

# 3

---

## Intra-Worldly Semantics

### 3.1 Introduction

In this chapter, I'll consider semantic theories that take the form of what I'll call "intra-worldly semantics." In the paradigmatic case, such a semantic theory will be one in which we think of the meaning of a sentence as a structured proposition which ascribes properties or relations to objects, representing objects as instantiating properties or standing in relations. As a variant of worldly semantics, an intra-worldly semantics requires us to try to comprehend our knowledge of meaning sentences and predicates as asymmetrically dependent on our knowledge worldly entities and their relations. Specifically, an intra-worldly semantics thinks of this knowledge as knowledge of objects, properties, relations, and primitive modal relations between properties and relations. I will argue that we cannot comprehend our knowledge of meaning in this way. The core problem, which will be familiar in its basic form, is that the intra-worldly knowledge to which an intra-worldly semantics appeals can only be understood *in terms of* our knowledge of the semantical rules governing the correct use of predicates, but this knowledge of semantical rules, on an intra-worldly semantics, is understood as *depending on* our knowledge of primitive modal relations between properties and relations.

### 3.2 The New Non-Primitivist Actualism

Let us introduce intra-worldly semantics by returning to the issue of defining possible worlds. In the previous chapter, I considered two "primitivist" views of possible

worlds, the genuine realism of David Lewis and the “modest realism” of Robert Stalnaker. The key difference between these two views is that, whereas the Lewis takes possible worlds to be genuine other worlds of exactly the same sort as the actual one, Stalnaker takes possible worlds to be properties—ways for the world as a whole to be—that exist in the actual world. In my eyes, the greater similarity, however, is that neither try to define what possible worlds are. At least for the purposes of the semantic theory, they take possible worlds to be basic. This contrasts with views that try to give a constitutive account of possible worlds, saying what they are in terms of more basic entities such as states of affairs, propositions, or properties. These are non-primitive, or, views according to which propositions are constituted by more ontologically more basic entities. Often, non-primitivists align themselves with Stalnaker as “actualists,” but the difference between the primitivist and the non-primitivist is, actualists, on the other hand, such as Robert Adams (1974) and Alvin Plantinga (1976) of the old days, and Jeff King and Scott Soames of the new days, do.

I’ll focus here on the new non-primitivist actualism, which has emerged most prominently in the work of Jeff King (2007) and Scott Soames (2010). On this view, possible worlds are “big uninstantiated properties that are complex and have as parts other properties and relations,” (King 2007, 447). That is, they are complex properties that the actual world could have instantiated (and would have instantiated had things been otherwise), that have, as constituents, other simpler properties and relations. To get this conception of possible worlds into view, first consider the thought that there can be the properties that the world as a whole might instantiate—for instance, the property of being such that *a* is black. If *a* is black, the world as a whole is such that *a* is black. That is, it instantiates the property of being such that *a* is black.<sup>1</sup> Now, consider with the thought that properties can be joined together to form conjunctive properties. For instance, the property of being white and the property of being round can be conjoined to form the conjunctive property of being white and round. Since there are properties

---

<sup>1</sup>For Soames (2010), the relevant property is spoken of as the property of making the proposition that *a* is black true, but this may plausibly be taken to be the very same property as the property of being such that *a* is black.

that the world as a whole might instantiate and properties can be conjoined to form complex properties, there is, for instance, the property of being such that *a* is black, *b* is white, and *c* is gray. This is a property that the world could have instantiated and would have instantiated if it were actually such that *a* is black, *b* is white, and *c* is gray. This property is a way the world could have been, a possible world, or, as we should say if our terminology is not to be misleading, a “possible world-state.”

King and Soames, taking possible world-states to be complex properties that the world could have instantiated, understood in this way, both say that, in addition to possible world-states, there are impossible world-states, complex properties of the same sort that the world could not have instantiated. It’s not hard to see why this is a natural conclusion to draw on a view of this sort. As we’ve already said, properties can be conjoined to form conjunctive properties. We’ve also said that there is a property of being such that *a* is black, and there is a property of being such that *a* is white. So why shouldn’t we think that there is the conjunctive property of being such that *a* is black and such that *a* is white? As King says, “if you hold that properties exist, but deny that properties of a certain sort exist, you should provide a principled reason why properties of that sort don’t exist,” (448), and it’s hard to see what our principled reason could be here. If there is this property, then clearly there is a property of being such that *a* is black, *a* is white, *b* is white, and *c* is gray. If we think that possible worlds are properties of this sort, would could that property be other than an impossible world? King follows this thought through and claims that impossible worlds exist. It’s just that, unlike (non-actual) possible world-states, which are only *contingently* uninstantiated, impossible world-states are *necessarily* uninstantiated. Whereas possible world-states are ways that the world could have been, impossible world-states are ways that the world could not have been.

It is important for primitivists about possible worlds such as Lewis and Stalnaker to maintain that there are no worlds that are impossible, for Lewis and Stalnaker want to understand what it is for some state of affairs to be impossible in terms of the fact that there is no world in which it obtains. If there are impossible worlds, no analysis of this

sort can be maintained. However, King and Soames do not wish to provide a reductive analysis of modal notions in terms of possible worlds. According to Soames, one of Lewis's main errors consists in his "thinking that modal notions can be analyzed away, rather than taken as primitive," (2010, 110). If we take modal relations that obtain between properties as primitive, then we can demarcate the set of possible worlds from the set of impossible ones in terms of whether or not it's possible for the world to be instantiated. Cashing this out, if we think possible worlds are complex properties that have simpler properties as their constituents, we can demarcate the possible world-states from the impossible ones by specifying whether simple properties that cannot possibly be co-instantiated by some object would have to be co-instantiated by some object in order for a world-state to be instantiated by the world as a whole. By taking the compatibility or incompatibility of simple properties to be explanatorily prior to the possibility or impossibility of world-states in this way, we can demarcate the possible from the impossible world-states. So, since the property of being black is incompatible with the property of being white in the sense that the two properties cannot possibly be co-instantiated by a single object, the world cannot instantiate the property of being such that  $a$  is black and being such that  $a$  is white; instantiating this property would require a single object (namely,  $a$ ) to instantiate both the property of being black and the property of being white, and that is not possible. So, any world-state that includes the property of being such that  $a$  is white and being such that  $a$  is black is thus not a possible world-state.

The basic metaphysical idea here, underlying the new non-primitivist actualism, is, as Michael Jubien (2008) puts it, that "relations among properties are the real source of our intuitions about necessity and possibility," (104). Spelling out this account, Jubien tells us that we can provide the following explanation for the fact that, if something's a horse, then it must be an animal, or to use our example, that something black (all over) cannot be white (all over):

*It's that the two properties' intrinsic natures together guarantee it. We may therefore see this connection as an 'intrinsic relation'—one that holds between the two properties strictly as a result of their individual intrinsic natures. Here*

is the locus of the needed ‘modal oomph’. Differences between properties’ own intrinsic properties establish modal connections between them, (93).

So, on Jubien’s account, modal facts obtain in virtue of the “intrinsic natures” of properties these facts involve. For instance, the fact that, if *a* is black, then it can’t be the case that *a* is white is explained by the intrinsic natures of the properties of being black and being white which jointly establish the modal relation of incompatibility obtains between them. Similarly, Hale (2013) appeals to the “essences” of properties—*what it is to be* these properties—in order to explain modal facts. Hale tells us

No matter what entity or kind of entity is in question—be it a kind of object, or property, or relation, or function, or thing of some other kind—there will be some facts about what it is to be that entity (or an entity of that kind), and these will give rise to corresponding necessities.

So, on Hale’s preferred way of putting things, the fact if *a* is black, then it can’t be the case that *a* is white is explained what it is to be black and what it is to be white, where part of what it is to be black is to be non-white. However one prefers to spell out the details, the basic idea, in both of these accounts, is that the properties, in virtue of their natures or essences, bear certain modal relations to one another, and it is in virtue of these modal relations, and our grip on them, that we are able to say what states of the world are possible or not.

### 3.3 A Simple Intra-Worldly Semantics

Thinking of possible worlds in this way lends itself to a very different sort of semantic theory. If one explains which worlds are possible by appealing to relations of compatibility and incompatibility between properties, one can’t then turn around and analyze these relations in an extra-worldly framework. Rather, properties and their modal relations are taken as primitive, for the purposes of the semantic theory. We thus get a different kind of semantic theory—an intra-worldly semantics. Both King and Soames propose versions of an intra-worldly semantics. The details of their theories differ, but those differences don’t matter much for our purposes here. For King, names are

assigned objects as semantic values, 1-place predicates are assigned properties, and  $n$ -place predicates are assigned  $n$ -place relations. Sentences are assigned propositions, which are composed out of these semantic values and represent objects as instantiating properties or standing in relations.<sup>2</sup> Soames assigns names, 1-place predicates, and  $n$ -place predicates, *acts of cognizing* objects, properties, and relations as semantic values, and takes propositions to be acts of predicating cognized properties and relations of cognized objects, but everything basically works out the same. For simplicity, I will consider a view along the lines of that King proposes in which names are assigned objects and predicates are assigned properties and relations.

How does a semantic theory that appeals to intra-worldly facts of this sort work? At the most basic level, such a semantic theory assigns objects to names, properties to 1-place predicates, and  $n$ -place relations to  $n$ -place predicates. For our simple toy language, for example, the assignment of basic semantic values might be the following:

<b>[[a]]</b> = $a$	<b>[[is black]]</b> = the property of being black
<b>[[b]]</b> = $b$	<b>[[is gray]]</b> = the property of being gray
<b>[[c]]</b> = $c$	<b>[[is white]]</b> = the property of being white
<b>[[is darker than]]</b> = the relation of being darker than	
<b>[[is lighter than]]</b> = the relation of being lighter than	
<b>[[is the same color as]]</b> = the relation of being the same color as	

The semantic value of a sentence is a proposition. A sentence consisting in a name concatenated with a 1-place predicate expresses a proposition that represents the object that is the semantic value of that name as instantiating the property that is the semantic value of that predicate. A proposition that represents an object as instantiating a property is true if that object instantiates that property, false if that object does not instantiate that property. For instance, the sentence “ $a$  is gray” expresses the proposition

---

<sup>2</sup>King’s account of just what the relation is that binds together an object and a property in a proposition such that the proposition represents that object as instantiating that property is very complicated, and, in my view, utterly confused. For a criticism, see Simonelli (M.S.). For our purposes here, the details of the theory don’t matter.

that  $a$  is gray, which represents  $a$ , the semantic value of “ $a$ ,” as instantiating the property of being gray, the semantic value of “is gray,” and so is true just in case  $a$  is gray.

For logical operators, different intra-worldly semanticists have somewhat different proposals, and the details of any such proposal one particularly matter for our purposes here. One simple proposal, suggested by King (2009, 114 n. 28) and also adopted by Hanks (2011), is to take the semantic values of the logical operators to also be properties or relations, but ones that are ascribed to and instantiated by propositions rather than objects. So, we might assign to the logical operators the following properties and relations instantiatable by propositions:

[[**It’s not the case that**]] = the property of being false.

[[**and**]] = the relation of both being true

[[**or**]] = the relation of at least one being true

So, the proposition “It is not the case that  $a$  is gray” expresses the proposition that it is not the case that  $a$  is gray, which represents the proposition that  $a$  is gray, the semantic value of “ $a$  is gray,” as instantiating the property of being false, the semantic value of “It is not the case that,” and so true just in case it is not the case that  $a$  is gray. And so on. We thus have a simple compositional semantic theory for our toy language which appeals only to entities in the world—objects, properties, and relations—and no worlds as a whole.

Assigning names objects as semantic values and 1-place predicates properties as semantic values enables us, according to King, to “give a simple, direct explanation” (785) of facts such as those consisting of the “robust judgments about entailment relations between sentences” (784) that speakers make. Consider, for instance (F2), the fact that the sentence “ $a$  is black” and the sentence “ $b$  is white” jointly entail the sentence “ $a$  is darker than  $b$ .” An intra-worldly semantics of the sort proposed by King takes this fact to be explained in part by the natures of the following three entities in the world: the property of being black, the property of being white, and the relation of being darker than. These three entities stand in a certain relation: if something instantiates the property of being black, and something else instantiates the property of being white, then

the first thing stands in the relation of being darker than to the second thing. The fact that these three entities stand in this relation is not a semantic fact but a worldly fact; it is a fact consisting in three entities in the world (two properties and a relation) standing in a certain relation. Furthermore, the fact has a certain sort of modal robustness. It doesn't just happen to be the case that, if one thing is black and another thing is white, then the first thing is darker than the second thing. Rather, if one thing is black and another thing is gray, then the first thing *must* be darker than the second thing.

The modal robustness of this fact consisting in this relation obtaining between the property of being black, the property of being white, and the relation of being darker than, is not to be analyzed in terms of the fact that, in every possible world, every thing that is black is darker than every thing that is white. Rather, the order of explanation goes the other way around. There is no possible world in which one thing is black, another thing is white, and the first thing is not darker than the second thing in virtue of the fact that these properties and this relation, in virtue of being what they are, stand in this relation. That is, the modal robustness of the fact is grounded in the essences of the properties and relations it involves.<sup>3</sup> It follows from what it is for something to be black, what it is for something to be gray, and what it is for one thing to be darker than another that, if something is black, and something else is gray, the first thing is darker than the second. From the fact that the property of being black, the property of being white, and the relation of being darker than stand in the relation specified above, and the fact that the expressions "is black," "is white," and "is darker than," have these properties and relations as semantic values, it follows that the sentences "*a* is black" and "*b* is white" jointly entail the sentence "*a* is darker than *b*." It seems, then, that our simple intra-worldly semantics gives us a "simple, direct explanation" of the behavior we set out to explain, just as that consisting of "judgments about entailment relations between [atomic] sentences." Once again, however, things are not how they seem.

The intra-worldly semanticist appeals to properties at the base level of their semantic theory. Accordingly, it is these entities, and our grasp of them, that needs to be

---

<sup>3</sup>This is explicitly endorsed by Hale (2013), 146-147 and Jubien (2008), 92-94.



investigated if we are to investigate the foundation of an intra-worldly semantics. What are these entities and how do we grasp them?

### 3.4 Properties, Appealed to and Unaccounted for

Though many semanticists appeal to properties at some level in their semantic theory, most don't take them to be semantic primitives. If you ask a working semanticist what a property is, they're likely to answer this question in an extra-worldly framework, saying, for instance, that a property is a function mapping each possible world to a set of objects. This is, of course, a definition of a property as mathematically constructed from more primitive entities—objects and possible worlds, which are taken as basic from the point of view of the semantic theory. It is not a definition of properties that takes properties themselves as basic. In contrast to the extra-worldly semanticist, the intra-worldly semanticist does not want to think of properties as constructed from objects and possible worlds. This may well be because they rightly see that the particular objects that populate those possible worlds cannot be understood as the distinctive objects that they are independently of the properties they instantiate, and possible worlds cannot be understood as genuinely *possible* rather than *impossible* apart from thinking about the properties that would have to be co-instantiated by objects if some world were actual. So properties can be no less conceptually basic than objects and, indeed, must be more conceptually basic than possible worlds. But what *are* properties? They are, of course, the things that serve as the semantic values of predicates in the semantic theory. So, they are what speakers grasp in grasping the meaning of a predicate. But, once again, what *do* speakers grasp in grasping the meaning of a predicate? The intra-worldly semanticist, insofar as they are proposing an account of speakers' knowledge of meaning, ought to have *something* to say in response to this question, so they ought to be able to say something about the properties that are theorized to be the semantic values of predicates.

There are really two questions here. The first is what are properties *in general*? That is, what constitutes the ontological category of properties? This first question, I

don't think, is too difficult. Consider first Van Inwagen's (2006) identification of the notion of a property with the notion of a "thing that can be said of something," (27). According to Van Inwagen, properties, relations, and propositions all belong to the same fundamental type: they are all assertables. Properties are assertables with one place for a thing for the property to be asserted of, *n*-place relations are assertables with *n*-places for things for the relation to be asserted of, and propositions are assertables with no place for any things to fill—so, they are asserted not *of* anything, but *full-stop*. Things that can be asserted *of* things, whether just one or many, are *predicables* or *ascribables*. So properties, on this account, are things that can be predicated of or ascribed to things, said of them. This is one quite common construal of what properties are. Consider now Hale's (2013) construal of what properties are:

A property, on this account, just is a condition which things may or may not satisfy. [...] A property, we might say, is a way for things to be—perhaps a way some things are or could be, but perhaps a way nothing could be, (165-166).

This is, in fact, how Soames answers the question of what properties are, when pressed on the issue, and this is also quite a common construal. Bringing these two conceptions together I think there is a relatively philosophically unobjectionable specification of what, in general, properties are: they are ascribables and instantiateables. That is to say, they are things that can be both ascribed *to* objects and instantiated *by* objects. This is, indeed, how they are characterized in the opening paragraph of the *Stanford Encyclopedia* entry on properties (Orilla and Paolini Paoletti, 2020), and, of course, this is just what they need to be if they are to play the role in an intra-worldly semantic theory that they are supposed to play. If a proposition ascribes a certain property to a certain object, then that proposition is true just in case that object instantiates that property, false if that object doesn't instantiate that property.<sup>4</sup>

What else should we say of properties in general? Theorists like King and Soames, who appeal to properties at the base-level of their semantic theories, typically say very

---

<sup>4</sup>Note that, for Soames, a proposition is an act type that ascribes a property to an object only in a derivative sense that an agent that tokens that act in thought or speech ascribes that property to an object.

little.<sup>5</sup> However, we can find at least a few remarks in King's work about properties, further specifying what they are. King's clear that he thinks of them as "entities in the external world" (2018, 784) "existing quite independently of minds and languages" (2007, 450), that some of them are complex in the sense of being "made up of other properties and relations" (1998, 157; 2007, 447), and that they stand in relations of entailment to other properties (1998, 173 n8). So, though King never explicitly provides a general answer the question of what properties are, he is happy to specify that they have just the features he needs them to have in order for them to play just the role he wants them to play in the sort of explanation of semantic competence that he wants to give, serving as the contents of predicates. We might reasonably immediately wonder whether this sort of move is really justifiable. One might be inclined to liken it to postulation in other areas of scientific inquiry. For instance, of course, in the course of doing astronomy, we may posit an object with a certain gravitational force to make sense of the observed behavior of other celestial objects. One might think of properties as theoretical posits along the same lines.<sup>6</sup> However, the assumption implicit in a scientific postulation is that further scientific inquiry will eventually (or at least in principle *could*) tell us *what this thing is*, be it a gaseous planet, a large asteroid, or what have you. In this case, however, we don't even have so much as a gesture at what a proper account of these things would look like.

This brings us to our second question, which will prove more decisive for our project here: what are the *particular* properties that figure into the semantic theory? What, for

---

<sup>5</sup>One notable exception is Peter Hanks (2015, 2017) who, though he appeals to properties in much the way that King and Soames do, worries about these appeals in the context of a semantic theory and thinks that we need some account of how our grip on properties is achieved such that we understand the "standards of correctness" they provide for acts of predication. Though the framing of the issue here is quite different than that of Hanks, and so I do not explicitly engage with him here, the positive account that will be provided in chapters four and five will answer these concerns though they do so only at the cost of the explanatory use to which Hanks hopes to put properties in the context of his semantic theory.

<sup>6</sup>Indeed, Jubien (2008), whose work King (1998) cites on the nature of properties, explicitly makes this comparison, saying

Properties have central roles to play and we speak every day as if they are playing them right there. As long as there are no genuine problems with properties, we should welcome them as entirely sensible theoretical posits, (2008, 42-43).

instance, is the property of being black, appealed to in the intra-worldly semantics for our toy language? Or, to turn to a real language for a moment, what are the various properties that we would appeal to in providing an intra-worldly semantic theory for a natural language such as English? To get a sense of what an account of the properties that figure in an actual semantic theory would have to encompass, consider all the properties we would need to be able to appeal to in order for such a theory to be explanatory adequate. We need not just what one might regard as more fundamental properties, such as the property of being red, being green, being square, being round, and so on; being positively charged, negatively charged, and so on. We need *all* the properties corresponding to predicates of the language. So, we need the property of being a reptile, a bird, a cardinal, a penguin; the property of being a chair, being a table, a cup, a flask; the property of dancing, of swimming, of skiing, of smoking; the property of being milk, coffee, tea, chocolate cake; the property of being a novel, a novella, a screenplay; and many *many* more. As competent speakers of English we conceptually grasp this *vast network* of properties and their inter-relations, each one the content of a predicate of the language we speak. Without any story about what the constituents of this vast realm of properties are, or how speakers come to have a grip on them, an intra-worldly semantic theory, based on the assumption that speakers *do* have a grip on them, hangs in the air.

Now, perhaps theorists like King and Soames simply think that, whatever story is to be told about the structured space of properties on which we have a grip, there surely is *some* story to be told, and it's simply not their job to tell it. Who's job is it? Presumably, one would say that it's the job of the metaphysicians. King and Soames don't explicitly say this, since they don't explicitly address this question. Very few semantic theorists actually do. One exception is Herman Cappellan and Ernie Lepore. They do address this question, and this is the line they take. Since it seems like the sort of line the intra-worldly semanticist would presumably want to take, let us consider their explicit defense of it. To provide some context, Cappellen and Lepore endorse a "minimalist" semantic theory according to which, for instance, the sentence "A is

tall" expresses the proposition that *A* is tall and is true just in case *A* is tall. On their account, far as semantics is concerned, the analysis ends there. This is analogous, they claim, for how the sentence "*a* is red" expresses the proposition that *a* is red and is true just in case *a* is red or how the sentence "*A* dances" expresses the proposition that *A* dances and is true just in case *A* dances. Now, though their minimalist view is quite controversial in contemporary philosophy of language when applied to expressions like "tall," which are widely taken to be context-sensitive and so in need of some sort of further semantic analysis, the line they take with respect to "red" is supposed to be quite uncontroversial, and taking this line with respect to "dances" is supposed to be completely uncontroversial. Their argumentative strategy, then, is to first to argue why it is patently unreasonable to demand that the semanticist give an account of what it is to be *red* or what it is to *dance*, and then to extend this reasoning to the question of what it is to be tall. They assume their reader will regard the sort of objection they're arguing against, concerning the properties of being red and dancing, as "borderline silly" (156), but the arguments they give in response to this "silly" objection concerning "red" and "dances" are supposed to extend the philosophically respectable objection regarding "tall." My aim here, of course, is not to get into the debate within intra-worldly semantics concerning the correctness of minimalism or contextualism. Rather, it is to show that this "silly" objection, which would apply to both minimalism and contextualism, is really quite serious for the project of intra-worldly semantics as a whole.

The first point to make is that, insofar as these properties play an essential explanatory role in an intra-worldly semantic theory, these questions about *what they are* are certainly *legitimate* ones. Division of labor is fine, but, somewhere along the line, *someone's* got to do the labor. At the very least, even if the labor is not *actually* going to get done for each case, we ought to be sure that it's *possible* to actually do it, and, moreover, have some idea of how it *would be* done *were* we to actually do it. One might think, then, that Cappellen and Lepore would have some remarks about how this division of labor is supposed to work, how the respective disciplines of semantics and meta-

physics can ultimately be connected in a complete theory of linguistic and conceptual competence. Yet, when Cappellen and Lepore actually discuss the sort of metaphysical work that *would* give answers to these questions about the properties that figure in the sort of semantic theory they propose, they are clearly quite pessimistic regarding the possibility that it can be productively done or, indeed, done at all. Indeed, their rhetoric makes it quite clear that they don't take this sort of work seriously in the slightest. They tell us that, if you seriously ask these questions about what the properties that figure into the semantic theory are, "you'll regret it because it'll just turn into a rather large metaphysical mess (*not* a mess of our making, just the regular mess metaphysicians inevitably like to throw themselves into)" (158). *Good thing we're not metaphysicians!* is what we are presumably supposed to think as readers. The rhetoric here suggests not just that the sort of metaphysics that would seek to answer these questions is *hard*, like, say, topology is hard, but, rather, that it's *hopeless*. This dismissive rhetoric, however, is in tension with any real talk of a "division of labor" between semantics and metaphysics. Consider the following analogy. If one is proposing a biological theory, and, in response to some question concerning the things to which one appeals to in that theory, says that it's a matter of chemistry how those questions are to be answered, it's fine if one admits that chemistry is hard, perhaps even too hard for one to do oneself, but one should surely not say that chemistry is hopeless!

Now, of course, even if Cappellen and Lepore have this pessimistic attitude towards the metaphysical enterprise of providing an account of the properties that figure into an intra-worldly semantic theory, it need not be the case that every intra-worldly semanticist has this attitude. Perhaps intra-worldly theorists have or at least ought to have a more optimistic attitude towards the metaphysical enterprise that would provide in accounts of the properties that figure in the intra-worldly semantic theory they propose. I will eventually show that this enterprise is not hopeless at all, but in fact quite doable. However, what I'll now argue is that this enterprise is indeed hopeless, given the constraints put on it by the explanatory use to which properties are put in an intra-worldly semantic theory. There *is* an account of the properties that are grasped by

speakers of a natural language to be given. By the end of this chapter, we will say just what that account is, and, by the end of the dissertation, we will have developed the key tools needed to actually fill it in. The problem, however, is that this account of properties is simply not available to a proponent of an intra-worldly semantic theory, given the explanatory use to which properties are put in such a theory. Before actually getting to this account and explaining its incompatibility with an intra-worldly semantic theory, let us first go through the failure of various proposals that might be considered viable from the perspective of intra-worldly semantics.

### 3.5 The Problem of Defining Properties

Sticking with the strategy of examining semantic theories by considering how they fare with respect to a very simple toy language, let us consider how someone who is proposing an intra-worldly semantics for our toy language might try to specify the meanings of one of the predicates belonging to it—the predicate “black,” say. According to the intra-worldly semantic theory such an individual proposes, the predicate “black” expresses the property of being black. This property figures in at the base level of the semantic theory. How should one say what this property is? Let us consider some initial attempts.

Of course, it would be absolutely hopeless to try to define the property of being black as follows:

**[[black]]** = the property of being black =  
 the property that an object instantiates just in case that object is  
 black.

This, of course, says nothing. It is indeed the case that the property of being black is the property that an object instantiates just in case it is black, but one should not be tempted to hear this as a substantive definition of what the property of being black is. What such a statement is, really, is just a substitution instance of a grammatical remark expressing how property-talk in general is to be used. We can introduce

property-talk into a language into a language by specifying, among other things, that we are entitled to say “*a* instantiates the property of being *F*” just in case we are entitled to say “*a* is *F*.” Given this schema, it is of course, true of the property of being black that it is instantiated by an object just in case that object is black, and it is likewise true for any property, but that’s because it’s a completely empty description, one that says absolutely nothing about any particular property at all. Accordingly, such descriptions can’t function to specify what the particular properties that figure in the base level of our intra-worldly semantic theory actually are. They are, quite literally, without content. Since what we are supposed to be specifying is, of course, the semantic *content* of the predicate “black,” such a definition will not do.

A different way to try to specify what the property of being black is would be to try to do so along the following lines:

**[[black]]** = the property of being black =  
the property that all and only the black things instantiate.

Of course, unless we appeal to possible worlds, this isn’t going to work. After all, it might just so happen that all and only the black things are spherical, say, rather than cubical. In such a case, this definition and the definition of “spherical” would specify exactly the same property, since the exact same things would instantiate them. Clearly, however, the property of being black and the property of being spherical are not the same property. So this definition is blind to the difference between properties that just happen to be instantiated by the same things. Now, one might try to get around the issue with possible worlds by adding a primitive necessity operator, transforming the above definition into the following:

**[[black]]** = the property of being black =  
the property that, necessarily, all and only the black things instantiate.

However, trying to make any sense of the sort of necessity at play here simply brings us back to the first definition. Similarly, one could add to the end of this definition “in



virtue of being black," but, once again, this just brings us back to the first definition. So, any definition along these lines is not going to do.

Given the failures of the above two definitions, one might think that the answer to the question of what the property of being black is not a conceptual question at all, but, rather, an empirical question, something to be answered by empirical investigation into the nature of black things. Suppose, upon conducting such an investigation, we come to the following conclusion:

[[**black**]] = the property of being black =  
the property of completely absorbing light.

This does seem to give us a substantive specification of what the property of being black is. The obvious issue here, however, is that it doesn't give us a substantive specification of what the speakers of our toy language grasp in grasping the meaning of the predicate "black." The hypothetical speakers of our toy language, we may suppose, grasp the property of being black in grasping the meaning of the predicate "black" which belongs to their language without having any grip on the property of absorbing light. They have no words for the property of absorbing light and so we have no reason to think that they know what it is for something to absorb light at all. So, this specification of what the property of being black is cannot be a specification of what that speakers of our toy language grasp in grasping the meaning of the predicate "black."<sup>7</sup>

Now, perhaps a defender of the above definition will want to say that it's *really* this property on which they have a grip on, even though they don't have a grip on the essence of this property. Recall, however, the explanatory role of properties in the context of the semantic theories. Speakers are supposed to have *some* grip on the essences of these properties, since their grip on these essences is supposed to explain their grip on the modal relations that the properties bear to one another. Perhaps, if we

---

<sup>7</sup>Even if *we* want to say that the property of being black just is the property of absorbing light, now we're deploying a conceptual content in specifying the property that. Our concern is the property of being black *qua conceptual content*.

take this line we can explain these modal relations *ourselves*. For instance, if we take this line with respect to the property of being black, then, taking the same line with respect to the property of being, we might define it as follows:

[[white]] = the property of being white =  
the property of completely reflecting light.

It may well be an aspect of our conceptual framework that nothing can both completely absorb and completely reflect light, and so we can account for *our* grasp of the incompatibility of these properties in this way. Our task as semantic theorists however, is to give a grasp of the *speakers of the language for which we are constructing a semantic theory*. Any such definition, which is blind to the distinction between what is grasped by the speakers of the language for which we are giving an intra-worldly semantics and what we grasp ourselves as theorists, won't do.<sup>8</sup>

In response to the failure of these last three definitions, it might seem that the problem is the very idea that the property of being black can be captured in words. Perhaps, because of the particular sort of property that the property of being black is, a simple qualitative property, words will not do. If that's so, then perhaps the right way to specify the property of being black is to do so as follows:

[[black]] = the property of being black =  
the shade instantiated by the following object:



Here, one *shows* the reader the property of being black, by showing the reader an instance of it, rather than trying *say* what it is. Upon being shown something that visibly instantiates of the property of being black, the reader is supposed to know

<sup>8</sup>Hale (2007) suggests this sort of answer, saying "If what is in question is being red as a property of surfaces (as distinct from the property of light, or the property of sense-impressions), *being coloured* consists in reflecting light in the visible spectrum (roughly 390–750 nm), and the 'more' is that what is red reflects light of wavelengths of roughly 630–740 nm." In the context of semantics, this fails for the reason specified here. We will also see in Chapter Six (Section 6.5) that this even fails in the context of an attempt at scientific specification of the property of being red, understood as a theoretical property.

the specific property that figures in the semantic theory by simply being shown that property. Now, if one goes this route for the property of being black, then surely one would go the same route for the property of being gray:

[[**gray**]] = the property of being gray =  
the shade instantiated by the following object:



It might seem as if this is the way to go, at least for very simple properties such as color properties that don't seem like they can be constructed in any way from other properties.

As Wittgenstein pointed out, however, such "ostensive definitions" are not going to work. To see this, consider how one might attempt to ostensively define the relation of being darker than. This is supposed to be the semantic value of "is darker than" that figures in the semantic theory. Accordingly, we should be able to specify what it is as well. Attempting the same strategy here, however, yields obvious problems. Consider the attempt to try to define this relation as follows:

[[**darker than**]] = the relation of being darker than =  
the relation instantiated by the following two objects, with the one on the left occupying the first place of this relation and the one on the right occupying the second place in this relation:



The problem here, of course, is that there are indefinitely many relations instantiated by these two objects. For all that is said here, the demonstrated relation could be the relation of being to the left of, the relation of being the same shape as, or any one of a great number of relations. The demonstration itself does nothing to ensure that the reader takes it to be relation of being darker than that is demonstrated rather than any one of a number of other relations that these two objects stand in. Now, presumably,

you *did* take it to be the relation of being darker than that was demonstrated here rather than one of these other relations. But that's only because you read the text above the demonstration and, knowing what "darker than" means, you knew it was the relation of being darker than that was supposed to be demonstrated!<sup>9</sup>

At this point, one might think that the very idea of a *public* definition—something that articulates, in public language, what it is for something to be black or even publicly shows what it is—is problematic. The problem, one might think, is that the property of being black, as each of us grasp it, is essentially tied to a certain phenomenal quality that each of us is able to know, in our own case, but which cannot describe with public language nor can we even publicly demonstrate, since we cannot know that it is instantiated by the experience that someone else has when they look at something that we communally call "black." So, each of us knows, considering our own experience, what *we* mean when we say that something is black, since we each know the quality of the experience *we* have when we see something black. It is *this quality*, understood in terms of its intrinsic phenomenal character, that we principally mean when we speak of "blackness." The property of being black, as a property of objects in the world, might be understood, in a secondary sense, as the propensity of objects to produce experiences that have this quality. Once again, there can be no *public* expression guaranteed to pick out this quality, for it may well be this is not the quality that your experiences instantiate when you look at objects that we both call "black." It could be, for instance, that, when you look at an object that we both call "black," your experiences instantiate the quality that my experiences instantiate when I look at an object we both call "white." Still, though there is no public expression that we can be sure to pick out this quality, we can nevertheless each coin a term for ourselves that directly picks out this quality in terms of its intrinsic character.<sup>10</sup> Thus, I might coin the term "X" to pick out the phenomenal quality instantiated by my experience when I see something we call "black," thereby providing the following "private definition" which specifies what *I*, at least, mean by

<sup>9</sup>This is what Wittgenstein (1956, §257) speaks of as the "stage setting" required for a successful ostensive demonstration.

<sup>10</sup>This view is most explicitly spelled out by David Chalmers (2010), Chapter 8.

“black”:

[[**black**]]<sup>R.S.</sup> = the property of being black (as I grasp it) =

The propensity of objects to produce experiences with quality X.

I might likewise coin the term “Y” to pick out the phenomenal quality intantiated by my experience when I see something

[[**white**]]<sup>R.S.</sup> = the property of being white (as I grasp it) =

The propensity of objects to produce experiences with quality Y.

I will leave it as an exercise for the reader to define such semantic values for themselves and show, to their own satisfaction, that they cannot actually make determinate sense of what is purportedly expressed by their terms “X” and “Y,” at least, not without appealing to their grasp of what is expressed by the public expression “black” and “white,” thus bringing us right back to the problem with which we started.<sup>11</sup>

We have gone down quite a path in our attempt to say what the properties that figure into our simple intra-worldly semantic theory for our toy language actually are. By this point, one might have come to the conclusion that any attempt to provide a proper specification of what these properties are, at least for the basic ones like the property of being black, is bound to fail. But this would be too hasty. Let me now define these properties, thereby showing that they really can be defined.

### 3.6 The Way to Define Properties

There is, I think, a way that properties can be defined, though I take it that these definitions will always be relative to the rules of a linguistic practice which a particular structure. To see what I mean here, consider the properties grasped by the hypothetical speakers of our toy language. If we can imagine the speakers of our toy language as

---

<sup>11</sup>If the reader would like some help with this exercise, I point them, first, to Wittgenstein’s *Investigations* §28 – §39, §239 – §304, and, second, if help is still needed, to Stroud’s (2002) very helpful guide, “Wittgenstein’s ‘Treatment’ Of the Quest for ‘A Language Which Describes My Inner Experiences and Which Only I Myself Can Understand.’”

cognizers at all, we must suppose that there is some sense in which they grasp the property of being black, the property of being gray, the relation of being darker than, and on. Furthermore, I take it that, if we can imagine the speakers of our toy language as cognizers at all, then, for any property or relation that they grasp, there must be some specification of that property in the very terms which they themselves grasp it. Here is a proposal, based on this idea:

**[[black]]** = the property of being black =

The property such that, if something instantiates it, then, necessarily, it is darker than anything gray or white, nothing is darker than it, everything is either the same shade as it or lighter than it, and so on.

Here, we've supplemented the vocabulary of the speakers of our toy language with some additional logical vocabulary: words like "if," "then," "necessarily," "anything," "nothing," and "everything."<sup>12</sup> With this additional logical vocabulary, the speakers of our toy language are able to specify not only the objects that instantiate the property of being black (though, importantly, they can also do that for at least some of them) but also the modal relations that this property stands in to other properties and relations expressible in their language. The proposal is that the property of being black just is what is expressed by the above sentence of the logically enriched toy language, namely, a bit of metaphysical structure. By "metaphysical structure" I mean nothing but that structure which can aptly be expressed with a metaphysical "necessity" operator, the sort of structure that, once we introduce the toolkit of possible worlds, we'll be able to articulate by universally quantifying over them. According to Sellars, this metaphysical structure nothing but a codification of the exceptionless semantic norms governing the use the predicate "is black."

In the next chapter, I'll give an official account of these "semantic norms," and, in the chapter after that, I'll give an official account of how logical vocabulary can function

---

<sup>12</sup>Along with, among other things, the capacity for anaphoric reference. This is a bit of natural language that, though certainly essential for a full account of conceptual contents expressed by singular terms, we'll end up ignoring here for simplicity as our focus will be on properties, the conceptual contents expressed by predicates.

to make the semantic norms governing the use of predicates explicit. Here, however, I want to consider what the conception of the property of being black that is yielded by this final definition is, and why it is unavailable to our intra-worldly semanticist. For starters, on this definition, the property of being black is identified partly in terms of the modal relations that it bears to members of a family of related properties and relations. For instance, it is partly constitutive of what it is for something to be black on this definition that, for any objects  $x$  and  $y$ , if  $x$  is black and  $y$  is gray, then, necessarily,  $x$  is darker than  $y$ . This is a modal relation that the property of being black stands in to the property of being gray and the relation of being darker than, and, on this definition, it is partly constitutive of what the property of being black is. Accordingly, if we opt for this definition, we can't appeal to what that property is—its “essence”—in order to explain the modal relations it bears to other properties and relations. But that, of course, is just what the intra-worldly semanticist proposes we do. Opting for this final definition constitutes a radical turn—the move from an *atomist* semantics for predicates, in which one explains the relations of entailment and incompatibility that obtain between the entities that are assigned to predicates as semantic values by appeal to independently intelligible features of these entities, to a *holist* semantics for predicates, in which the entities that are assigned to predicates as semantic values are only intelligible in virtue of the relations of entailment and incompatibility that they bear to one another.

To accept a holist semantics for predicates is a radical divergence from the sort of semantic theory we considered in a last chapter, in which the semantic relations obtain between predicates in virtue of these predicates' independently intelligible semantic values. Recall, on an extra-worldly semantic theory, the semantic values of “black” and “gray” are functions that map each possible world to a certain set of entities, the first function mapping each possible world to the set of black things in that world and the second function mapping each possible world to the set of gray things in that world. Simply given what these two functions are, it follows that, for each world, the set of entities to which the semantic value “gray” maps that world and the set of entities to which the semantic value of “black” maps that world are disjoint. Accordingly, given

the definition of incompatibility provided by the extra-worldly semanticist, the extra-worldly semanticist can maintain that “black” and “gray” are incompatible in virtue of the specific semantic values of these two expressions. Now, of course, we raised a problem for semantic values of these sorts being advertised as models of properties, but we can nevertheless note that semantic values of these sorts do accord with a basic methodological principle: semantic relations that obtain between expressions of a language obtain in virtue of the semantic values of those expressions. Now, the intra-worldly semanticist recognizes the problem with semantic values of these sorts, seeing that our grasp of properties must be more fundamental than the grasp of these functions. However, if, in attempting to define properties, the intra-worldly semanticist opts for this final definition, he cannot maintain this basic principle.

If one opts for this final definition, the meanings of predicates are understood, at least partly, in terms of their relations of implication and incompatibility that they bear to the meanings of other predicates. If one goes this route, it is a short step to the view that the properties that are taken to be the meanings of predicates are in part constituted by the semantic relations that those predicates bear towards other predicates. After all, it is clear from the failure of the third definition that the modal relations that we are permitted to appeal to in providing the fourth definition are relative to the vocabulary of the speakers of the language for which we are providing a semantic theory. It is this specific class of modal relations that is partly constitutive of the properties that figure into the semantic theory. But what could this class of modal relations be other than the semantic relations that obtain between the predicates of the language? That, according to this Sellarsian proposal just what properties are: codifications, in alethic modal terms, of semantic relations between predicates, where these semantic relations between predicates are just the relations of entailment and incompatibility that they stand to one another. This, I believe, is the correct theory of properties. The intra-worldly semanticist, however, cannot accept this theory, for, on his theory, properties are supposed to *explain* the relations of entailment and incompatibility that obtain between predicates of the language. So, we have here a familiar problem: if properties



*explain* the relations of entailment and incompatibility that obtain between predicates, the relations of implication and incompatibility that obtain between predicates cannot *constitute* the properties. The intra-worldly semanticist must give a different account of the their properties figure in their semantic theory. But *what could that account be?*

At this point, it is worth recalling the basic dilemma faced by someone who has fallen prey to the Myth of the Given: they are stuck with a conception of our knowledge of some aspect of the structure of reality according to which it is either *unintelligible* or *incoherent*. It seems to me that the intra-worldly semanticist is stuck in just such a dilemma here. On the one hand, if they *don't* accept the account of properties I've just given, then, since the account of properties I've just given is the only account that really can be given (because it is the correct one), then they have no account of properties. Accordingly, they're stuck with a semantic theory that is, at its base level, *unintelligible*. On the other hand, if they *do* accept the account of properties that I've given, then they appeal to the rules governing the use of predicates in order to account for the entities that are supposed to explain this use. Accordingly, they're stuck with a semantic theory that is, at its core, *incoherent*. These are the only two options for the intra-worldly semanticist. Since, both options are unacceptable, so too is intra-worldly semantics.

### 3.7 Conclusion

Appeals to speakers' grasp of properties and relations is nearly universal in semantic theorizing. We saw, in the last chapter, that attempts to define such entities as constructions from possible worlds either make it impossible to understand how speakers grasp such entities, ending up with an account that's unintelligible, or appeal to speakers' grasp of properties in explaining this grasp, ending up with an account that's incoherent. In this chapter, we considered theories that don't define properties in terms of possible worlds, but, rather, take properties as basic, appealing to speakers' grasp of properties, primitively construed, to explain their knowledge of linguistic meaning. We have now argued that such theories face a similar problem: either one is left with no

account of these properties at all or an account in which they are understood in terms of the very thing that they are supposed to help explain. Though we haven't considered all possible forms of worldly semantics, we have considered enough, I take it to license the main negative conclusion of this dissertation: worldly semantics is committed to the Myth of the Given. Given the constraints of these theories, the worldly knowledge to which such theories appeal cannot be understood except as *simply given* to speakers of a language, and that is not actually a way to understand knowledge at all. To avoid the Myth, we must turn to a very different sort of semantic theory, one that does not presuppose this sort of worldly knowledge but, rather, actually enables us to account for it. That is the positive task to which we now turn.