

{ 1 }

**On What There Isn't
(But Might Have Been)**

The problem of ontology, Quine told us in his classic essay “On what there is,”¹ can be put in a simple question, “what is there?” and answered in a word: “everything.” My question should be equally simple, and its answer should follow from Quine’s: there is nothing that isn’t. But of course as Quine went on to say, the problem gets harder when one tries to be more specific about what there is and what there isn’t. Quine’s concern was mainly with the problem of expressing disagreement about ontology—if I believe there are more things in Heaven and Earth than are dreamt of in your philosophy, how can you talk about what it is that I believe in, but you do not? But even when we agree about what there is, we may want to acknowledge that things might have been different—not just that things might have been differently arranged but that there might have been different things than there actually are. If we ask not just “what is there” but “what might there have been,” the answer “everything” does not seem sufficiently inclusive. But what else is there that might be included?

The problem is sufficiently daunting to have driven many philosophers, in different ways, to deny there could have been anything other than what in fact exists, or that anything that exists could have failed to exist. (Three examples of philosophers who develop this idea in very different ways: Wittgenstein of the *Tractatus*, David

¹ Quine 1948.

Lewis, and Timothy Williamson.) Others have hypothesized actual surrogates for the nonexistent things—individual essences that are themselves necessary existents and that correspond one-to-one with all the “things” (as we are inclined to put it) that might exist.² Still others think that because taking modality seriously forces us to such metaphysical extravagance, we should reject modal discourse as anything more than a *façon de parler*. But I think modal concepts are central to our understanding of the world—the actual world—and that understanding them should not require extravagant metaphysical commitments. My aim in this book is to sketch a framework that allows us to avoid extravagant metaphysical commitments and that is also compatible with intuitively natural beliefs about the way things might have been.

There are some philosophers who want to take modality seriously, and seek a theoretical account of modal discourse, but who think that we cannot take possible-worlds semantics, as an account of modality, seriously without making extravagant metaphysical commitments. Christopher Peacocke, for example, holds that “it is an unstable, indeed incoherent, position to think that you can at the same time use the Kripke-style semantics in the metalanguage to give absolute truth-conditions for modal sentences, count . . . [the proposition that there could have been something that doesn’t actually exist] as true, yet avoid commitment to the existence of nonactual objects.”³ But I want to defend the metaphysical innocence not only of modal concepts but also of a theoretical account of them in terms of possible worlds. Whether my construal of possible-worlds semantics counts as a realistic one or not is open to debate, and I will concede that on one of the several ways of construing the term “possible world,” the possible worlds posited by these semantic models are artifacts of the model and not entities

²This is the response to the problem developed and defended by Alvin Plantinga. See the papers collected in Plantinga 2003.

³Peacocke 2002, 121.

whose existence is affirmed. But I will argue that on another way of understanding the term, we can affirm the existence of possible worlds, as well as the claim that the semantic theory provides “absolute truth conditions for modal sentences” and “avoids commitment to the existence of nonactual objects.”

Here is my plan for this chapter: I will start, in section 1, with some preliminary methodological remarks—about the aim and value of reduction in philosophical analysis, about thinking of the evaluation of philosophical theses in terms of costs and benefits, and about the contrast between realistic and anti-realistic accounts of a philosophical theory. In section 2, I will say what I take possible worlds to be, and what, from the perspective of this account of possible worlds, the problem is about merely possible individuals. Possible worlds, on the account I want to defend, are (to a first approximation) *properties*, and the main point I want to make in this section is that properties (and so possible worlds) are not *representations*. In section 3, I take an extended look at some examples of properties that are simpler and easier to think about than possible worlds but that share some of the features of possible worlds, construed as properties. In this section and section 4, I will use the analogy I develop to motivate what I hope is a metaphysically innocent account of the domains of other possible worlds.

The view I will be defending is committed to making sense of the contingent existence of individuals and properties, of propositions, and even of possible worlds themselves. I will conclude, in section 4, by sketching a problem that an account of this kind faces, a problem that I will respond to in chapter 2.

1. Methodological Preliminaries

According to John Divers in his useful survey of the range of alternative philosophical accounts of possible worlds, “the primary question of conceptual application of the species of AR [actualist

realism] is whether any affords a thoroughly non-modal analysis of the family of modal and intensional concepts.”⁴ Divers acknowledges that “the proponents of AR typically do not claim that the favored version of AR affords thoroughly nonmodal analysis of the modal concepts,”⁵ but he seems to assume that it would be a benefit (in the cost-benefit evaluation of the general view) if it did provide such an analysis. But my view is that if an account of modality were to meet this condition, that would be a sure sign that it was on the wrong track. Necessity and possibility are fundamental concepts, like truth and existence. What would you say to a philosopher who was seeking a thoroughly nonexistential analysis of quantificational concepts, or a thoroughly non-alethic analysis of truth, and related concepts? It is not that philosophers have not proposed such analyses (substitutional quantification, truth as warranted assertability or as what ideal believers will believe at the end of inquiry, for example). But even if an analysis of this kind were to be extensionally correct, at least according to someone’s philosophical theory, it would only blur the distinction between semantic analysis and a substantive metaphysical thesis about what exists or what is true. Consider the nominalist who defines existence as *having spatio-temporal location*. Platonists will agree that if that is what you mean by “exist,” then numbers, sets, and properties do not exist. They will need to find alternative means of describing their ontological beliefs.

I do not want to suggest that one can distinguish, on some pre-theoretical a priori ground, which concepts are fit subjects for some kind of reductive analysis. It may be a contentious philosophical question, not only how to answer substantive questions but also which questions are substantive and which are semantic. So, for example, I am inclined to think not only that what is *actual* coincides with what exists but that this is because “actual” just means

⁴Divers 2002, 181.

⁵Ibid., 301n.

(more or less) real, or existent. The modal realist disagrees, and he might complain that by understanding “actual” in this way, I am blurring the line between metaphysical and terminological questions. I agree that my disagreement with the modal realist is a mix of semantic disagreement and disagreement about what there is in the world, and that to be clear, it is important to try to separate semantic from substantive questions, but it is not always easy to do so.⁶ I will discuss this issue in more detail and make some claims about how the two kinds of issues should be separated in chapter 4.

I have alluded to the cost-benefit, reflective equilibrium methodology that Lewis articulated and made fashionable, but I have my reservations about this way of thinking about the way philosophical alternatives are evaluated. This picture may be fine if it is taken simply as a reminder that in philosophy, as in science, political theory, or any other enterprise, everything is potentially criticizable; there are no absolute unquestionable dogmas. One should add that even judgments about what is a cost and what a benefit might be a proper subject of debate. But beyond the bland truism, the reflective equilibrium method does not offer much guidance. Even though anything might be epistemically relevant to anything else, one important task, in deciding between alternative philosophical views, is to isolate considerations of different kinds. There may be no absolutely neutral conceptual standpoint, but it is a virtue of a theoretical account of some concept or family of concepts (a benefit in the cost-benefit analysis) if it is able to fashion some tools that manage to remain neutral on issues in dispute—to provide resources to formulate alternative substantive views as coherently as possible. A more neutral account (of truth, existence, properties and relations, modality) may seem disappointing (it would be nice to have an account of truth that gave us a lot of information about

⁶For an excellent discussion of some of the problems of sorting out semantic from metaphysical questions when discussing fundamental ontology, see Lewis 1990.

what is true), but I think we should be suspicious of an account of modality that tells us too much about what there is or about what there might have been.

Consider this parody of the cost-benefit methodology, run amok: X says, “I have a beautiful, austere, and crystal clear theory of properties: they are just sets—no more and no less. The relation between a property and its exemplifications is just the relation between a set and its members.” Y responds: “It is a beautiful theory, I agree, but unfortunately it is false—there are many obvious counterexamples. We don’t need to consider exotic examples like renates and cordates. Consider any two uninstantiated properties like being a talking donkey and being a philosophizing cat [two of David Lewis’s favorite examples]. It follows from your view that these two properties are one, which is obviously wrong.” X replies: “Some who like my theory—perhaps a Quinean—would reply by rejecting your intuition that there are distinct properties here. But I am a commonsensical chap [as David Lewis liked to describe himself], and I agree with you that the properties you have described are distinct. Nevertheless, I am reluctant to give up my beautiful theory, since its benefits are great. I prefer to give up instead the belief that there are no talking donkeys and no philosophizing cats.”

There is much more that X needs to say, for example, about how these ontological hypotheses are to be reconciled with apparent evidence to the contrary, but however X goes on, I think most of us will find this response suspicious, not just because the benefit of the beautiful theory is outweighed by the cost of the ugly facts, but because there is something suspect about using this kind of theoretical virtue to reach this kind of conclusion about what there is. (This is a parody, but it can be argued that my story of X and Y is just an uncharitable spin on the kind of consideration that actually motivated Lewis’s modal realism. Lewis does hold this theory of properties, and getting the identity condition for properties right is a prime motivation for the hypothesis of a plurality of worlds. And

it is the identification of properties with sets that rules out, for him, an actualist account of possible worlds.)

One final methodological remark before getting down to business: it is common to distinguish between “ontologically serious” applications of modal semantics and purely instrumental uses; the latter includes mathematical uses (for example, the construction of models to show the satisfiability, in a technical sense, of certain sets of sentences of a formal language) and heuristic uses that treat possible worlds discourse as “a vivid shorthand for sentences containing modal operators.”⁷ It is often suggested that if we are to take possible-worlds semantics to be a theory that contributes to a project of philosophical explanation of modality, then we must specify a particular model—the intended model of metaphysical possibility. Jon Barwise and John Perry wrote, in criticizing possible-worlds semantics:

If the model-theoretic structures of possible worlds semantics, the ones that include a set of all possible worlds, are supposed to be a model of something, say super-reality, under some correspondence or other, then there ought to be *one* that is an intended or standard model, *the* one that really corresponds to super-reality.⁸

I think this contrast is oversimplified. Taking possible-worlds semantics seriously as an explanatory account need not require the belief that there is one intended model any more than taking quantification theory, and its semantics, seriously requires the belief that *the* intended interpretation is in terms of a single domain of absolutely everything. There is controversy about whether it makes sense to quantify over absolutely everything, but whether it does or not, I think all should agree that we can take quantification theory as more than a mathematical tool or a heuristic device even while

⁷ Sider 2002, 280.

⁸ Barwise and Perry 1985, 120.

rejecting the idea of an absolute domain. And whether or not it makes sense to talk of an absolute, context-independent domain of all possible worlds, it is useful to separate the project of clarifying the framework for doing modal metaphysics from the project of saying, within that framework, what is really necessary and possible. (Just as Quine distinguished the project of getting clear about what ontological commitment is and how it is to be represented from the project of stating what one's ontological commitments are.) I agree that if we are to take possible-worlds semantics seriously, we must say something about the *kind* of thing that a possible world is and justify the claim that it is reasonable to think that there are such things as possible worlds. But my aim will be to vindicate the possible-worlds theory while making minimal commitments about substantive metaphysical questions, for example, about whether there are things, or properties, that exist only contingently, whether there are individual essences that are irreducible to qualitative properties, whether there could be distinct but qualitatively indiscernible worlds. In the balance of costs and benefits, I give positive weight to this kind of neutrality.

2. What Are Possible Worlds?

So what, on my view, are possible worlds (in the sense in which it is reasonable to say that there is a plurality of such things)? I take them to be *properties*—ways a world might be. Of course this leaves a lot open, since there are many different accounts of what properties are, but I take the significance of the categorization of possible states of the world as properties to be that it implies at least these two things. First, a possible world is the kind of thing that is, or can be, instantiated or exemplified. An actualist needs the distinction between existing and being exemplified in order to be able explain the sense in which a merely possible world exists (a property the world might have had exists) and the sense in which it does not (no

world that is that way exists). But second—and this is the point I want to emphasize—if possible worlds are properties, they are not *representations*—not mental or linguistic entities. So the account of possible worlds I will defend rejects what Lewis calls linguistic ersatzism, as well as the other forms of ersatzism that Lewis considers, all of which treat possible worlds as representations. The significance of this point is that possible worlds are not the kind of thing that faces a problem of intentionality. About a representation (a name, a predicate, a picture, a scale model, a sentence) we can intelligibly ask, what is it about it in virtue of which it represents what it represents? (The following, for example, is a perfectly reasonable question: “What is it about the inscription or vocable ‘triangular,’ as it is used in a certain linguistic community, that makes it a word for the property of being triangular, rather than for the property of being square?”) David Lewis, taking the actualist to be giving some kind of representational account of possible worlds, asks: What is it about a world in which there are talking donkeys that makes it a world in which there are talking donkeys, rather than a world in which there are philosophizing cats? But if possible worlds are properties, this is like the question, what is it about the property of being triangular that makes it *that* property rather than the property of being square? This I take to be an unintelligible question.

The assumption that possible worlds are representations is widespread. Brad Skow, for example, gives voice to the following remark: “Possible worlds are *representations*. All theories of possible worlds agree about this.”⁹ I can say with confidence that the second of these two statements is false, since my own account of possible worlds rejects the first statement, and I don’t think I am alone in rejecting the idea that possible states of the world are representations.

⁹Skow 2008, 103. It is not clear that Skow endorses this remark, since it is attributed to a critic of a point that he is making. But his response to the critic does not reject the claim.

But Skow's claim has an appearance of plausibility that I think rests on the fact that while most philosophers reject David Lewis's modal realism, most have accepted his way of framing the debate about possible worlds. A theorist of possible worlds (many follow Lewis in assuming) is either a modal realist or a believer in ersatz substitutes for the possible worlds that the modal realist believes in. But is it appropriate to describe a property of individuals as an ersatz individual? Is, for example, the property of being a king an ersatz king? Does the property of being a king *represent* something as being a king? What does it represent as being a king? One can *use* properties to represent: the colors, red and blue, for example, are used to represent Republican and Democratic voting patterns, respectively. (It is not just the predicates "red" and "blue" that do the representing, as in the expressions "red state" and "blue state." One also uses the colors themselves, on maps, to do the representing.) Properties such as color properties might be used to represent themselves, as when one colors a part of a scale model of something red in order to represent that the corresponding part of the thing being modeled is red. It might be perverse, but one could also use different colors for this purpose—red to represent blue. So a color might represent a color; nevertheless, the relation between a property and what exemplifies it is not itself a representational relation. By painting the wall blue one does not thereby represent the wall as being blue, nor does the wall itself represent itself as being blue simply by being blue.

Why does it matter that the relation between a property and its exemplifications is not a representational relation? It matters because if one thinks of this relation, or the relation between a proposition and the world in virtue of which the proposition is true or false, as a special case of a representational relation—a particularly intimate one—then one creates the illusion of a problem. When properties or other things are used to represent, one explains the representation relation in terms of the intentions of the users. But

properties and propositions are thought of as mind-independent objects that are *intrinsic* representations: they represent without our help; how do they do it? Jeffrey King characterizes the classical view of propositions as the view that propositions are “eternal abstract entities that *by their very nature and independently of all minds and languages* represent the world as being a certain way and so have truth conditions.”¹⁰ Nothing could do this, which is why King rejects the classical view. One central problem for any theory of propositions, King argues, is to explain their capacity to “represent how the world is,”¹¹ to explain “how propositions have truth conditions.”¹² But on the account of propositions I will defend, propositions *are* truth conditions. What needs to be explained is how things that express propositions—that represent the world as being some way—can express the propositions that they express. But in giving a theory of possible worlds and propositions themselves, we are not addressing this question.

If one tries to say just a little about what properties, in general, are, it becomes clear the extent to which, in classifying possible worlds as properties, we are not explaining modal notions in terms of something more basic. I take the notions of property and relation to be themselves modal notions. Properties are to be understood in terms of what it would be for them to be exemplified, which means we understand what a particular property is in terms of a range of possible situations in which it would be exemplified. But possible situations, we are saying, are themselves properties—ways a situation, or a world, might be. It is not reduction but regimentation that the possible-worlds framework provides—a procedure for representing modal discourse, using primitive modal notions, in a way that helps reveal its structure.

¹⁰ This quotation is from a handout of a talk, but the general view is expounded in King 2007.

¹¹ King 2007, 3–4.

¹² *Ibid.*, 58.

What are possible worlds properties *of*? They are properties of the total universe. One may question whether there is such a thing as the total universe to be what has these properties,¹³ but I will assume that one can intelligibly speak of a universe that is (in the sense of “exemplifies”) a way things might be. (If there is no such entity, perhaps we can speak of possible states of the world as being exemplified, or not, but not *by* anything.)

Possible worlds, on the actualist construal, are usually said to be *complete* or *maximally specific* in some sense. The idea seems to be that they are properties that are as specific as the things that might exemplify them, but it is not easy to say exactly what this means. There are properties that are defined in terms of their exemplifications (like the property of being identical to Osama bin Laden), but of course there is only one possible state of the total universe that is exemplified, and so the others cannot be defined in terms of the universes that exemplify them. What is it for a property that is not so defined to be as specific as what would exemplify it?

Rather than trying to explain what this might mean, I will define maximality in a different way: a possible state of the world must be maximal in the sense that it decides every proposition. But propositions (in the possible-worlds theory) are identified with sets of possible worlds (or equivalently, functions from possible worlds to truth values), and on this account, the claim that possible worlds are maximal puts no constraints on the character of the worlds. If one explains propositions independently of possible worlds (perhaps the propositions are all the properties that are either exemplified or not by the total universe), then we would have an account of what it is for a world to be maximal, at least relative to the domain of propositions. But one might be suspicious of an absolutely complete domain of all propositions. What matters for the applications of possible-worlds semantics is that the possible states of the

¹³Robert Adams raised this question, in correspondence, in 1974 in response to my original paper on possible worlds.

world be maximal with respect to all questions that are of concern in the application at hand. I prefer to think of the worlds not as the *points* in logical space but as the cells of a relatively fine-grained partition of logical space—a partition that makes all of the distinctions we need. If the partition is fine enough for the purposes at hand, then we can understand the propositions as sets of the partition cells. We do not thereby foreclose the possibility that in some other context, one might cut the space up more finely. The question of whether there is an absolutely finest partition or whether the space is best understood as an atomless algebra, rather than a set of points, is a controversial metaphysical question we can set aside: taking possible-worlds semantics seriously does not require a commitment to an interpretation in which the possible worlds are absolutely specific, in some metaphysical sense.¹⁴

I have been suggesting that possible-worlds semantics need assume only that possible states of the world are as specific as is needed for the purposes at hand, but where the purposes at hand involve understanding talk about what might exist, but does not, we have a problem. The problem is that it seems that in this case, our purposes may require that we carve up logical space more finely than we have resources for. Since we are actualists, we have only the resources that the actual world provides for representing possibilities. We can represent a purely existential possibility (for example, that there is a purple cow) if we can understand the property of being a world in which there is a purple cow.¹⁵ But we understand a property in terms of what it would be for it to be instantiated, and

¹⁴ Cf. Saul Kripke: a “counterfactual situation’ could be thought of as a mini-world or a ministate, restricted to features of the world relevant to the problem at hand” (1980, 18). See also Stalnaker 1986, where I distinguished internal from metaphysical completeness (unconsciously echoing Hilary Putnam’s terminology for two kinds of realism) and argued that possible-worlds semantics was committed only to the former.

¹⁵ It is controversial whether we can understand the possibility of a purple cow; Peter van Inwagen (1998) has suggested that such a beast may be impossible.

this general property would be instantiated only if a more specific property, being a world in which a particular x exists, where x is a purple cow. We have a problem if we want to say that while there might have been purple cows, there are no particular things that might have been purple cows.

I am going to approach this problem indirectly by looking at some examples of properties that are exemplified by things that are less grand than total universes but that illustrate some of the problematic features of such properties.

3. Containment Properties

I want to consider a range of properties that an envelope (for example, one of those large envelopes that are recycled in the campus mail) might have, properties that concern what is inside the envelope. Start with these three examples:

- (1) the property of containing three sheets of blank white paper, size A4
- (2) the property of containing a reprint of a critical notice, published in *Mind*, of David Lewis's *On the Plurality of Worlds*
- (3) the property of containing two photocopies of a handwritten letter from Ludwig Wittgenstein to Saul Kripke

Call these *generic containment properties*. One might also define *specific containment properties*, such as:

- (4) containing three *particular* sheets of blank, white paper, size A4 (in a particular order)
- (5) containing *this* reprint of a critical notice, published in *Mind*, of David Lewis's *On the Plurality of Worlds*

(4) and (5) might be construed in different ways. One might mean something like this by (4): *containing exactly a b and c, in that order* (which are in fact sheets of blank white paper, size A4); alternatively, one might mean *containing exactly three sheets of blank white paper, size A4, namely a b and c (in that order)*. (The difference is that on the first understanding, the three specific items might have the property in a possible world in which they are not blank white sheets of paper, while on the second they must be.) I will understand specific containment properties in the second way.

For every generic containment property that is instantiated, there is a corresponding specific containment property that is instantiated by the same thing. We could define this correspondence relation; it would be a second-order binary relation that relates two properties.

Now Saul Kripke (SK) was about twelve years old when Ludwig Wittgenstein (LW) died. We know that Kripke was a precocious child, but I am going to assume that these two philosophers never exchanged letters and thus that there are no photocopies (or things that might have been photocopies) of a handwritten letter from Wittgenstein to Kripke. If this is right, then (it seems reasonable to assume) there will be no specific properties corresponding to our third example of a generic containment property. Still, the general claim we made about correspondence still holds: (3) would be exemplified by the envelope only if a corresponding specific property were exemplified by that envelope.

In terms of this second-order correspondence relation, we can define a second-order property of properties—being a specific containment property corresponding to the generic containment property (3). This is an uninstantiated second-order property but a perfectly good property nonetheless.

One can define properties that are partly specific, partly generic, such as:

(6) containing a certain specific sheet of paper, plus two others (all blank, white, size A4)

One can define negative and disjunctive containment properties:

(7) *not* containing a reprint of a critical notice, published in *Mind*, of David Lewis's *On the Plurality of Worlds*

(8) containing either three sheets of blank white paper, size A4, or a certain specific reprint of a critical notice, published in *Mind*, of David Lewis's *On the Plurality of Worlds*

And one can define additional second-order relations, for example, a *permutation* relation that might hold between two specific containment properties: say that two specific containment properties are *permutations* of each other if they involve the same specific objects, but in a different order.

So suppose we had a specific property corresponding to (3)—the property of containing two specific photocopies of a letter from LW to SK. Then there would be a different specific property that permutes these two specific photocopies. Of course there are no specific properties of that kind, since there are no letters from LW to SK, but there is still no problem with the second-order relation.

This game could go on, but it is time to connect our exercise back to possible worlds. Before doing this, let me point out just one general fact about *negative* generic and specific properties. We noted that if a generic (positive) containment property is exemplified, then some corresponding specific property must be exemplified (by the same thing). So if there are no specific properties corresponding to a generic property, it follows that the generic property is un-instantiated. It is also the case that if a *negative* generic property is exemplified by something, then *every* corresponding negative specific property is exemplified by that thing. For example, if the envelope does not contain three blank white sheets of paper, size A4,

then for every a b and c the envelope does not contain three blank white sheets of paper, size A4, which are a b and c, in that order.

Since there is no possibly instantiated specific property of the form *containing x and y, which are photocopies of a handwritten letter from LW to SK*, there also are no negative specific properties of this kind. But we can see that since the envelope *actually* has the negative generic property, not containing two photocopies of a handwritten letter from LW to SK, if there were a specific negative property of this kind, it would be a property that it seems would be exemplified *in the actual world*. The point is (if I may put it in this loose way) that there are merely possible properties (such as the specific negative properties that would exist if there were any photocopies of handwritten letters from LW to SK) that are *actually* instantiated. We will return to this point.

Possible worlds, we said, are properties, and I hope the way that they are like containment properties is clear. Worlds, like envelopes, have things in them, and they might have contained things other than those they in fact contain. A possible (state of the) world is like a mixed generic/specific containment property. A counterfactual world might be specified as one containing a certain specific thing (Saul Kripke, for example) and a thing of a certain kind that is not any actual thing (for example, SK's seventh son). If the property of being a world containing SK and his seventh son were exemplified, then there would be a more specific property that would also be exemplified (a property of the form *containing SK and x, x being the seventh son of SK*). There are, it seems reasonable to believe, no persons who might have been SK's seventh son, or anything that might have been a person who was SK's seventh son, and so no properties of this form that might be exemplified. But we can still generalize, using second-order properties and relations, about properties of this kind that involve specific individuals.¹⁶

¹⁶In defending the contingency of properties and propositions, I am following Kit Fine, who has long argued for this. See the postscript in Prior and Fine 1977. The

4. Kripke's Dice: Let Me Count the Ways

I will use an example Kripke used, slightly modified, to expand on the point that one can use second-order properties and relations to talk about the possibility of specific properties that do not in fact exist. Kripke, to underline the modest and commonsensical character of his conception of possible worlds, and to help dissolve what he regarded as a pseudo-problem about the identification of individuals across possible worlds, asked us to consider a simple school probability exercise—a problem about a pair of dice and the thirty-six possible ways that they might have landed. Kripke tells us that “one of these miniworlds—the one that corresponds to the way the dice in fact come up—is the ‘actual world.’”¹⁷ So Kripke is assuming (fictively) that we are talking about an actual pair of dice, which he labels die A and die B. His main point was that it would be silly to ask, about the possibility in which A lands 6 and B 5, how we know that it is A, rather than B, that was the 6. But suppose our dice are a merely possible, generic pair of dice. There is a possible state of the world in which two such dice are thrown, one lands 6 and the other 5, but there is not a different state or property in

account I want to defend is also very close to the view developed in Adams 1981, but there are some differences between our views. Adams distinguishes more sharply than I would between qualitative properties and properties that are ontologically dependent on particular individuals, and he seems to be assuming that while properties of the latter kind may exist contingently, purely qualitative properties will be necessary existents. Perhaps there are some very abstract properties that exist necessarily, but I would argue that qualitative properties such as color and shape, like the things that exemplify them, exist only contingently. But my most important disagreement with Adams is that he holds that a metaphysical view that accepts the contingency of propositions and possible states of the world requires, or at least motivates, a serious modification of modal logic he describes as “metaphysically satisfying though formally inconvenient” (1981, 29). I will argue in chapter 4 that the formal inconvenience is both greater than Adams suggests, and unnecessary; the standard logic and formal semantics can, I think, be reconciled with the austere and satisfying metaphysics.

¹⁷ Kripke 1980, 16.

which two such dice are thrown, one lands 5 and the other 6. Of course we might add some detail to distinguish the two dice: we might, for example, stipulate that one has a scratch on the face with one spot, while the other does not. Then we could distinguish the possible state in which the one with the scratch lands 6 (the other 5) from the situation where the one with the scratch lands 5 (the other 6). But suppose there is no such detail. How do we distinguish the 6-5 situation from the different 5-6 situation? What do we mean when we call one of the dice A and the other B? (That one is A and the other B is not a fact about the possible states that can be used to distinguish them. The A and the B are our labels for describing the situation.) Perhaps we should say, in the generic case, that there are really just twenty-one possible states of the dice but that if one of them had been realized, then there would have been thirty-six possible states of the specific dice that would then have existed. We cannot distinguish specific die A from specific die B, from the perspective of the actual world, where neither exists, but we can talk, in a general way, about specific properties of the form *A lands 5, and B 6*, and we use the second-order permutation relation to talk about pairs of specific properties, both of this form, but with the A and B reversed. We need to talk, in a general way, about the possible specific properties in order to represent facts about the generic situation, such as the fact that in the possible situation in which one lands 5 and the other 6, it is also true of the one that landed 6 that it might have landed 5 while the one that landed 5 landed 6.

5. The Problem of Iterated Modality

Let me conclude this chapter by summarizing the general idea and pointing to a problem with it that I will develop and respond to in the next chapter.

Possible worlds are maximal properties that a universe might have or, equivalently, maximal propositions. Each such proposition

is maximal in the sense that for every (actual) proposition, either it or its contradictory is entailed by it. But a proposition might be maximal in this sense while failing to be fully specific, where a proposition is fully specific only if for every existential proposition that it entails, it also entails a singular proposition that is a witness to that existential proposition.

Just saying this is a step toward reconciling simple modal claims about merely possible things with actualism. We can give truth conditions for statements such as “Saul Kripke might have had seven sons” without committing ourselves to the existence of anything that might have been one of Saul Kripke’s seven sons. The statement is true if and only if there is a maximal proposition that entails the existential proposition that Saul Kripke had seven sons. This gives truth conditions for the possibility statement as a function of the inner proposition that is said to be possible, but we need our recursive semantics also to give the conditions under which the clause that expresses this inner proposition would be true *relative to a nonactual possible world*. It seems, however, that “Saul Kripke had seven sons” can be true, relative to a given possible world, only if seven singular propositions of the form “*x* is Saul Kripke’s son” are also true with respect to that world, which is to say: only if seven singular propositions of this form are entailed by the maximal proposition that is that possible state of the world. But if maximal propositions can fail to be maximally specific, this condition will not be met.

Alan McMichael made the problem clear and precise in a classic criticism of actualist possible-worlds semantics¹⁸ by focusing on the problem of iterated modal propositions—for example, that Saul Kripke might have had seven sons, the last of whom was a plumber who might instead have become a lawyer. McMichael proved, using premises that he argued the actualist should accept, that an

¹⁸ McMichael 1983.

iterated modal claim such as this couldn't be true (assuming that no actual thing could have been Saul Kripke's seventh son). But as McMichael emphasized, one of his premises was that the abstract objects we are calling propositions exist necessarily, a premise the account I am promoting rejects. In response to this way out of the problem, McMichael argues that "to acknowledge [that the possible worlds which exist from the point of view of one world are distinct from those that exist from the point of view of another] is to give up the extensionality of possible worlds semantics. . . . But if we have to give up the extensionality of the possible worlds approach, we might as well do without it."¹⁹

McMichael also considers a response to the problem that gives a nonrealistic interpretation of possible-worlds semantics—one that rejects "the idea of there *really being* nonactual possibles" but employs "a semantics which includes *so-called* nonactual possibles," and he raises some problems for nonrealistic semantics. The problems he raises for reconciling actualism with possible-worlds semantics are serious and on target. I think they can be overcome, but doing so will require that I be more explicit about the way I want to use and interpret possible-worlds semantics. I will try to do this in the next chapter.

¹⁹ Ibid., 55.