

9 Ramseyan Humility

David Lewis

1 The Thesis of Humility

Rae Langton has recently argued that Kant is not the thoroughgoing idealist we mistook him for. Our irremediable ignorance of noumena is not ignorance of some remarkable things that are hidden from us behind a veil of appearances. Rather, it is our ignorance of the intrinsic properties of substances. The substances that bear these intrinsic properties are the very same unhidden substances that do indeed affect us perceptually. But they affect us, and they affect other things that in turn affect us, in virtue of their causal powers, which are among their relational properties. Thereby we find out about these substances as bearers of causal powers, but we find out nothing about them as they are in themselves (Langton 1998).

Langton makes an impressive case for her interpretation. But in doing so, she disagrees extensively with previous commentators—who, in turn, disagree extensively with one another. I am not entitled to an opinion about these disputed exegetical questions. In any case, my interest is not in whether the thesis of Humility, as she conceives it, is Kantian, but rather in whether it is true. I shall argue that it is, or at least that something very like it is.

Many of us nowadays believe that things do indeed affect other things in virtue of their intrinsic properties. Something is disposed under certain circumstances to cause a response in something else because it has some intrinsic ground which, in accordance with the contingent laws of nature, would under those circumstances cause the response.¹ If that is true, it might seem to demolish the case for Humility. Langton disagrees:

[If] the ground is distinct from the power, and contingently connected with it, then our orthodoxy is faced with a conclusion surprisingly similar to Kantian Humility. . . . [Our name for the ground] becomes the name for a something-we-

know-not-what—ominously similar to a Kantian thing in itself. Our contemporary orthodoxy must concede that it, too, is faced with a conclusion similar to Kant's—that there are intrinsic features of the world with which we can never become acquainted. (Langton 1998, 176)

To be the ground of a disposition is to occupy a role, but it is one thing to know that a role is occupied, another thing to know what occupies it.

The point generalizes. Being the ground of a certain disposition is only one case among many of role-occupancy. There are a variety of occupied roles, among them nomological roles and others as well. Quite generally, to the extent that we know of the properties of things only as role-occupants, we have not yet identified those properties. No amount of knowledge about what roles are occupied will tell us which properties occupy which roles. I do not agree with Langton that the predicament is “ominous” but I do agree that we are in it. That is what I shall argue here. First I shall argue from fairly weak assumptions that Humility applies to at least some of the fundamental properties. Next I shall argue from somewhat more contentious assumptions that it applies to all of them—or if it doesn't, we are in no position to know that it doesn't, which is no less of a setback for our aspirations to knowledge. Finally I shall argue that if Humility applies to most or all of the fundamental properties, then it spreads to a great range of less than fundamental properties, intrinsic and extrinsic alike.

2 Fundamental Properties

Fundamental properties are those properties that I have elsewhere called ‘perfectly natural’.² They are not at all disjunctive, or determinable, or negative. They render their instances perfectly similar in some respect. They are intrinsic; and all other intrinsic properties supervene on them. They are not conjunctive or structural.³

Unfortunately, all that I've just said about fundamental properties would, if analyzed, lead us back in a definitional circle to the notion of a fundamental property. If we had no initial understanding of such notions as ‘disjunctive’, ‘determinable’, and so on, we would gain no understanding by running around the circle. If, on the other hand, we do have some initial understanding—and I think we do, as witness our abilities to classify—then the circle is benign. When we run around it, our initial understandings of the several notions we meet reinforce one another.⁴

Properties are abundant, as abundant as classes of *possibilia* (which is what I take them to be). But the great majority of properties are utterly

miscellaneous disjunctions, not at all natural. Properties that are even somewhat natural are sparse, and perfectly natural (that is, fundamental) properties are sparser still. I do not know whether some properties correspond to universals or to tropes, but it offends against Occam's razor to imagine that all the abundant unnatural properties do. Rather, it is the fundamental properties, or maybe the near-enough fundamental properties, that are plausible candidates for corresponding to universals or tropes.

Fundamental properties figure in a minimal basis on which all else supervenes. No two possible worlds just alike in their patterns of instantiation of fundamental properties could differ in any other way.

I speak of 'fundamental properties' for short, but they fall into several categories. There are all-or-nothing monadic properties. There are all-or-nothing n -adic relations, at least for smallish n . There are properties that admit of degree, that is, magnitudes; more generally, there are scalar-valued, vector-valued, tensor-valued, . . . magnitudes. There are relational magnitudes. Maybe my list is too long; maybe the magnitudes could somehow be reduced to all-or-nothing properties and relations, but that is a question I shall not take up here.

Scientific theorizing and the discovery of fundamental properties have gone hand in hand. For instance the discovery of the phenomena of electromagnetism and the laws governing them was inseparable from the discovery of the previously unknown, and very likely fundamental, properties of positive and negative charge. So if we had a true and complete "final theory," it ought to deliver a true and complete inventory of those fundamental properties that play an active role in the actual workings of nature.⁵

That inventory might omit two kinds of fundamental properties: idlers and aliens. *Idlers* are those fundamental properties, if any, that are instantiated within the actual world, but play no active role in the workings of nature. *Aliens* are those fundamental properties, if any, that are instantiated within unactualized possible worlds but not within the actual world. We shall consider later whether we have any reason to believe, or any reason to disbelieve, in idlers or in aliens.

3 Ramsification

Let T be the true and complete final theory just considered. The language of T contains T -terms: theoretical terms implicitly defined by T . And there is all the rest of our language, call it O -language. 'O' stands for 'old'; it is the

language that is available to us without benefit of the term-introducing theory T. 'O' does not stand for 'observation'. O-language is not meant to be a 'pure observation language', and indeed I doubt that there could be any such thing. I'll assume, however, that O-language does suffice to express all possible observations, whatever else it may also be able to do. I'll assume also that O-language is interpreted—never mind how. I do not assume it to be first-order, or extensional, or finitary, or free of indexicality. Nor do I assume it to be unmetaphysical, suited only to talk of everyday matters.

All fundamental properties, except for idlers and aliens if such there be, are mentioned in the theory T. If T is the limit (perhaps never reached) of a process in which theorizing and the discovery of fundamental properties go hand in hand, then the fundamental properties mentioned in T will be named by T-terms. I assume that no fundamental properties are named in O-language, except as occupants of roles; in which case T will name them over again, and will say that the property named by so-and-so T-term is the occupant of such-and-such role.

I'll assume, until further notice, that at least two fundamental properties fall in the same category. They might be two monadic properties: unit positive charge and unit negative charge, say. Or they might be two relational magnitudes: space-like distance and time-like distance, say. An inconclusive reason to assume this is the expectation that current physics is not entirely wrong in its inventory of fundamental properties. Fundamental properties belonging to one-membered categories, if such there be, fall outside the scope of my present argument for Humility.

The theory T consists of all the logical consequences of a sentence we shall call the *postulate* of T. This is not a substantive assumption, since I have not said that the postulate is of finite length; but if it were not, it is hard to see how any of the theorems of T could deserve the name of laws. The postulate can be written as $T(t_1, \dots, t_n)$, where t_1, \dots, t_n are the theoretical terms of T, and all other language that appears in the postulate is O-language. We can assume without loss of generality that the T-terms are names, since given suitable copulas, predicates can be replaced by names of properties; and we can assume that no two of them name the same thing. (If they did, a true and complete theory would say so; but then it could be better formulated by using one name twice over.) Some T-terms might be names of other things (for instance the center of the universe, in the unlikely event that it plays a special theoretical role) but for the most part the T-terms will name fundamental properties. When they do, the postulate ought to say so, and we have assumed that O-language is indeed capable of saying so.

Replacing T-terms by variables, we get a formula $T(x_1, \dots, x_n)$. An n -tuple may or may not satisfy this formula, with respect to the actual world and the fixed interpretation of O-language. One which does is called an *actual realization* (for short, a *realization*) of T. One which might is called a *possible realization* of T.

We have assumed that a true and complete final theory implicitly defines its theoretical terms. That means that it must have a unique actual realization. Should we worry about symmetries, for instance the symmetry between positive and negative charge? No: even if positive and negative charge were exactly alike in their nomological roles, it would still be true that negative charge is found in the outlying parts of atoms hereabouts, and positive charge is found in the central parts. O-language has the resources to say so, and we may assume that the postulate mentions whatever it takes to break such symmetries. Thus the theoretical roles of positive and negative charge are not purely nomological roles; they are locational roles as well.

Prefixing existential quantifiers to the formula just considered we get the *Ramsey sentence* of T: 'For some x_1, \dots , for some x_n , $T(x_1, \dots, x_n)$ '. It says that T has at least one actual realization (see Ramsey 1990, chapter 6). What we need to know about the Ramsey sentence is that it logically implies exactly the O-language sentences that are theorems of T. Equivalently: exactly those that are logically implied by the postulate of T.⁶

O-language, we assumed, is rich enough to express all possible observations. Therefore any predictive success for T is equally a predictive success for the Ramsey sentence of T. Since the evidence for T consists in its record of predictive success, there is no way to gain evidence for T that is not equally evidence for the Ramsey sentence.

Though our theory T has a unique actual realization, I shall argue shortly that it has multiple possible realizations. Suppose it does indeed have multiple possible realizations, but only one of them is the actual realization. Then no possible observation can tell us which one is actual, because whichever one is actual, the Ramsey sentence will be true. There is indeed a true contingent proposition about which of the possible realizations is actual, but we can never gain evidence for this proposition, and so can never know it. If there are multiple possible realizations, Humility follows.

Or rather, Humility about some members of the actual realization follows. We have n T-terms, so possible realizations will be n -tuples. Suppose all possible realizations of T have the same 12th and 14th members, but differ elsewhere. Then we can gain evidence that will uniquely identify the 12th and 14th members of the actual realization, but not the other members.

How *does* the theory T go beyond its Ramsey sentence? Exactly by adding the *Carnap sentence*: the truth-functional conditional with the Ramsey sentence as antecedent and the postulate as consequent (Carnap 1963, section 24 D, 963–966). The two things we need to know about the Carnap sentence are, first, that the Ramsey and Carnap sentences together are logically equivalent to the postulate;⁷ and, second, that the Carnap sentence logically implies no O-sentences except for logical truths.⁸ The Carnap sentence of a theory constrains the reference of the theoretical terms, and that is all it does. If the theory is uniquely realized, it says that they refer to the corresponding members of the unique realization. If the theory is multiply realized, it says that they refer to the members of some realization, but it doesn't say which one. Nor does it say what happens if the theory is unrealized. But neither of those silences need concern us here: our true and complete final theory T is uniquely realized.⁹

4 Combinatorialism and Quidditism

Suppose we have the actual realization of T. Maybe some members of the *n*-tuple that realizes T are not fundamental properties, or maybe some belong to single-membered categories. Hold those ones fixed. Permute the rest within their categories to obtain a new *n*-tuple.¹⁰ It too would realize T.

Suppose, for instance, that we start with the actual world, and we permute two fundamental monadic properties F_1 and F_2 , these being the actual referents of the T-terms t_1 and t_2 , leaving all else fixed. Exchange F_1 and F_2 throughout the world. (Suppose for simplicity that possibilities are worlds like ours; though if instead they were some sort of representations of worlds, my argument would go through *mutatis mutandis*.) Then F_2 will be found in exactly those places in space and time (or, more generally, in the pattern of instantiation of fundamental relations and magnitudes) that correspond to the places where F_1 was found originally; and vice versa. And the laws of nature governing F_1 in the permutation will be just the same as the laws governing F_1 originally (more precisely, the laws governing F_2 *vis-à-vis* F_1 in the permutation will be the same as those governing F_1 *vis-à-vis* F_2 originally); and *vice versa*.¹¹ But what could T tell us about besides the locational and nomological roles of the referents of its T-terms? So if this permutation really is a possibility, it is a possible realization of T.

Possibility is governed by a combinatorial principle (see Lewis 1986a, sec. 1.8; Armstrong 1989). We can take apart the distinct elements of a possibility and rearrange them. We can remove some of them altogether.

We can reduplicate some or all of them. We can replace an element of one possibility with an element of another. When we do, since there is no necessary connection between distinct existences, the result will itself be a possibility. How much this means depends on what we take the distinct elements to be. Here, let us take them to include not only spatiotemporal parts, but also abstract parts—specifically, the fundamental properties.

Combinatorialism tells us that the laws of nature are contingent.¹² Let it be a law that every F is a G; combinatorialism generates a possibility in which an F is not a G, so that this law is violated. Combinatorialism therefore collides head on with the view that the laws of nature are necessary.¹³ That is the principal reason why some doubt combinatorialism. For myself, I find combinatorialism far more compelling than the alleged necessity of laws, but I shall not try to adjudicate that question here.

Combinatorialism tells us that possibility is preserved under permutations of items—at least if they are items from the same category. If it is possible that $\neg A\neg B\neg C$ -, and if A, B, and C, are, say, all-or-nothing monadic properties, then it is also possible that $\neg C\neg A\neg B$ -. The actual realization of T is a possible realization; we permute items within more-than-one-membered categories; and what we get is also a possible realization of T. But is it different from the one we had before?

Quidditism is the premise that tells us that the permutation is indeed a different possibility (see Black 2000). Two different possibilities can differ just by a permutation of fundamental properties. They do not differ in whether T is realized, or in what we observe. Given combinatorialism and quidditism, our argument for Humility is complete. We have the actual realization of T, permutations thereof yield different possible realizations of T, and we have no way to tell which one is the actual realization.

Quidditism is to properties as haecceitism is to individuals. If we start with a possibility and permute individuals, combinatorialism says that we get a possibility; haecceitism says that it is a different possibility. Haecceitism says that two possibilities can differ just by a permutation of individuals (see Lewis 1986a, sec. 4.4). I accept quidditism. I reject haecceitism. Why the difference?

It is not, I take it, a difference in *prima facie* plausibility. In both cases alike, haecceitistic or quidditistic distinctions between possibilities seem offhand to make sense. In both cases alike, however, we can feel an uncomfortable sense that we are positing distinctions without differences. To dislike distinctions without differences is to feel some sympathy toward (nontrivial) principles of identity of indiscernibles. To reject haecceitism

is to accept identity of qualitatively indiscernible worlds; to reject quidditism is to accept identity of structurally indiscernible worlds—that is, worlds that differ just by a permutation or replacement of properties. Such principles, applied within a single world, prove too much: for instance, they force us to reject the possibility of a world of two-way eternal recurrence.¹⁴ Are they any better if applied to two different worlds?¹⁵

In both cases alike, this conflict of intuitions is inconclusive, but it seems to me to favor both haecceitism and quidditism. However, haecceitism leads to trouble in a way that quidditism does not. Suppose we allegedly have two indiscernible worlds that differ haecceitistically: at the place where individual A is in W_1 , B is in W_2 , and *vice versa*. We ask: in virtue of what does B rather than A occupy that place in W_2 ? Because the worlds are indiscernible, it seems we can have no informative answer. We can only say “It just does”. That is, we need to invoke the transworld identity of B in W_1 with B in W_2 .

Similarly, *mutatis mutandis*, if we have two structurally indiscernible worlds that differ quidditistically: where property F is in W_1 , G is in W_2 and *vice versa*. In virtue of what does G rather than F occupy that place in W_2 ? It just does. That is to say, we invoke the transworld identity of G in W_1 with G in W_2 . Unmysterious haecceitism demands transworld bilocation of individuals; unmysterious quidditism demands transworld bilocation of properties.

But bilocation of individuals, whether between worlds or times or places, is trouble. For bilocated individuals are apt to have different intrinsic properties at their different locations. For instance, A may be standing and B sitting in W_1 , whereas in the indiscernible world W_2 it is B who is standing and A who is sitting. How can we make sense of an intrinsic difference between something and itself?¹⁶ The problem is serious enough that we should not believe in bilocated individuals at all, in which case haecceitistic differences are utterly mysterious. Bilocated properties raise no similar problem: I can think of no plausible example of an intrinsic (higher-order) property which a bilocated property has at one but not another of its locations.¹⁷ Hence it is unproblematic to think of properties as bilocated in time and space. Indeed, we often do think of them that way, whether we believe they are immanent universals wholly present in different locations, or whether we believe they are classes of possible individuals or of resembling tropes with members at different locations. Transworld bilocation of properties is equally unproblematic. We should be quidditists and not haecceitists because haecceitism, unlike quidditism, carries an unaffordable price.

It would be possible to combine my realism about possible worlds with antiquidditism. I could simply insist that, despite the undoubted ability of properties to indulge in bilocation, no property is ever instantiated in two different worlds. That would get in the way of transworld similarity, which we might have thought was transworld sharing of properties. But I could say for world-bound properties, as I do for world-bound individuals, that a counterpart relation is a substitute for transworld identity: two world-bound properties in two worlds are counterparts if the role of one in one world approximately matches the role of the other in the other. Then we could say that individuals in the two worlds are alike to the extent that the properties of one are counterparts of the properties of the other. But although I could treat individuals and properties alike in this way, why should I want to? It could be for the sake of upholding identity of structurally indiscernible worlds, but I see no good reason for wanting to uphold that principle. Or it could be for the sake of blocking the argument for Humility. But why should I want to block that argument? Why is Humility “ominous”? Who ever promised me that I was capable in principle of knowing everything?

The situation changes radically if we are *actualists*: if we suppose, as many do, that possible worlds are constructions out of actual materials. We might say, for instance, that they are state descriptions in a “Lagadonian” language in which each actual individual and each actually instantiated fundamental property serves as a name for itself (see Lewis 1986a, 145ff.). On the one hand, for an actualist, properties are not really instantiated within worlds; rather a world represents them as being instantiated. So the bilocation of properties within the actual world is now no precedent for transworld bilocation. On the other hand, we have an easy way to account for haecceitistic distinctions involving actual individuals and actually instantiated properties. If worlds mention individuals by (Lagadonian) name, W_2 can simply say that B rather than A occupies a certain place. Quidditistic distinctions are just as easy: W_2 can simply say that G rather than F occupies a certain place. If it is this easy to capture haecceitistic and quidditistic distinctions, the questions of haecceitism and quidditism can be settled—if they can be settled at all—just by deciding whether we’re moved more by the *prima facie* intelligibility of such distinctions or by our dislike of distinctions without a difference.

But what an actualist cannot do is to capture haecceitistic distinctions involving alien individuals—unactualized individuals which cannot be described as recombinations of actually existing parts; or quidditistic distinctions involving alien fundamental properties. The trouble is that alien

individuals and properties are not available within actuality to serve as Lagadonian names for themselves, nor can we construct any other names for them. An actualist faces a choice: he can reject aliens, or he can reject haecceitism and quidditism. (Or he can reject combinatorialism, without which *prima facie* examples of haecceitistic and quidditistic distinctions would not arise.)¹⁸

5 A Second Argument for Humility

Our first argument for Humility—call it the *permutation argument*—has an advantage and a drawback. The advantage is that it makes use of only the most unquestionable of fundamental properties: those that play an active part in the actual workings of nature. The drawback is that it establishes Humility only for those fundamental properties that fall in more-than-one-membered categories. (In the unlikely event that no two fundamental properties fell in the same category, none would be eligible for permutation, so the argument would not establish any sort of Humility at all.) Now I shall give a second argument that has neither the advantage nor the drawback. Whether that is on the whole an improvement, I shall not say.

The skeleton of the second argument—call it the *replacement argument*—is the same as before. We start with the unique actual realization of T; all fundamental properties except idlers and aliens are members of it. If we replace those properties by others, we get a possible realization by combinatorialism, and it is different from the one we had before by quidditism. Since no possible observation gives us evidence that goes beyond the Ramsey sentence, we have no way to tell which possible realization is actual.¹⁹

The difference lies in where we find replacements for the members of the original realization. Instead of permuting those members, we now draw instead on idlers or aliens. Given an abundant supply of idlers or aliens or both, we can replace every fundamental property that is a member of the original realization, whether or not it is the sole member of its category. But should we believe in abundant idlers or aliens?

Begin with aliens. One who believes that ours is one world among many has two reasons to believe in abundant aliens. Both rest on the plausible claim that it is a contingent matter which fundamental properties are instantiated, so it ought to be possible for either more or fewer of them to be instantiated than actually are. So, first, once we appreciate that it is a contingent matter which of them are instantiated, it seems that there

ought to be a world where more of them are instantiated than are instantiated in our world. So properties are instantiated in that world which are alien to ours. And if we take a world where more are instantiated, it seems possible at that world that still more should be instantiated. And so on. What might be questioned is whether there is *any* world where all fundamental properties are instantiated. Is there room for all of them in any one world? But surely our world is nowhere near running out of room. Second, start with our actual world. Any of the actually instantiated fundamental properties have been uninstantiated. So properties that are instantiated at our world are aliens at other worlds. But if there are worlds where some properties are aliens, and ours is not one of them, then our world has a special distinction that some other worlds lack. Why should the world we happen to live in be special in this way? Very likely it isn't. I think this second reason is weaker than the first, but still a strong reason for me to believe in aliens.

My view that ours is one world among many is contentious, to say the least. What should you think about aliens if you are an actualist? You *do* think that our world is special, so my second reason vanishes. My first retains at least some force. However, we saw earlier that an actualist faces a choice between rejecting aliens and rejecting quidditism (or combinatorialism). Perhaps, on balance, rejecting aliens is the wiser choice. Anyway, it doesn't matter which choice he makes. To run the argument from replacement using aliens, he needs a supply of aliens and he also needs combinatorialism and quidditism. So the argument is not for him.

Next, idlers. Surely the laws of nature might have been different in such a way that some of the fundamental properties which actually play an active part in the workings of nature would have been idlers instead. So idlers are possible. At this point, I can say that if there are no idlers in our world, then our world has a special distinction that some other worlds lack. Why should the one world among many that we happen to live in be special in this way? Very likely it isn't. This is at best an inconclusive reason for me, since after all it is a contingent question whether our world has idlers and I have no real evidence. If you are an actualist, it is no reason at all for you, since you *do* think that our world is special. I don't see that an actualist has any reason at all to believe in idlers.

But how much reason has he to *disbelieve* in idlers? I grant that Occam's razor justifies him in having a low degree of belief in idlers. But can he claim to *know* that there are no idlers? I think not. To say they don't exist because we can have no reason to believe they do exist seems nothing better than an appeal to verificationism.

The upshot is that whatever you think about the metaphysics of modality, you should at least agree that for all you know, there may be an abundant supply of idlers. So for all you know, the replacement argument using idlers may succeed. *But that is good enough.* Humility is, after all, a thesis of irremediable ignorance. If we are irremediably ignorant about whether the replacement argument using idlers establishes Humility, then either the argument does establish it, and it is true; or else the argument does not establish it, but we cannot know that. In both cases alike, we are irremediably ignorant about the identities of the fundamental properties that figure in the actual realization of the true final theory.

6 Spreading Humility

So far, we have been concerned with Humility about the fundamental properties. But the great majority of properties, including most of the fairly natural properties, most of the properties we are acquainted with and have names for, and most of the intrinsic properties that Langton's Kant thought were unknown to us, are not fundamental.

I said that the fundamental properties (including relations and magnitudes) are a minimal basis on which all else supervenes. 'All else' should at least include all qualitative properties of things, natural or not, intrinsic or not. (Perhaps this should be taken as the definition of 'qualitative property'.)

First, take the case of a qualitative property instantiated by the entire world. It will supervene on the pattern of instantiation of fundamental properties by the parts of the world. Unless it is itself a fundamental property, it will be either a structural property constructed out of the fundamental properties, or else a truth-functional compound, finite or infinite, of such structural properties.

Next, take the case of an intrinsic property instantiated by some part of the world. Since it is intrinsic, this same property could have been instantiated by something that did not coexist with anything distinct from itself—in other words, by the entire world (see Langton and Lewis 1998). So this case reduces to the previous one. The property is either fundamental, or structural, or a truth-functional compound of structural properties.

Third, take the case of an extrinsic property of something. There will be some more inclusive thing such that the smaller thing has the extrinsic property because the larger thing has some intrinsic property. (In some cases the more inclusive thing will be the entire world, in other cases not.)

For instance, I have the extrinsic property of being an uncle in virtue of an intrinsic property of my family. So this case reduces to the previous one.

Suppose we have a structural property: say, the property of being composed of an F bearing relation R to a G, where F, R, and G are fundamental. If we do not know the identities of the properties we have called 'F', 'R', and 'G' we also do not know the identity of the structural property constructed out of them. If the so-called F might for all we know be either F_1 or F_2 , and R might be either R_1 or R_2 , and G might be either G_1 or G_2 , there are eight different possibilities for what the structural property might be. The same point applies to more complicated structural properties, and to at least some truth-functional compounds of structural properties. So if Humility applies to some or all of the fundamental properties, then it applies also to very many more properties. We may know them as role-occupants, including both their roles in scientific theory and their roles in daily life, and that might pass for "knowing what they are" by lax and commonplace standards.²⁰ But if they are structural properties and we don't know the recipe for their construction out of fundamental properties, or if they are certain truth-functional compounds of structural properties and we don't know the recipes for those structural properties (or we don't know, even in a lax sense, what structural properties they are compounded from), or if they are unidentified fundamental properties, then in a more demanding sense we do not know what they are, no matter how familiar we may be with them.

There are some exceptional truth-functional compounds, however, to which Humility does not apply. We saw that being composed of an F bearing R to a G might be any of eight different structural properties. We can identify the disjunction of all eight, even if we can't identify any one of its disjuncts. Likewise, even if we cannot identify the fundamental or structural property that actually occupies a certain role, we can identify the property of having whatever property it is that occupies that role (in the case in question). Perhaps it is thanks to such properties that O-language manages to be fully interpreted.

7 Ineffable Ignorance

Suppose we know of some property only as the occupant of a role, perhaps the role of being the seventeenth member of the actual realization of T, or perhaps some role that is less cumbersome to describe. Suppose Humility applies to this property. We cannot answer the question: which property

occupies that role? But worse: not only can we not answer that question, we can't even ask it.

How's that? Didn't I ask it just now? Yes, but I didn't ask it properly. To ask it properly is to ask it as a multiple-choice question: is it this, that, or the other? It is to ask it in such a way that, apart from limitations of finitude, the addressee could list all the alternative possible answers to it (see Belnap and Steel 1976). That is what I did not do, and that is what we cannot do.

There are alternative *answer-propositions*, to be sure. For each fundamental property F (of the right category), there is a contingent proposition true at all and only the worlds where F occupies the role in question. But we do not have alternative *answer-sentences* that express those alternative answer-propositions and do so in such a way that we can know which sentence expresses which proposition.

"Which property occupies the role?—The occupant of the role, whatever that is." A true answer, sure enough, but not an answer to the question we meant to be asking. Indeed, not an answer to any question we'd be likely to ask: the only information it conveys is that the role is uniquely occupied.²¹

Reference-fixing will not help. To make it easy for ourselves, assume we have descriptions to rigidify. Assume we are—rashly—convinced of the nonexistence of both idlers and aliens, so the only fundamental properties that exist are those we can describe as 'the so-and-so member of the actual realization of T'—for short: 'the occupant of the so-and-so role'. Assume for simplicity that we have only one category and that all members of the actual realization of T are fundamental properties. Use 'actual' as a rigidifier. Now can we ask our multiple-choice question thus?

Which property occupies the seventeenth role? Is it the actual occupant of the first role? Is it the actual occupant of the second role? . . . Is it the actual occupant of the seventeenth role? . . . Is it the actual occupant of the *n*th role?

No sooner do we ask our question this way than we seem to know the answer: it is the actual occupant of the seventeenth role. But we have cheated. Each of our answer-sentences does indeed express one of the alternative contingent answer-propositions, but we do not know which sentence expresses which proposition. We know that 'It is the occupant of the seventeenth role' expresses the true answer-proposition, but that is no help in knowing which one is true. We have fixed the reference blindly, not knowing which rigidified descriptions were rigid designators of which fundamental properties.²²

8 Humility about Qualia

A friend of phenomenal qualia might speculate that all the actually instantiated fundamental properties are qualia. That would not preclude them from also occupying physical roles—a sort of panpsychism. So, even if our true final theory is a physical theory, they might be the members of its actual realization. He might also accept the *Identification Thesis*: anyone acquainted with a quale knows just which property it is. Now it may seem that we can know the identities of the fundamental properties after all—we need only become acquainted with them.²³

Since he rejects the conclusion of our arguments for Humility, the friend of qualia had better reject some assumption that went into those arguments. He has two choices for what to reject. (1) I assumed that O-language names fundamental properties only as occupants of roles, and that T tells us that the referent of so-and-so T-term occupies such-and-such role. The friend of qualia may say that all fundamental properties are named in O-language, and not as occupants of roles, provided we take O-language to be the inner private language of someone who has become acquainted with certain qualia. Or (2) if he would prefer to have nothing to do with the dubious notion of an inner private language, he could instead take issue with our assumption that O-language suffices to express all possible observations. He could say that it does not suffice to express those observations that consist in gaining acquaintance with qualia.

As a materialist, I reply predictably: the Identification Thesis is false (see Lewis 1995). If the Identification Thesis is built into the very definition of qualia, there are no qualia. If not, there may well be qualia, but they are known to us only as role-occupants, and there are multiple possible occupants of the roles in question. Their roles are psychological: for instance, qualia confer on us abilities to recognize and imagine what we have previously experienced. But we cannot tell just by acquaintance which property occupies such a role. Obviously not, if the occupant of the role is some complicated neural property. But not if it is a fundamental property either. In that case, Humility applies to qualia just as it would to any other fundamental properties that were known to us only as role-occupants.

Even if the Identification Thesis is false for us, might it be true for some other being, say God? Indeed, might God have the supernatural power to become acquainted with *all* fundamental properties, qualia or not, and to identify each of them by acquaintance? And if Humility did not apply to God, would He then be able to tell us just which fundamental property it is that occupies any given role?

If there were a God, who knows what supernatural powers He might have. But no matter what knowledge He might gain by acquaintance with the fundamental properties, He could not share it with us. Since we cannot express any of the answer-propositions to the question which fundamental property occupies a given role—not in such a way that we know which of the propositions we are expressing—God would have no way to communicate His knowledge to us. If He wanted to remedy our ignorance, His only recourse would be to impart to us His own power to identify properties by acquaintance.

Acknowledgments

Thanks are due to Karen Bennett, Robert Black, Cian Dorr, Allan Gibbard, Charles Goodman, John Hawthorne, Anja Jauernig, Rae Langton, Jonathan Schaffer, Ted Sider, and Dean Zimmerman.

Notes

1. See, *inter alia*, Jackson, Pargetter, and Prior 1982; Prior 1985; Armstrong 1997, 69ff.; and Lewis 1997 (reprinted as chapter 7 in Lewis 1999). These authors disagree about whether a disposition should be identified in each particular case with its categorical ground in that case, or whether instead it is the disjunctive property of having some categorical ground or other. But that is a merely terminological question and need not concern us here.
2. See Lewis 1983a. Neither name has proved fully satisfactory. ‘Fundamental properties’ suggests to some that they must be properties of fundamental particles. They might be; or they might instead be properties of space-time points, or of extended things, or even of unlocated things. ‘Natural properties’ suggests to some that it is nature that makes them natural: in other words, that it is contingent which properties are natural. That was not my intention. The properties that figure in the fundamental laws of nature are natural, but that is not because figuring in the laws makes them natural. Rather it is because regularities are fit to compete for the status of lawhood only when formulated in terms of perfectly natural properties (*ibid.*, 367ff.).
3. That creates a problem about the properties instantiated in an infinitely complex world of “structures all the way down.” I take infinite complexity to be a genuine possibility, though a far-fetched one. (I don’t know what to say to someone who thinks it is not far-fetched.) For it may be that all of the properties instantiated in such a world are structural. If so, and if fundamental properties are never structural, we must conclude that in such a world no perfectly fundamental properties are instantiated, but only near-enough fundamental properties. We could then

conclude that some of the privileges otherwise reserved for perfectly fundamental properties can belong instead to the near-enough fundamental properties that are instantiated in infinitely complex worlds: for instance, the privilege of appearing in fundamental laws of nature, or the privilege of corresponding to universals or tropes, or the privilege of constituting a basis—not, however, a minimal basis—on which all else supervenes. I see no reason to object to these conclusions.

4. In my view, this is the proper conclusion to draw about the famous Quinean circle from analyticity to synonymy and back. Quine is right that there is a problem about analyticity, but the problem is not the circle. Rather it is that in most interesting cases, candidates for analyticity are analytic under some legitimate resolutions of semantic indeterminacy and not under others.

5. Optimists hope and expect that we will discover the final theory someday soon, or anyway someday. I share the hope and expectation, but I am not assuming it. Maybe scientific research will go out of fashion. Maybe the task of fully understanding the workings of nature is just too hard for us. But whether or not we will ever discover the final theory, it nevertheless exists at least in the way never-to-be-written poems do. Materialists hope and expect that the true final theory, whether or not we ever discover it, will be a physical theory: it will take over large parts of current physics, perhaps in corrected form, and it will be in the distinctive style of current physics. For instance, it will explain the workings of large things in terms of the workings of their very small parts; it will get by with very few fundamental properties and laws; and it will not posit different fundamental properties and laws for the conscious and the unconscious parts of nature, or for the living and the dead parts of nature. Again I share the hope and expectation, but again I am not assuming it.

6. *Proof.* Let S be any O-language sentence. If the Ramsey sentence implies S, so does the postulate, because the postulate implies the Ramsey sentence by Existential Generalization. Conversely, suppose the postulate implies S.

- (1) Then $[T(t_1, \dots, t_n) \text{ only if } S]$ is a logical truth.
- (2) Then $\text{All } x_1, \dots, \text{All } x_n [T(x_1, \dots, x_n) \text{ only if } S]$ is a logical truth.
- (3) Then $(\text{Some } x_1, \dots, \text{Some } x_n [T(x_1, \dots, x_n)])$ only if S is a logical truth.
- (4) Then the Ramsey sentence implies S.

The only step that requires comment is the step from (1) to (2). This step is an instance of the rule of Universal Generalization—a rule which, though it does not preserve truth, does preserve logical truth in any system of logic worthy of the name. Its rationale is that logic does not discriminate between things. The only way to have a logical truth, –t–, about a particular named thing (when the name ‘t’ is neither structured nor an item of logical vocabulary, as indeed our T-terms are not) is to have it as an instance of a general logical truth: $\forall x[-x-]$. Note that neither this nor any other step of the proof depends on any idiosyncrasies of first-order logic, or extensional logic, or finitary logic.

7. *Proof.* The Ramsey and Carnap sentences jointly imply the postulate by *Modus Ponens*. The postulate implies the Ramsey sentence by Existential Generalization, as already noted. The postulate implies the Carnap sentence because it is the consequent of the Carnap sentence.

8. *Proof.* We proved that the Ramsey sentence of $T(t_1, \dots, t_n)$ implies the same O-sentences as $T(t_1, \dots, t_n)$ itself. $T(t_1, \dots, t_n)$ needn't be the postulate of T; the proof works no matter what sentence it may be. So in particular the Ramsey sentence of the Carnap sentence implies the same O-sentences as the Carnap sentence itself. But the Ramsey sentence of the Carnap sentence is a logical truth, so it implies no O-sentences except logical truths. So the same is true of the Carnap sentence. Again, the proof does not depend on any idiosyncrasies of first-order logic, extensional logic, or finitary logic.

9. I once proposed adding that if a theory has no realization, or multiple realizations, its theoretical terms do not refer (Lewis 1970). I'd now say that if it is unrealized but almost realized, its theoretical terms refer to the members of its unique near-realization, if there is one; and that if it has multiple realizations (or near-realizations) its theoretical terms have indeterminate reference. My reason for saying that the theoretical terms of a multiply realized theory do not refer was that a theorist may be presumed to have intended to implicitly define the terms he introduces. But there is a simpler way to respect the theorist's presumed intention: we should write the postulate in such a way that his theory cannot be multiply realized.

10. Robinson 1993, 17ff., discusses a permutation of two fundamental properties.

11. If laws of nature are suitable regularities, the exchange of locations of P_1 and P_2 will guarantee the exchange of nomological roles. If laws of nature are certain special law-making relations of fundamental properties, P_1 and P_2 will have to be exchanged also as arguments of these law-making relations to guarantee the exchange of nomological roles. On these two conceptions of lawhood, see further endnote 12.

12. One way for the laws of nature to be contingent is for them to be those regularities that earn a place in the true system that achieves the best possible balance of simplicity and strength. See Mill 1843, Book III, chapter IV, section 1; Ramsey 1990, chapter 7A, "Universals of Law and of Fact," 140–144; and Lewis 1973, 73ff. Another way for laws to be contingent is for there to be a law-making relation N between fundamental properties, such that it is contingent which properties stand in this relation but necessary that whenever $N(F, G)$, every F is a G. (Or perhaps what's necessitated is something a little more complicated that allows laws to be defeasible.) See Dretske 1977; Tooley 1977; Armstrong 1983. Although the first-order law that every F is a G is contingent, the higher-order law that whenever $N(F, G)$, every F is a G is necessary. Whether that violates combinatorialism is a disputed matter.

13. Shoemaker 1980; Swoyer 1982; Ellis and Lierse 1994. Necessity of the laws would imply that combinatorialism was false. It would block our present argument for

Humility, but would not imply that it was false. Contingency of the laws implies nothing one way or the other about the truth of combinatorialism, and nothing one way or the other about Humility.

14. Further, principles of identity of indiscernibles lead to the following question. Nobody doubts that there could be two nonidentical *almost*-indiscernibles. If they had been just a little bit different, they would have been nonidentical indiscernibles. What stops them from being that little bit different? See Adams 1979.

15. Some seeming cases of haecceitistic difference are compelling enough that an antihaecceitist needs to explain them away. I do so by arguing that not all possibilities are possible worlds; rather, some are possibilities for individuals (or pairs, triples, . . . of individuals); and it is not haecceitism if individual possibilities within a single world differ by a permutation of individuals. See Lewis 1986a, 230ff.

16. I take it that explanations which replace intrinsic monadic properties by relations, or which take individuals not to have their properties simpliciter but only to stand in some relation to them, are unacceptable. See Lewis 1986a, 199ff., and Lewis 2002.

17. Armstrong (1989, 67) considers and rightly rejects a somewhat plausible case of intrinsic difference between a property in different locations: a universal that is allegedly simple at one world, complex at another.

18. Armstrong (1989, 54ff.) chooses to reject aliens. Black (2000, 100) chooses to reject quidditism. And I (Lewis 1986a, 158ff.), though I reject haecceitism, reject actualism in part because it cannot combine acceptance of alien properties with quidditism.

Plantinga (1974) hopes to side-step the problems of haecceitism about alien individuals and quidditism about alien properties by invoking allegedly actual surrogates for these aliens. Fine (1985) argues convincingly that this will not do because the surrogates presuppose the *possibilia* they are meant to be surrogates for.

19. Black (2000, 195) discusses a replacement of actually instantiated properties by alien properties.

20. By lax and commonplace standards, coming to know what something is means, roughly, acquiring a fuller description of it. You could come to know who Great Leader is by learning that he is the man on yonder balcony giving the clenched-fist salute. Equally, you could come to know who the man giving the salute is by learning that he is Great Leader. Discovering his haecceity, or even the necessary and sufficient conditions for something to be his counterpart, doesn't enter into it.

21. If that. The role might be multiply occupied, the improper description 'the occupant of the role' might be indeterminate in reference, and we might be

following the supervaluationist policy of coping with semantic indeterminacy by saying what's true on all resolutions of the indeterminacy.

22. The real lesson of Kripke's alleged examples of *a posteriori* necessity and *a priori* contingency (Kripke 1980, 97ff., 128ff.), is that when we have reference fixing, we have two different ways for a sentence to express a proposition. It may happen that the proposition expressed one way is necessary and knowable *a priori* but the proposition expressed the other way is contingent and knowable only *a posteriori*.

See Chalmers 1996, 56ff.; Stalnaker 1978; Jackson 1998a, 46ff. Let sentence S contain terms subject to reference fixing. (Ignore other forms of indexicality.) The *primary intension* of S (Chalmers), or the *diagonal proposition* (Stalnaker), or the *A-intension* (Jackson) is the proposition that is true at just those worlds W such that S is true at W with respect to reference fixing as it is at W. The *secondary intension* of S (Chalmers), or the *horizontal proposition* (Stalnaker), or the *C-intension* (Jackson) is the proposition that is true at just those worlds W such that S is true at W with respect to reference fixing as it is in actuality. If, for instance, the reference of 'heat' is fixed on molecular motion at the actual world but on other things at other worlds, the primary intension of 'Heat is molecular motion' is contingent but the secondary intension is necessary; and *vice versa* for 'Heat is the occupant of the heat-role'.

So far, when I have spoken of the propositions expressed by the answer-sentences, I've meant the secondary intensions. If instead we took the primary intensions, it would be just as if we'd omitted the rigidifier 'actual'. As already noted, that would not be a way to ask the question we meant to ask. We would know the primary intensions of the answer-sentences, but these would not be the answer-propositions that identify the properties that occupy the roles.

23. See Maxwell 1978; and for a sympathetic discussion of the view, see Chalmers 1996, 135ff. Just as our panpsychist thinks it is because fundamental properties *are* qualia that we can know their identities by acquaintance, so others think it is because (or insofar as) they are *not* qualia that we cannot know their identities. See, *inter alia*, Russell 1927, 497; Unger 1999.