to guarantee that SISTER is true. But the absence of w_2 does mean that the sparse worlds do not by themselves deliver the ontology that would be needed to give a Kripke-semantics for a sentence like SISTER. For, as emphasized above, a Kripke-semantics will only count SISTER as true if the range of one's metalinguistic quantifiers contains both w_1 and w_2 .

Fortunately, a sparse ontology of possible worlds is enough to guarantee the existence of a notion of admissibility relative to which an a -worlds semantics assigns the right truth-value to every sentence in the language (see footnote 7). So it is open to the sparse theorist to use admissible a -worlds, rather than possible worlds, for the purposes of semantic theorizing. The upshot is not, of course, that one has done away with one's specialized modal ontology, since possible worlds may be needed for the project of pinning down the crucial notion of admissibility (or for the project of explaining what modal truth consists in). But by using a -worlds as the basis of one's semantics, the requirements on one's modal ontology are confined to needs of these non-semantic projects.

A related point can be made with respect to mere possibilia. The availability of a -worlds means that there is no need to postulate mere possibilia (or specialized surrogates, such as Plantinga's individual essences) as far as the project of developing a Kripke-semantics for modal languages is concerned.¹⁰

2.3 Interlude: The Principle of Representation

In previous sections I have made informal remarks about the ways in which a -worlds represent possibilities. The purpose of this interlude is to be more precise. (Uninterested readers may skip ahead to section 2.4.)

When $\langle D, I \rangle$ is considered in isolation from other a -worlds, anything that can be said about the possibility represented by $\langle D, I \rangle$ is a consequence of the following principle:

REPRESENTATION PRINCIPLE (Isolated-World Version)

Let s be a sentence of L , and suppose s says that p . Then:

according to the possibility represented by $\langle D, I \rangle$, p

if and only if

s is true at $\langle D, I \rangle$.

Accordingly, when considered in isolation from other a -worlds, $\langle D_2, I_2 \rangle$ and $\langle D_3, I_3 \rangle$ represent the *same* possibility. It is the possibility that there be exactly one thing and that it be a cellist but not a philosopher.

We shall normally assume that L (and therefore L^\diamond) contains a name for every object in the domain of L . With this assumption in place, the following is a consequence of the Representation Principle:

Suppose z is in the domain of L . Then z exists according to the possibility represented by $\langle D, I \rangle$ just in case $\langle z, \text{'actual'} \rangle$ is in D .

In particular, one gets the result that none of the objects in the domain of L exists according to the possibility represented by $\langle D_2, I_2 \rangle$ (since $\ulcorner \exists x(x = c) \urcorner$ is false at $\langle D_2, I_2 \rangle$ for any constant c in L), and that I exist according to the possibility represented by $\langle D_1, I_1 \rangle$ (since $\ulcorner \exists x(x = \text{AR}) \urcorner$ is true at $\langle D_1, I_1 \rangle$).

So much for considering a -worlds in isolation. When they are considered in the context of a *space* of a -worlds, linking plays a role. So there is slightly more to be said about the possibilities that they represent. Let A be a space of a -worlds and let $\langle D, I \rangle$ be in A . Then anything that can be said about the possibility represented by $\langle D, I \rangle$ in the context of A is a consequence of the following principle:

REPRESENTATION PRINCIPLE (Official Version)

1. Let s be a sentence of L^\diamond , and suppose s says that p . Then:

according to the possibility represented by $\langle D, I \rangle$ in the context of A , p

if and only if

s is true at $\langle D, I \rangle$ in the Kripke-model based on A .

2. Let $\langle D^*, I^* \rangle$ be an arbitrary a -world in A . Let $\ulcorner \exists x_1 \dots x_k \phi(x_1 \dots x_k) \urcorner$ and $\ulcorner \exists x_1 \dots x_k \gamma(x_1 \dots x_k) \urcorner$ be sentences of L^\diamond which say, respectively, that $x_1 \dots x_k$ are F and that $x_1 \dots x_k$ are G . Assume that $\ulcorner \exists x_1 \dots x_k \phi(x_1 \dots x_k) \urcorner$ is true at $\langle D, I \rangle$ in A and that $\ulcorner \exists x_1 \dots x_k \gamma(x_1 \dots x_k) \urcorner$ is true at $\langle D^*, I^* \rangle$ in A . Then:

as one is inclined to put it, some of the individuals that are F according to the possibility represented by $\langle D, I \rangle$ in the context of A are the very same individuals as some of the individuals that are G according to the possibility represented by $\langle D^*, I^* \rangle$ in the context of A

if and only if

one of the sequences of pairs that witnesses $\ulcorner \exists x_1 \dots x_k \phi(x_1 \dots x_k) \urcorner$ at $\langle D, I \rangle$ in A is identical to one of sequences of pairs that witnesses $\ulcorner \exists x_1 \dots x_k \gamma(x_1 \dots x_k) \urcorner$ at $\langle D^*, I^* \rangle$ in the Kripke-model

based on A .

When A includes both $\langle D_1, I_1 \rangle$ and $\langle D_2, I_2 \rangle$, the first clause yields the result that, according to the possibility represented by $\langle D_2, I_2 \rangle$ in the context of A , there is a cellist who might have been my sister. And the two clauses together yield the slightly stronger result that—as one is inclined to put it—the individual who is a cellist according to the possibility represented by $\langle D_2, I_2 \rangle$ in the context of A is the very same object as the individual who is my sister according to the possibility represented by $\langle D_1, I_1 \rangle$ in the context of A .

A consequence of the Principle of Representation is that the possibilities represented by a -worlds are not maximally specific. Suppose, for example, that the property of tallness is not expressible in L . Then the possibility represented by $\langle D_1, I_1 \rangle$ is compatible with a more specific possibility whereby my sister is tall and it is compatible with a more specific possibility whereby my sister is not tall. On the other hand, the possibilities represented by a -worlds *are* maximally specific as far as the language is concerned: one can only add specificity to the possibility represented by an a -world by employing distinctions that cannot be expressed in L^\diamond .

The Principle of Representation can be used to determine which properties of an a -world are essential to its representing the possibility that it represents, and which ones are merely artifactual. It entails, for example, that $\langle D_2, I_2 \rangle$ and $\langle D_3, I_3 \rangle$ represent the same possibility when considered in isolation, so any differences between them are merely artifactual. In particular, the use of $\langle \text{Socrates}, \text{'nonactual'} \rangle$ in $\langle D_2, I_2 \rangle$ is merely artifactual. On the other hand, $\langle D_2, I_2 \rangle$ and $\langle D_3, I_3 \rangle$ represent *different* possibilities when considered in the context of $\{\langle D_1, I_1 \rangle, \langle D_2, I_2 \rangle, \langle D_3, I_3 \rangle\}$. For whereas according to $\langle D_2, I_2 \rangle$ there is a cellist who might have been my sister, according to $\langle D_3, I_3 \rangle$ there is a cellist who couldn't have been my sister. So the use of $\langle \text{Socrates}, \text{'nonactual'} \rangle$ in $\langle D_2, I_2 \rangle$ is essential in the context of $\{\langle D_1, I_1 \rangle, \langle D_2, I_2 \rangle, \langle D_3, I_3 \rangle\}$. This is not to say, however, that a possibility whereby there is a cellist who might have been my sister can

only be represented by an a -world if the a -world contains $\langle \text{Socrates}, \text{'nonactual'} \rangle$. For the possibilities represented by $\langle D_1, I_1 \rangle$, $\langle D_2, I_2 \rangle$, and $\langle D_3, I_3 \rangle$ in the context of $\{\langle D_1, I_1 \rangle, \langle D_2, I_2 \rangle, \langle D_3, I_3 \rangle\}$ are precisely the possibilities represented by $\langle D_4, I_4 \rangle$, $\langle D_3, I_3 \rangle$, and $\langle D_2, I_2 \rangle$, respectively, in the context of $\{\langle D_4, I_4 \rangle, \langle D_3, I_3 \rangle, \langle D_2, I_2 \rangle\}$.

When I speak of the possibility that an a -world represents what I will usually have in mind is the possibility that is represented in the context of the space of all *admissible* a -worlds.

2.4 The dot-notation

I would like to introduce a further piece of notation: the dot. The dot is a bit like a function that takes objects represented to the objects doing the representing. Suppose I am seeing a play according to which I have a sister; applying the dot-function is like shifting my attention from a character in the play—my sister—to the actor who is representing my sister.

Consider the following two formulas:

$$[F(x)]_w \qquad [F(\dot{x})]_w$$

For w a fixed representation, the undotted formula is satisfied by all and only objects z such that w represents a possibility whereby z is an F ; the dotted formula, on the other hand, is satisfied by all and only objects z such that z is used by w to represent something as being an F . Thus, if p is a performance of a play according to which I have a sister, the actor playing my sister satisfies $[\text{Sister}(\text{AR}, \dot{x})]_p$ but not $[\text{Sister}(\text{AR}, x)]_p$ (since the performance uses the actor to represent someone as being my sister, but the performance does not represent a scenario whereby I have that actor as my sister). And I satisfy $[\exists y \text{Sister}(x, y)]_p$ but not $[\exists y \text{Sister}(\dot{x}, y)]_p$ (since the performance represents a scenario whereby I have a sister, but—unlike the actors and props—I am not used by

the performance to represent anything).

Now consider how the dot-notation might be cashed out from the perspective of a Lewisian. Let l_1 be a Lewisian world representing a possibility whereby I have a sister. Accordingly, l_1 contains an individual a_1 , who is my counterpart, and an individual s_1 , who is a_1 's sister. Now consider the following two formulas:

$$[\text{Sister}(\text{AR}, x)]_{l_1} \qquad [\text{Sister}(\text{AR}, \dot{x})]_{l_1}$$

From the perspective of the Lewisian, no inhabitant of the actual world satisfies the undotted formula. For no inhabitant of the actual world could have been my sister; so—on the assumption that Lewisian worlds track metaphysical possibility—no inhabitant of the actual world is such that l_1 represents a possibility whereby *she* is my sister. The dotted formula, on the other hand, is satisfied by s_1 , since she is used by l_1 to represent something as being my sister.

Here is a second pair of examples:

$$[\exists y \text{Sister}(x, y)]_{l_1} \qquad [\exists y \text{Sister}(\dot{x}, y)]_{l_1}$$

The undotted formula is satisfied by me, since l_1 represents a possibility whereby I have a sister. But it is *not* satisfied by a_1 . For although it is true that a_1 has a sister in l_1 , l_1 represents a possibility whereby *I* have a sister, not a possibility whereby my counterpart has a sister. The dotted formula, on the other hand, is satisfied by a_1 , since a_1 is used by l_1 to represent something as having a sister (i.e. me). But the dotted formula is *not* satisfied by me, since it is only the inhabitants of l_1 that do any representing for l_1 , and I am an inhabitant of the actual world.

Let me now illustrate how the dot-notation works from the perspective of the modal actualist, with a -worlds in place of Lewisian worlds. (A detailed semantics is given in Appendix A.) Here is the first pair of examples (where w_1 is the a -world $\langle D_1, I_1 \rangle$ from

section 2.1):

$$[\text{Sister}(\text{AR}, x)]_{w_1} \qquad [\text{Sister}(\text{AR}, \dot{x})]_{w_1}$$

Since $I_1(\text{'Sister'}) = \{\langle\langle\text{Agustín, 'actual'}\rangle, \langle\text{Socrates, 'nonactual'}\rangle\rangle\}$, w_1 represents a possibility whereby I have a sister who doesn't actually exist. Accordingly, from the perspective of a modal actualist, there is no z such that w_1 represents the possibility that z is my sister. So, from the perspective of the modal actualist, *nothing* satisfies the undotted formula. The dotted formula, on the other hand, is satisfied by $\langle\text{Socrates, 'nonactual'}\rangle$, since $\langle\text{Socrates, 'nonactual'}\rangle$ is used by w_1 to represent something as being my sister.

Now consider the second pair of examples:

$$[\exists y \text{Sister}(x, y)]_{w_1} \qquad [\exists y \text{Sister}(\dot{x}, y)]_{w_1}$$

Since the pair $\langle\langle\text{Agustín, 'actual'}\rangle, \langle\text{Socrates, 'nonactual'}\rangle\rangle$ is in $I_1(\text{'Sister'})$, w_1 represents a possibility whereby I have a sister. The undotted formula is therefore satisfied by me. But it is *not* satisfied by $\langle\text{Agustín, 'actual'}\rangle$ because w_1 does not represent a possibility whereby any ordered-pairs have sisters. The dotted formula, on the other hand, is satisfied by $\langle\text{Agustín, 'actual'}\rangle$, since $\langle\text{Agustín, 'actual'}\rangle$ is used by w_1 to represent something as having a sister (i.e. me). But it is *not* satisfied by me, since it is only ordered-pairs that do any representing in w_1 , and I am not an ordered-pair.

2.5 Inference in a language with the dot-notation

The semantics for a -worlds that is supplied in Appendix A guarantees the truth of every instance of the following schemas:

1. Validity

$$[\psi]_w \text{ (where } \psi \text{ is valid in a negative free logic)}$$

2. Conjunction

$$[\psi \wedge \theta]_w \leftrightarrow ([\psi]_w \wedge [\theta]_w)$$

3. Negation

$$[\neg\psi]_w \leftrightarrow \neg[\psi]_w$$

4. Quantification

$$[\exists y(\phi(y))]_w \leftrightarrow \exists y([\dot{y} = \dot{y}]_w \wedge [\phi(\dot{y})]_w) \text{ (where } y \text{ is a non-world-variable)}^{11}$$

$$[\exists w'(\phi)]_w \leftrightarrow \exists w'([\phi]_{w'}) \text{ (where } w' \text{ is a world-variable)}$$

5. Trivial accessibility¹²

$$[[\phi]_w]_{w'} \leftrightarrow [\phi]_w$$

6. Identity

$$[v = v]_w \leftrightarrow [\exists y(y = v)]_w$$

$$x = y \rightarrow ([\phi(\dot{x})]_w \rightarrow [\phi(\dot{y})]_w)$$

$$[\dot{x} = \dot{y}]_w \rightarrow ([\dot{x} = \dot{x}]_w \wedge x = y)$$

(where v may occur dotted or undotted)

7. Atomic Predication

$$[F_j^n(v_1, \dots, v_n)]_w \rightarrow ([v_1 = v_1]_w \wedge \dots \wedge [v_n = v_n]_w)$$

(where the v_i may occur dotted or undotted)

8. Names

$$[\psi(c)]_w \leftrightarrow \exists x(x = c \wedge [\psi(x)]_w) \text{ (for } c \text{ a non-empty name)}$$

Schemas 2–5 are enough to guarantee that any sentence in the actualist's language is equivalent to a sentence in which only atomic formulas occur within the scope of ' $[\dots]_w$ '.

For instance, the actualist rendering of ' $\diamond(\exists x(\text{Phil}(x) \wedge \diamond(\neg\text{Phil}(x))))$ ':

$$\exists w[\exists x(\text{Phil}(x) \wedge \exists w'([\neg\text{Phil}(x)]_{w'}))]_w$$

is equivalent to

$$\exists w \exists x ([\text{Phil}(x)]_w \wedge \exists w' (\neg [\text{Phil}(x)]_{w'})).$$

As a result, the dot-notation allows a language containing the modal operator ‘ $[\dots]_w$ ’ to have the inferential behavior of a (non-modal) first-order language.

2.6 The expressive power of the dot-notation

In this section we shall see that a suitably qualified version of the following claim is true: anything the Lewisian can say, the modal actualist can say too—by using the dot-notation.

Here is an example. The Lewisian can use her mighty expressive resources to capture a version of the following thought:

LINKING

There are possible worlds w_1 and w_2 with the following properties: according to w_1 , there is an individual who is a philosopher; according to w_2 , that very individual is a cellist.

It is done as follows:

$$\exists w_1 \exists w_2 \exists x_1 \exists x_2 (I(x_1, w_1) \wedge I(x_2, w_2) \wedge \text{Phil}(x_1) \wedge \text{Cellist}(x_2) \wedge C(x_1, x_2))$$

(*Read:* There are Lewisian worlds w_1 and w_2 and individuals x_1 and x_2 such that: (a) x_1 is an inhabitant of w_1 and x_2 is an inhabitant of w_2 , (b) x_1 is a philosopher and x_2 is a cellist, and (c) x_1 and x_2 are counterparts.)

How might this be emulated by a modal actualist equipped with the dot-notation? Consider what happens when one treats the variables in the Lewisian rendering of LINKING as ranging over (admissible) a -worlds rather than Lewisian worlds, and carries out the

following replacements:

$$\begin{aligned} I(x_n, w_n) &\longrightarrow [\dot{x}_n = \dot{x}_n]_{w_n} \\ \text{Phil}(x_n) &\longrightarrow [\text{Phil}(\dot{x}_n)]_{w_n} \\ \text{Cellist}(x_n) &\longrightarrow [\text{Cellist}(\dot{x}_n)]_{w_n} \\ C(x_n, x_m) &\longrightarrow x_n = x_m \end{aligned}$$

The result is this:

$$\exists w_1 \exists w_2 \exists x_1 \exists x_2 ([\dot{x}_1 = \dot{x}_1]_{w_1} \wedge [\dot{x}_2 = \dot{x}_2]_{w_2} \wedge [\text{Phil}(\dot{x}_1)]_{w_1} \wedge [\text{Cellist}(\dot{x}_2)]_{w_2} \wedge x_1 = x_2)$$

(*Read:* There are admissible a -worlds w_1 and w_2 and objects x_1 and x_2 such that: (a) x_1 is used by w_1 to represent something and x_2 is used by w_2 to represent something, (b) x_1 is used by w_1 to represent a philosopher and x_2 is used by w_2 to represent a cellist, and (c) $x_1 = x_2$.)

or equivalently:

$$\exists w_1 \exists w_2 \exists x ([\text{Phil}(\dot{x})]_{w_1} \wedge [\text{Cellist}(\dot{x})]_{w_2})$$

(*Read:* There are admissible a -worlds w_1 and w_2 and an object x such that: x is used by w_1 to represent a philosopher and x is used by w_2 to represent a cellist.)

What gives the actualist's method its punch is the fact that it generalizes: one can show that there is a systematic transformation of arbitrary Lewisian sentences into dotted actualist sentences which preserves truth-values and inferential connections.¹³ (See Appendix B for details.)

The actualist's transformation-method does not preserve meaning—where Lewisian sentences quantify over Lewisian possibilities, their actualist transformation quantify over

a -worlds and ordered-pairs. But meaning-preservation is not what the actualist wants, since she doesn't want to countenance Lewisian possibilities. What she wants is a way of enjoying the theoretical benefits of quantification over Lewisian possibilities within the sober confines of an actualist framework. Here are two examples of ways in which she is able to do so:

1. *Firstorderizing Modal Sentences*

By quantifying over Lewisian possibilities, the Lewisian is able to render any sentence in the language of first-order modal logic in (non-modal) first-order terms. The sentence

$$\diamond(\exists x(\text{Phil}(x) \wedge \diamond(\text{Cellist}(x))))$$

(*read*: there might have been a philosopher who might have been a cellist),

for example, gets rendered as the (non-modal) first-order sentence:

$$\exists w_1 \exists w_2 \exists x_1 \exists x_2 (\text{I}(x_1, w_1) \wedge \text{I}(x_2, w_2) \wedge \text{Phil}(x_1) \wedge \text{Cellist}(x_2) \wedge \text{C}(x_1, x_2))$$

And Lewis (1968) shows that it can be done in general.¹⁴

The (non-modal) firstorderizability of modal sentences brings two immediate advantages. The first is that it allows one to think of the inferential connections amongst modal sentences in terms of the inferential connections amongst the corresponding non-modal sentences; the second is that it allows one to read off a semantics for modal sentences from the semantics of the corresponding non-modal sentences.

The actualist transformation-method allows actualists equipped with the dot-notation to enjoy both of these advantages.

2. *Characterizing Intensions*

On a standard way of doing intensional semantics for natural languages, characterizing the semantic value of an expression calls for quantification over possibilia.¹⁵ Oversimplifying a bit, the semantic value of, e.g. ‘philosopher’ might be taken to be the set of pairs $\langle w, z \rangle$ where w is a possible world and z is an (actual or merely possible) individual who is a philosopher at w .

As emphasized in Lewis (1970), the Lewisian is able to do the job by quantifying over Lewisian possibilia:

$$\llbracket \text{‘philosopher’} \rrbracket = \{ \langle w, z \rangle : I(z, w) \wedge \text{Phil}(z) \}$$

But the actualist transformation-method allows actualists armed with the dot-notation to follow suit, by quantifying over a -worlds and ordered pairs:

$$\llbracket \text{‘philosopher’} \rrbracket = \{ \langle w, z \rangle : [\dot{z} = z]_w \wedge [\text{Phil}(\dot{z})]_w \}$$

or, equivalently,

$$\llbracket \text{‘philosopher’} \rrbracket = \{ \langle w, z \rangle : [\text{Phil}(\dot{z})]_w \}.$$

Since the actualist transformation-method preserves inferential connections, the actualist semantics is guaranteed to deliver the same theorems as its Lewisian counterpart. And since the transformation-method preserves truth-value the actualist’s axioms will be true just in case their Lewisian counterparts would count as true from the perspective of the Lewisian.

In particular, one can expect the actualist semantics to deliver every instance of the (world-relative) T-schema. For instance:

$$\text{True}(\text{‘}\exists x \text{Phil}(x)\text{’}, w) \leftrightarrow ([\exists x \text{Phil}(x)]_w)$$

or, equivalently,

$$\text{True}(\text{'}\exists x \text{ Phil}(x)\text{'}, w) \leftrightarrow \exists x([\text{Phil}(\dot{x})]_w)$$

(*Read:* The object-language sentence ‘ $\exists x \text{ Phil}(x)$ ’ is true at admissible a -world w just in case there is an individual which is used by w to represent a philosopher.)

2.7 Limitations of the proposal

A -worlds are subject to an important limitation.¹⁶ Whereas differences between a -worlds are no more fine-grained than is required to make distinctions expressible in one’s language, there might be differences amongst Lewisian worlds too fine-grained to be expressed in one’s language.¹⁷

Because of this limitation, some of the metaphysical work that the Lewisian gets out of Lewisian possibilities cannot be replicated by an actualist equipped with the dot-notion. Here is an example. Lewis (1986) treats properties as sets of worldbound individuals. Up to a certain point, the actualist is able to follow suit. When the Lewisian claims that the property of being a philosopher is to be identified with the set of philosophers inhabiting actual or non-actual Lewisian worlds, for instance, the actualist could claim that the set $\{\langle z, w \rangle : [\text{Phil}(\dot{z})]_w\}$ is to be used as a *surrogate* for the property of being a philosopher. But the strategy breaks down when it comes to properties making finer distinctions than can be expressed in one’s language.¹⁸

In general, whether or not the actualist’s limitation turns out to get in the way will depend on whether the job at hand calls for using possibilities to make finer distinctions than can be expressed in one’s language. When the job at hand is a piece of semantic theorizing the extra resources are unnecessary: since a semantic theory is ultimately an effort to explain how language is used, it need not be concerned with distinctions too fine-grained to figure in our explanations. But when the job at hand is metaphysical

reduction, matters are otherwise.¹⁹

Appendices

A A semantics

I give a formal semantics for a language \mathcal{L}^w which allows for empty names and contains both the intensional operator $[\dots]_w$ (read ‘according to w , ...’) and the dot-notation. \mathcal{L}^w consists of the following symbols:

1. for $n > 0$, the n -place (non-modal) predicate letters: $\ulcorner F_1^n \urcorner, \ulcorner F_2^n \urcorner, \dots$ (each with an intended interpretation);
2. for $n > 0$, the one-place modal predicate letter $\ulcorner B_n \urcorner$ (each with an intended interpretation);
3. the identity symbol ‘=’;
4. for $n > 0$, the individual non-empty constant-letter $\ulcorner c_n \urcorner$ (each with an intended referent);
5. for $n > 0$, the individual empty constant-letter $\ulcorner e_n \urcorner$;
6. the individual constant ‘ α ’
7. the dot ‘.’;
8. the monadic sentential operator ‘ $[\dots]$ ’;
9. the monadic sentential operator ‘ \neg ’;
10. the dyadic sentential operator ‘ \wedge ’;
11. the quantifier-symbol ‘ \exists ’;
12. the modal variables: ‘ w ’, ‘ v ’ ‘ u ’ with or without numerical subscripts;

13. the non-modal variables: ‘ x ’, ‘ y ’, ‘ z ’ with or without numerical subscripts;
14. the auxiliaries ‘(’ and ‘)’.

Undotted terms and formulas are defined as follows:

1. any modal variable is an undotted modal term;
2. ‘ α ’ is an undotted modal term;
3. any non-modal variable or individual constant-letter is an undotted non-modal term;
4. if τ_1, \dots, τ_n are undotted non-modal terms, then $\lceil F_i^n(\tau_1, \dots, \tau_n) \rceil$ is an undotted formula;
5. if τ_1 and τ_2 are either both undotted non-modal terms or both undotted modal terms, then $\lceil \tau_1 = \tau_2 \rceil$ is an undotted formula;
6. if w is an undotted modal term, then $\lceil B_i(w) \rceil$ is an undotted formula;
7. if ϕ is an undotted formula and w is an undotted modal term, then $\lceil [\phi]_w \rceil$ is an undotted formula;
8. if v is an undotted (modal or non-modal) variable and ϕ is an undotted formula, then $\lceil \exists v(\phi) \rceil$ is an undotted formula;
9. if ϕ and ψ are undotted formulas, then $\lceil \neg\phi \rceil$ and $\lceil (\phi \wedge \psi) \rceil$ are undotted formulas;
10. nothing else is an undotted term or formula.

A *non-modal term* is either an undotted non-modal term or the result of dotting a non-modal variable; a *modal term* is an undotted modal term; a *formula* is the result of dotting any free or externally bounded occurrences of non-modal variables in an undotted formula.²⁰

Next, we characterize the notion of an a -world and of a variable assignment. An a -world is a pair $\langle D, I \rangle$ with the following features:

1. D is a set of individuals in the range of the non-modal variables, each of which is either of the form $\langle x, \text{'actual'} \rangle$ or of the form $\langle x, \text{'nonactual'} \rangle$.²¹
2. I is a function assigning a subset of D to each 1-place predicate-letter, and a subset of D^n to each n -place predicate letter (for $n < 1$). In addition, if e is an empty name, I may or may not assign a referent to e (and if a referent is assigned, it may or may not be in D).

The *actualized* a -world $\langle D_\alpha, I_\alpha \rangle$ will be singled out for special attention. D_α is the set of pairs $\langle z, \text{'actual'} \rangle$ for z an individual in the range of the non-modal variables; and $I_\alpha(\text{'F}_j^n)$ is the set of sequences $\langle \langle z_1, \text{'actual'} \rangle, \dots, \langle z_n, \text{'actual'} \rangle \rangle$ such that $z_1 \dots z_n$ are in the range of the non-modal variables and satisfy F .

A variable assignment is a function σ with the following features:

1. σ assigns an a -world to each modal variable.
2. σ assigns an individual to each non-modal variable.

This puts us in a position to characterize notions of quasi-denotation and quasi-satisfaction. (With a suitable notion of admissibility is on board, one can characterize truth and satisfaction for \mathcal{L}^w . Satisfaction is the special case of quasi-satisfaction in which attention is restricted to admissible a -worlds, and truth is the special case of quasi-truth in which attention is restricted to admissible a -worlds.) For v a non-modal variable, σ a variable assignment, ϕ a formula and w an a -world, we characterize the quasi-denotation function $\delta_{\sigma,w}(v)$ and the quasi-satisfaction predicate $Sat(\phi, \sigma)$. In addition, we characterize an auxiliary (a -world-relative) quasi-satisfaction predicate $Sat(\phi, \sigma, w)$. We proceed axiomatically, by way of the following clauses:

- If v is a (modal or non-modal) variable, $\delta_{\sigma,w}(v)$ is $\sigma(v)$;
- If v is a non-modal variable, w is an a -world and $\sigma(v)$ is an ordered pair of the form $\langle z, \text{'actual'} \rangle$, then $\delta_{\sigma,w}(\ulcorner v \urcorner)$ is the first member of $\sigma(v)$; otherwise $\delta_{\sigma,w}(\ulcorner v \urcorner)$ is undefined;
- if c is a non-empty constant-letter and w is an a -world, $\delta_{\sigma,w}(c)$ is the intended referent of c .
- if e is an empty constant-letter and w is an a -world, $\delta_{\sigma,w}(e)$ is the w -referent of e if there is one, and is otherwise undefined;
- $\delta_{\sigma,w}(\ulcorner \alpha \urcorner)$ is $\langle D_\alpha, I_\alpha \rangle$;
- if τ_1 and τ_2 are terms (both of them modal or both of them non-modal) and neither of them is an empty constant-letter, then $Sat(\ulcorner \tau_1 = \tau_2 \urcorner, \sigma)$ if and only if $\delta_{\sigma,w}(\tau_1) = \delta_{\sigma,w}(\tau_2)$ for arbitrary w ;
- if τ_1 and τ_2 are non-modal terms at least one of which is an empty constant-letter, then $\text{not-}Sat(\ulcorner \tau_1 = \tau_2 \urcorner, \sigma)$;
- if τ_1, \dots, τ_n are non-modal terms none of which is an empty constant-letter, then $Sat(\ulcorner F_i^n(\tau_1, \dots, \tau_n) \urcorner, \sigma)$ if and only if $F_i^n(\delta_{\sigma,w}(\tau_1), \dots, \delta_{\sigma,w}(\tau_n))$, where w is arbitrary and $\ulcorner F_i^n \urcorner$ is intended to express F_i^n -ness;
- if τ_1, \dots, τ_n are non-modal terms at least one of which is an empty constant-letter, then $\text{not-}Sat(\ulcorner F_i^n(\tau_1, \dots, \tau_n) \urcorner, \sigma)$;
- if v is a modal variable, $Sat(\ulcorner B_i(v) \urcorner, \sigma)$ if and only if $B_i(\delta_{\sigma,w}(v))$ for arbitrary w , where B_i is intended to express B_i -ness;
- if v is a non-modal variable, $Sat(\ulcorner \exists v(\phi) \urcorner, \sigma)$ if and only if there is an individual z in the range of the non-modal variables such that $Sat(\ulcorner \phi \urcorner, \sigma^{v/z})$, where $\sigma^{v/z}$ is just

like σ except that it assigns z to v ;

- if v is a modal variable, $Sat(\ulcorner \exists v(\phi) \urcorner, \sigma)$ if and only if there is an a -world z such that $Sat(\ulcorner \phi \urcorner, \sigma^{v/z})$, where $\sigma^{v/z}$ is just like σ except that it assigns z to v ;
- $Sat(\ulcorner \neg \phi \urcorner, \sigma)$ if and only if it is not the case that $Sat(\ulcorner \phi \urcorner, \sigma)$;
- $Sat(\ulcorner \phi \wedge \psi \urcorner, \sigma)$ if and only if $Sat(\ulcorner \phi \urcorner, \sigma)$ and $Sat(\ulcorner \psi \urcorner, \sigma)$;
- $Sat(\ulcorner [\phi]_w \urcorner, \sigma)$ if and only if $Sat(\ulcorner \phi \urcorner, \sigma')$, where $\sigma'(x) = \sigma(x)$ for x a modal variable and $\sigma'(x) = \langle \sigma(x), \text{'actual'} \rangle$ for x a non-modal variable;
- if τ_1 and τ_2 are non-modal terms neither of which is an empty constant-letter without a w -reference, then $Sat(\ulcorner \tau_1 = \tau_2 \urcorner, \sigma, w)$ if and only if $\delta_{\sigma, w}(\tau_1)$ is in the domain of w and is identical to $\delta_{\sigma, w}(\tau_2)$;
- if τ_1 and τ_2 are non-modal terms at least one of which is a constant-letter without a w -reference, then $\text{not-}Sat(\ulcorner \tau_1 = \tau_2 \urcorner, \sigma, w)$;
- if τ_1 and τ_2 are modal terms, then $Sat(\ulcorner \tau_1 = \tau_2 \urcorner, \sigma, w)$ if and only if $\delta_{\sigma, w}(\tau_1) = \delta_{\sigma, w}(\tau_2)$ for arbitrary w ;
- if τ_1, \dots, τ_n are non-modal terms none of which is a constant-letter without a w -reference, then $Sat(\ulcorner F_i^n(\tau_1, \dots, \tau_n) \urcorner, \sigma, w)$ if and only if $\langle \delta_{\sigma, w}(\tau_1), \dots, \delta_{\sigma, w}(\tau_n) \rangle$ is in the w -extension of $\ulcorner F_i^n \urcorner$;
- if τ_1, \dots, τ_n are non-modal terms at least one of which is a constant-letter without a w -reference, then $\text{not-}Sat(\ulcorner F_i^n(\tau_1, \dots, \tau_n) \urcorner, \sigma, w)$;
- if v is a modal term, $Sat(\ulcorner B_i(v) \urcorner, \sigma, w)$ if and only if $B_i(\delta_{\sigma, w}(v))$, where $\ulcorner B_i \urcorner$ is intended to express B_i -ness;

- if v is a non-modal variable, $Sat(\ulcorner \exists v(\phi) \urcorner, \sigma, w)$ if and only if there is an individual z in the domain of w such that $Sat(\ulcorner \phi \urcorner, \sigma^{v/z}, w)$, where $\sigma^{v/z}$ is just like σ except that it assigns z to v ;
- if v is a modal variable, $Sat(\ulcorner \exists v(\phi) \urcorner, \sigma, w)$ if and only if there is an a -world z such that (a) any empty constant-letter which is assigned a referent by w is assigned the same referent by z , and (b) $Sat(\ulcorner \phi \urcorner, \sigma^{v/z}, w)$, where $\sigma^{v/z}$ is just like σ except that it assigns z to v ;
- $Sat(\ulcorner \neg \phi \urcorner, \sigma, w)$ if and only if it is not the case that $Sat(\phi, \sigma, w)$;
- $Sat(\ulcorner \phi \wedge \psi \urcorner, \sigma, w)$ if and only if $Sat(\phi, \sigma, w)$ and $Sat(\psi, \sigma, w)$;
- $Sat(\ulcorner [\phi]_u \urcorner, \sigma, w)$ if and only if $Sat(\phi, \sigma, \sigma(u))$.

Finally, we say that a formula ϕ is *quasi-true* if and only if $Sat(\phi, \sigma)$ for any variable assignment σ .

B A transformation

The purpose of this appendix is to explain how the transformation from Lewisian sentences to dotted actualist sentences works in general. I shall assume that the Lewisian language is a two-sorted first-order languages, with *world-variables* ' w_1 ', ' w_2 ', etc. ranging over Lewisian worlds, and *individual-variables* ' x_1 ', ' x_2 ', etc. ranging over world-bound individuals in the Lewisian pluriverse. The language contains no function-letters; there is a world-constant ' α ' referring to the actual Lewisian world and individual-constants ' c_1 ', ' c_2 ', etc. referring to world-bound individuals. The only atomic predicates are 'I' (which takes an individual-variable and a world-variable), 'C' (which takes two individual-variables), '=' (which takes (i) two world-variables, (ii) two individual-variables, (iii) a

world-constant and a world-variable, or (*iv*) an individual-constant and an individual-variable) and, for each j , $\lceil P_j^n \rceil$ (which takes n individual-variables). (If one likes, one can also take the language to contain set-theoretic vocabulary.) Finally, I shall assume that universal quantifiers are defined in terms of existential quantifiers in the usual way, and that logical connectives other than ‘ \wedge ’ and ‘ \neg ’ are defined in terms of ‘ \wedge ’ and ‘ \neg ’ in the usual way.

The plan is to proceed in two steps. The first is to get the Lewisian sentence into a certain kind of normal form; the second is to convert the normal-form sentence into a dotted actualist sentence. Here is a recipe for getting an arbitrary Lewisian sentence into normal form:

1. Start by relabeling variables in such a way that no world-variable has the same index as an individual-variable;
2. next, replace each occurrence of $\lceil I(x_j, w_k) \rceil$ by $\lceil w_j = w_k \rceil$;
3. then replace each occurrence of the atomic formula $\lceil P_j^n(x_{k_1}, \dots, x_{k_n}) \rceil$ by

$$(P_j^n(x_{k_1}, \dots, x_{k_n}) \wedge w_{k_1} = w_{k_2} \wedge \dots \wedge w_{k_1} = w_{k_n}),$$

each occurrence of $\lceil c_j = x_k \rceil$ by $\lceil (c_j = x_k \wedge \alpha = w_k) \rceil$, and each occurrence of $\lceil x_j = x_k \rceil$ by $\lceil (x_j = x_k \wedge w_j = w_k) \rceil$.²²

4. finally, replace each occurrence of $\lceil \exists x_j(\dots) \rceil$ by $\lceil \exists x_j \exists w_j (I(x_j, w_j) \wedge \dots) \rceil$.

On the assumption that $\lceil P_j^n \rceil$ is projectable (and, hence, that $\lceil P_j^n(x_{k_1}, \dots, x_{k_n}) \rceil$ can only be true if x_{k_1}, \dots, x_{k_n} are world-mates), it is easy to verify that this procedure respects truth-value. (For a characterization of projectability, see footnote 13.)

Here is an example. The Lewisian sentence

$$\exists w_{17} \exists x_2 \exists x_5 (I(x_2, w_{17}) \wedge \text{Sister}(x_2, x_5))$$

(*Read:* There is a Lewisian world w_{17} an individual x_2 and an individual x_5 such that x_2 is an inhabitant of w_{17} and x_2 has x_5 as a sister.)

gets rewritten as:

$$\exists w_{17} \exists x_2 \exists w_2 (\mathbf{I}(x_2, w_2) \wedge \exists x_5 \exists w_5 (\mathbf{I}(x_5, w_5) \wedge w_2 = w_{17} \wedge \text{Sister}(x_2, x_5) \wedge w_2 = w_5))$$

(*Read:* There is a Lewisian world w_{17} , an individual x_2 inhabiting Lewisian world w_2 , and an individual x_5 inhabiting Lewisian world w_5 such that w_2 is identical to w_{17} , x_2 has x_5 as a sister and w_2 is identical to w_5 .)

Once one has a Lewisian sentence in normal form, it can be transformed into a dotted actualist sentence by carrying out the following replacements:

$$\begin{aligned} \mathbf{I}(x_j, w_j) &\longrightarrow [\exists y(y = \dot{x}_j)]_{w_j} \\ \mathbf{P}_j^n(x_{k_1}, \dots, x_{k_n}) &\longrightarrow [\mathbf{P}_j^n(\dot{x}_{k_1}, \dots, \dot{x}_{k_n})]_{w_{k_1}} \\ \mathbf{C}(x_j, x_k) &\longrightarrow x_j = x_k \\ c_j = x_k &\longrightarrow [c_j = \dot{x}_k]_{w_k} \end{aligned}$$

Here is an example. The Lewisian rendering of ‘ $\diamond(\exists x_1 \exists x_2 \text{Sister}(x_1, x_2))$ ’ is

$$\exists w_3 \exists x_1 \exists x_2 (\mathbf{I}(x_1, w_3) \wedge \mathbf{I}(x_2, w_3) \wedge \text{Sister}(x_1, x_2))$$

whose normal form

$$\begin{aligned} \exists w_3 \exists x_1 \exists w_1 (\mathbf{I}(x_1, w_1) \wedge \exists x_2 \exists w_2 (\mathbf{I}(x_2, w_2) \wedge w_1 = w_3 \wedge w_1 = w_3 \wedge \\ \text{Sister}(x_1, x_2) \wedge w_1 = w_2)) \end{aligned}$$

gets transformed by the actualist into

$$\exists w_3 \exists x_1 \exists w_1 ([\exists y (y = \dot{x}_1)]_{w_1} \wedge \exists x_2 \exists w_2 ([\exists y (y = \dot{x}_2)]_{w_2} \wedge w_1 = w_3 \wedge w_1 = w_3 \wedge [\text{Sister}(\dot{x}_1, \dot{x}_2)]_{w_1} \wedge w_1 = w_2))$$

which boils down to:

$$\exists w_3 \exists x_1 \exists x_2 ([\text{Sister}(\dot{x}_1, \dot{x}_2)]_{w_3})$$

(*Read:* There is are objects x_1 and x_2 and an admissible a -world w such that x_1 and x_2 are used by w to represent someone's having a sister.)

and is guaranteed by the Appendix A semantics to be equivalent to

$$\exists w_3 ([\exists x_1 \exists x_2 \text{Sister}(x_1, x_2)]_{w_3})$$

(*Read:* There is an admissible a -world w according to which there are individuals x_1 and x_2 such that x_1 has x_2 as a sister.)

which is the actualist's rendering of ' $\diamond(\exists x_1 \exists x_2 \text{Sister}(x_1, x_2))$ '.

Here is a slightly more complex example. The actualist's rendering of

There are possible worlds w_1 and w_2 with the following properties: according to w_1 , there is an individual who is my sister and a philosopher; according to w_2 , that very individual is a cellist rather than a philosopher.

is

$$\begin{aligned} &\exists w_4 \exists w_5 \exists x_1 \exists x_2 \exists x_3 (\\ &\quad \text{I}(x_1, w_4) \wedge \text{I}(x_2, w_4) \wedge \text{I}(x_3, w_5) \wedge \\ &\quad \exists x_6 (\text{AR} = x_6 \wedge \text{C}(x_6, x_1)) \wedge \text{C}(x_2, x_3) \wedge \\ &\quad \text{Sister}(x_1, x_2) \wedge \text{Phil}(x_2) \wedge \text{Cellist}(x_3) \wedge \neg \text{Phil}(x_3)) \end{aligned}$$

whose normal form

$$\begin{aligned} & \exists w_4 \exists w_5 \exists x_1 \exists w_1 (\mathbf{I}(x_1, w_1) \wedge \exists x_2 \exists w_2 (\mathbf{I}(x_2, w_2) \wedge \exists x_3 \exists w_3 (\mathbf{I}(x_3, w_3) \wedge \\ & \quad w_1 = w_4 \wedge w_2 = w_4 \wedge w_3 = w_5 \wedge \\ & \quad \exists x_6 \exists w_6 (\mathbf{I}(x_6, w_6) \wedge \mathbf{AR} = x_6 \wedge \alpha = w_6 \wedge \mathbf{C}(x_6, x_1)) \wedge \mathbf{C}(x_2, x_3) \wedge \\ & \quad \mathbf{Sister}(x_1, x_2) \wedge \mathbf{Phil}(x_2) \wedge \mathbf{Cellist}(x_3) \wedge \neg \mathbf{Phil}(x_3)))) \end{aligned}$$

gets transformed by the actualist into

$$\begin{aligned} & \exists w_4 \exists w_5 \exists x_1 \exists w_1 ([\exists y (y = \dot{x}_1)]_{w_1} \wedge \exists x_2 \exists w_2 ([\exists y (y = \dot{x}_2)]_{w_2} \wedge \exists x_3 \exists w_3 ([\exists y (y = \dot{x}_3)]_{w_3} \wedge \\ & \quad w_1 = w_4 \wedge w_2 = w_4 \wedge w_3 = w_5 \wedge \\ & \quad \exists x_6 \exists w_6 ([\exists y (y = \dot{x}_6)]_{w_6} \wedge [\mathbf{AR} = \dot{x}_6]_{w_6} \wedge \alpha = w_6 \wedge x_6 = x_1) \wedge x_2 = x_3 \wedge \\ & \quad [\mathbf{Sister}(\dot{x}_1, \dot{x}_2)]_{w_1} \wedge w_1 = w_2 \wedge [\mathbf{Phil}(\dot{x}_2)]_{w_2} \wedge [\mathbf{Cellist}(\dot{x}_3)]_{w_3} \wedge \neg [\mathbf{Phil}(\dot{x}_3)]_{w_3})) \end{aligned}$$

which boils down to

$$\exists w_4 \exists w_5 \exists x_2 ([\mathbf{Sister}(\mathbf{AR}, \dot{x}_2)]_{w_4} \wedge [\mathbf{Phil}(\dot{x}_2)]_{w_4} \wedge [\mathbf{Cellist}(\dot{x}_2)]_{w_5} \wedge [\neg \mathbf{Phil}(\dot{x}_2)]_{w_5})$$

(*Read:* there are admissible a -worlds w_4 and w_5 and an individual x_2 such that: (i) x_2 is used by w_4 to represent my sister and a philosopher, and (ii) x_2 is used by w_5 to represent a cellist rather than a philosopher.)

Finally, we show that there is a notion of a -world admissibility which guarantees that the actualist transformation of an arbitrary Lewisian sentence has the truth-value that the Lewisian sentence would receive on its intended interpretation. (The proof relies on the assumption that the Lewisian language is rich enough—and the space of Lewisian worlds varied enough with respect to predicates occurring in the language—that the set of true sentences has a model in which any two worlds are such that some atomic predicate

is satisfied by a sequence of objects inhabiting one of the worlds, but not by the result of replacing each object in the sequence by its counterpart in the other world.) Start by enriching the Lewisian language with a standard name for each inhabitant of the actual world, and let S be the set of true sentences in the extended language. One can use the Completeness Theorem to generate a model M of S in which the domain consists of ‘=’-equivalence classes of terms and in which the assumption above is satisfied. If a_1 is an object in the individual-variable domain of M , let a_1^* be the set of individuals a_2 such that ‘ $C(x_1, x_2)$ ’ is satisfied by a_1 and a_2 . (I assume that ‘ C ’ is an equivalence relation, and therefore that $*$ partitions the individual-variable domain of M into equivalence classes.) If a_1 is an object in the individual-variable domain of M , let a_1^\dagger be $\langle z, \text{‘actual’} \rangle$ if the standard name of z is in the transitive closure of a^* , and $\langle a^*, \text{‘nonactual’} \rangle$ otherwise.

For each c in the world-variable domain of M , we construct an a -world w_c , as follows: the domain of w_c is the set of a^\dagger such that ‘ $I(x_1, w_1)$ ’ is satisfied by a and c in M ; the w_c extension of ‘ $\lceil P_j^n \rceil$ ’ is the set of sequences $\langle a_1^\dagger, \dots, a_n^\dagger \rangle$ such that ‘ $I(x_1, w_1) \wedge \dots \wedge I(x_n, w_1) \wedge P_j^n(x_1, \dots, x_n)$ ’ is satisfied by a_1, \dots, a_n, c in M . In addition we let individual constants receive their intended interpretations, and let ‘ α ’ be w_c , where c is the M -referent of ‘ α ’. Say that an a -world is admissible just in case it is a w_c for some c in the world-variable domain of M . If σ is a variable-assignment function in M , let σ^\dagger be such that $\sigma^\dagger(\lceil x_j \rceil)$ is $\sigma(\lceil x_j \rceil)^\dagger$, $\sigma^\dagger(\lceil w_j \rceil)$ is w_c if ‘ $\lceil x_j \rceil$ ’ occurs in ϕ (where c is such that ‘ $I(x_1, w_1)$ ’ is satisfied by $\sigma(\lceil x_j \rceil)$ and c in M), and $\sigma^\dagger(\lceil w_j \rceil)$ is $w_{\sigma(\lceil w_j \rceil)}$ if ‘ $\lceil x_j \rceil$ ’ does not occur in ϕ . An induction on the complexity of formulas shows that a Lewisian formula ϕ is satisfied by an assignment function σ in M just in case the actualist transformation of its normal form is satisfied by σ^\dagger when the world-variables range over admissible a -worlds and the individual-variables range over the union of the domains of admissible a -worlds.

Notes

¹See, for instance, Plantinga (1976) and chapter 3 of Lewis (1986). A recent ersatzist proposal is discussed in Fine (2002) and Sider (2002) (but embraced only by Sider). The sort of proposal that Fine and Sider discuss has a more ambitious objective than the proposal developed here, since it is intended to capture finer-grained distinctions amongst possibilities. It also relies on more substantial expressive and ontological resources. (It relies, in particular, on an infinitary language and an ontology of properties.)

²By emulating the benefits of quantification over sets of possibilia the proposal also allows us to go some way towards addressing the expressive limitations identified in chapter 7 of Williamson (2010). In particular, one can use the dot-notation introduced in section 2.4 to express a version of the claim that impossible objects form a set:

$$\exists\alpha\exists x\exists y(\alpha = \{x, y\} \wedge \diamond(\exists z(z = x)) \wedge \diamond(\exists z(z = y)) \wedge \neg\diamond(\exists z\exists v(z = x \wedge v = y))).$$

³See McMichael (1983).

⁴Lewis (1986), p. 194. (Lewis writes ‘genuine’ where I have substituted ‘Lewisian’.) Lewis takes the counterpart relation to be context-dependent, but here I shall treat it as constant for the sake of simplicity. Also for the sake of simplicity, I shall assume that the counterpart relation is an equivalence.

⁵It is worth emphasizing that by availing oneself of an ontology of Lewisian worlds and a Lewisian counterpart relation one does not immediately do any better. What one gets is a way of making clear when two Lewisian worlds are to be counted as linked, not a characterization of linking. A friend of the Lewisian ontology can, however, give a proper characterization of linking by making an additional *reductionist* claim; namely, that there is no more to possibility than the existence of the relevant representation.

There is, for example, no more to the fact that I might have had a sister than the fact that someone who is like me in certain respects has a sister.

⁶I assume, for simplicity, that the accessibility relation is trivial.

⁷More precisely, what one can show is this: provided there is a determinate fact of the matter about which sentences of L^\diamond are true, there is a notion of admissibility relative to which an a -worlds semantics assigns the right truth-value to every sentence in L^\diamond . (I assume that the first-order language L on which L^\diamond is based has a set-sized domain, and a name for every object in this domain.) *Proof:* Where S is the set of true sentences in L^\diamond , use Kripke's completeness theorem for modal languages to construct a Kripke-model for S in which the domain consists of equivalence-classes of terms. Then transform the Kripke-model into an a -world semantics by substituting the pair $\langle x, \text{'actual'} \rangle$ for each equivalence class in the domain of the actual world of the Kripke-semantics containing a name for x , and the pair $\langle x, \text{'nonactual'} \rangle$ for each object x in the domain of some non-actual world of the Kripke-semantics but not in the domain of the actual world of the Kripke-semantics. (The proof relies on the reasonable assumption that S is consistent relative to a normal logic. To avoid talking about accessibility relations, I have also assumed that S is consistent relative to S5. It is worth noting that—unless one assumes that the language is countable—the Completeness Theorem assumes a weak version of the Axiom of Choice, so the resulting characterization of admissibility will be non-constructive.)

⁸On the Principle of Recombination, see Lewis (1986) §1.8. For other approaches to grounding admissibility see Fine (1994) and Peacocke (1999), ch. 4.

⁹For a sparse conception of possible worlds, see Stalnaker's 'On what there isn't (but might have been)'. Stalnaker makes clear that he does not see sparseness as an obstacle for doing Kripke-semantics.

¹⁰See Plantinga (1976). For a critique of Plantinga, see Fine (1985).

¹¹Keep in mind that self-identity may be used as an existence predicate.

¹²If the accessibility relation is non-trivial, what one gets is the following instead:

$$[[\phi]_w]_{w'} \leftrightarrow ([\phi]_w \wedge \text{Accessible}(w', w)).$$

¹³ When I say that the transformation preserves truth-value what I mean is that there is a notion of a -world admissibility which guarantees that the actualist transformation of an arbitrary Lewisian sentence is true just in case the original Lewisian sentence would count as true from the perspective of the Lewisian. When I say that the transformation preserves inferential role, what I mean is that a Lewisian sentence ϕ follows from a set of Lewisian sentences Γ just in case ϕ 's transformation follows from the the transformations of sentences in Γ .

The result assumes that atomic predicates in the Lewisian's language other than 'I', 'C' and '=' (and any set-theoretic vocabulary) be *projectable*. For a monadic predicate P to be projectable is for it to be the case that a Lewisian world represents a possibility whereby something is P by containing an inhabitant who is P . (And similarly for many-place predicates.) Thus, 'Philosopher' is projectable because a Lewisian world represents a possibility whereby something is a philosopher by containing an inhabitant who is a philosopher; but 'inhabits a Lewisian world which is part of a pluriverse containing many Lewisian worlds' is not projectable because a Lewisian world does *not* represent a possibility whereby something inhabits a Lewisian world which is part of a pluriverse containing many Lewisian worlds by containing an inhabitant who inhabits a Lewisian world which is part of a pluriverse containing many Lewisian worlds.

¹⁴As Lewis observes, a feature of the 1968 translation is that ' $\forall x \Box (\exists y (x = y))$ ' turns out to be true. For this reason, I prefer a modification of the translation whereby

$(\Box\phi)^\beta$ is $\forall\beta_1(W(\beta_1) \rightarrow (\phi)^{\beta_1})$, $(\Diamond\phi)^\beta$ is $\exists\beta_1(W(\beta_1) \wedge (\phi)^{\beta_1})$ and $(P(x_1, \dots, x_n))^\beta$ is $\exists\gamma_1 \dots \exists\gamma_n(I(\gamma_1, \beta) \wedge C(\gamma_1, x_1) \wedge \dots \wedge I(\gamma_n, \beta) \wedge C(\gamma_n, x_n) \wedge P(\gamma_1, \dots, \gamma_n))$ (for P atomic). The modified translation delivers the same truth-values as a version of Kripke-semantics in which atomic formulas (including identity-statements) can only be satisfied at a world by objects that exist at that world.

¹⁵For a representative textbook, see Heim and Kratzer (1998), ch. 12.

¹⁶For discussion of related issues, see Lewis (1986) §3.1. It is worth noting that Lewis's critique is not meant to apply to the sort of project that is the focus of the present essay. Here is Lewis: "When I complain, as I shall, that there are various ways for different possibilities to get conflated in their linguistic descriptions, that may be harmless when we want to use ersatz *possibilia* to characterize the content of thought for a subject who has no way to distinguish the conflated possibilities in his perception and conduct." (p. 144, footnote).

¹⁷More precisely, there might be different Lewisian worlds such that every inhabitant of the one world is a counterpart of an inhabitant of the other, and every predicate in the language which is projectable in the sense of footnote 13 is satisfied by inhabitants of one world just in case it is satisfied by the counterparts of those individuals at the other. The distinct possibilities represented by such Lewisian worlds would both be compatible with the less-specific possibility represented by an a -world in which the behavior of the predicates mirrors the behavior of the predicates at the Lewisian worlds.

¹⁸Any set of worldbound individuals containing an inhabitant of one of the Lewisian worlds described in footnote 17 but not its counterpart at the other corresponds to a Lewisian property with no actualist surrogate.

¹⁹For their many helpful comments, I am grateful to Elizabeth Barnes, Ross Cameron, Michael Glanzberg, Ephraim Glick, Caspar Hare, Sally Haslanger, Allen Hazen, Michael

Rescorla, Graham Priest, Peter Simons, Bob Stalnaker, Robbie Williams, Steve Yablo, audiences at the University of Leeds, the University of Melbourne and MIT's Work in Progress Seminar, and especially to two helpful referees for *Critica*.

²⁰An occurrence of a non-modal variable in an undotted formula is *free* iff it is not bound by a quantifier; an occurrence of a non-modal variable in an undotted formula is *externally bounded* iff it is bound by a quantifier which is not within the scope of '[...]'.

²¹I assume that D is a *set* for the sake of simplicity. The assumption can be avoided by characterizing the notion of an a -world in second-order terms. This can be done by employing the technique in Rayo and Uzquiano (1999).

²²These three replacements are needed to secure the base clause in the induction below.

References

- Fine, K. (1985) “Plantinga on the Reduction of Possibilist Discourse.” In Tomberlin and van Inwagen (1985), 145-86. Reprinted in Fine (2005), 176-213.
- Fine, K. (1994) “Essence and Modality,” *Philosophical Perspectives* 8, 1–16.
- Fine, K. (2002) “The Problem of Possibilia.” In Zimmerman (2002), 161-179. Reprinted in Fine (2005), 214-231.
- Fine, K. (2005) *Modality and Tense*, Oxford University press, Oxford.
- Heim, I., and A. Kratzer (1998) *Semantics in Generative Grammar*, Blackwell Textbooks in Linguistics, Blackwell, Oxford.
- Lewis, D. (1968) “Counterpart Theory and Quantified Modal Logic,” *The Journal of Philosophy* 65, 113–126. Reprinted in Lewis (1983).
- Lewis, D. (1970) “General Semantics,” *Synthese* 22, 18–67. Reprinted in Lewis (1983).
- Lewis, D. (1983) *Philosophical Papers, Volume I*, Oxford.
- Lewis, D. (1986) *On the Plurality of Worlds*, Blackwell, Oxford and New York.
- Loux, M., ed. (1979) *The Possible and the Actual*, Cornell University Press, Ithaca, NY.
- McMichael, A. (1983) “A Problem for Actualism about Possible Worlds,” *The Philosophical Review* 92, 49–66.
- Melia, J. (2001) “Reducing Possibilities to Language,” *Analysis* 61, 19–29.
- Peacocke, C. (1999) *Being Known*, Clarendon Press, Oxford.
- Plantinga, A. (1976) “Actualism and Possible Worlds,” *Theoria* 42, 139–60. Reprinted in Loux (1979).
- Rayo, A., ed. (forthcoming) *The Construction of Logical Space*, Oxford University Press, Oxford.
- Rayo, A., and G. Uzquiano (1999) “Toward a Theory of Second-Order Consequence,” *The Notre Dame Journal of Formal Logic* 40, 315–325.

- Roy, T. (1995) “On Defence of Linguistic Ersatzism,” *Philosophical Studies* 80, 217–42.
- Sider, T. (2002) “The Ersatz Pluriverse,” *Journal of Philosophy* 99, 279–315.
- Tomberlin, J., and P. van Inwagen, eds. (1985) *Alvin Plantinga*, Reidel Profiles Series 5, Reidel, Dordrecht.
- Williamson, T. (2010) “Necessitism, Contingentism, and Plural Quantification,” *Mind* 119, 657–748.
- Zimmerman, D. (2002) *Handbook of Metaphysics*, Oxford University Press, Oxford.