7

Innateness and Ontology, Part II: Natural Kind Concepts

> [It is] a matter quite independent of . . . wishing it or not wishing it. There happens to be a definite intrinsic propriety in it which determines the thing and which would take me long to explain.
>
> - Henry James, *The Tragic Muse*

HERE's how we set things up in Chapter 6: suppose that radical conceptual atomism is inevitable and that, atomism being once assumed, radical conceptual nativism is inevitable too. On what, if any, ontological story would radical conceptual nativism be tolerable?

However, given the preconceptions that have structured this book, we might just as well have approached the ontological issues from a different angle. I've assumed throughout that informational semantics is, if not self-evidently the truth about mental content, at least not known to be out of the running. It's been my fallback metaphysics whenever I needed an alternative to Inferential Role theories of meaning. But now, according to informational semantics, content is constituted by some sort of nomic, mind world relation. Correspondingly, *having* a concept (concept possession) is constituted by *being in* some sort of nomic, mind world relation. It follows that, if informational semantics is true, then there must be *laws about* everything that we have concepts of. But how could there be laws about *doorknobs*?

The answer, according to the present story, is that there is really only one law about doorknobs (qua doorknobs); viz. that we lock to them in consequence of certain sorts of experience.[1] And this law isn't really about doorknobs because, of course, it's really about us. This is quite a serious point. I assume that the intuition that there aren't laws about doorknobs (equivalently, for present purposes, the intuition that doorknobs aren't a 'natural kind') comes down to the thought that there's nothing in the world

whose states are reliably connected to doorknobs qua doorknobs *except our minds*. No doubt, some engineer might construct a counter-example a mindless doorknob detector; and we might even come to rely on such a thing when groping for a doorknob in the dark. Still, *the gadget would have to be calibrated to us* since there is nothing else in nature that responds selectively to doorknobs; and, according to the present account, it's *constitutive of doorknobhood* that this is so. The point is: it's OK for there to be laws about doorknobs that are really laws about us. Doorknobs aren't a natural kind, but *we are*.

What with one thing and another, I've been pushing pretty hard the notion that properties like *being a doorknob* are mind-dependent. I needed to in Chapter 6 because, if doorknobs aren't mind-dependent, there is only one way I can think of to explain why it's typically *doorknob*-experiences from which the concept DOORKNOB is acquired: viz. that DOORKNOB is learned inductively. And I didn't want that because the Standard Argument shows that only *non*-primitive concepts can be learned inductively. And it's been the main burden of this whole book that all the evidence - philosophical, psychological, and linguistic - suggests that DOORKNOB is primitive (unstructured); and, for that matter, that so too is practically everything else. Likewise, in this chapter, I need *being a doorknob* to be mind-dependent because there is only one way I can think of to reconcile informational semantics, which wants there to be laws about doorknobs, with the truism that doorknobs aren't a natural kind; viz. to construe what appear to be laws about doorknobs as really laws about "our kinds of minds".

But all this stuff about the mind-dependence of *doorknobhood* invites a certain Auntie-esque complaint. Viz.:

> *I get it; the good news is that DOORKNOB isn't innate; the bad news is that there aren't any doorknobs. Aren't you ashamed of yourself?*

I am definitely sensitive to this criticism. For I'm a Realist about doorknobs, I am. I think there are *lots* of doorknobs, and I wouldn't consider for a moment holding a metaphysical view which denies that there are. So, one of the main questions I want to consider in this chapter is: what, if any, consequences would the (putative) mind-dependence of *doorknobhood* have for issues about Metaphysical Realism? My answer will be 'none', and this for two reasons: first, because being mind-dependent is perfectly compatible with being real; and second, and more important, because DOORKNOB isn't the general case. If there are lots of our concepts that express mind-dependent properties *there are also lots*

Doorknobs are Real because Minds are Real

The first of these considerations is entirely banal. Suppose, per hypothesis, that DOORKNOB expresses a property that things have in virtue of their effects on us. Suppose, in particular, that *being a doorknob* is just *having the property that minds like ours reliably lock to in consequence of experience with typical doorknobs*. Well, then, there *are* doorknobs iff the property that minds like ours reliably lock to in consequence of experience with typical doorknobs is instantiated. Which, of course, it is; every doorknob has it, and there are, as previously remarked, lots of doorknobs.

Look, there is simply *nothing wrong* with, or ontologically second-rate about, being a property that things have in virtue of their reliable effects on our minds. For we really do have minds, and there really are things whose effects on our minds are reliable. If you doubt that we do, or that there are, then whatever is the source of your scepticism, it can't be metaphysical considerations of the sort that I've been claiming bear on the nature of *doorknobhood*. Perhaps it's that you're worried about evil demons?

Fingers, I suppose, are, hand-dependent: if there were no hands, there could be no fingers; if you had your fingers on your feet they'd be your toes. This is all entirely compatible with the rigorous Metaphysical Realism about fingers which, surely, common sense demands. For, since there really are hands, such metaphysical conditions for the instantiation of *fingerhood* as its hand-dependence imposes are *ipso facto* satisfied. *Since there are hands, the metaphysical dependence of fingers on hands is not an argument for there not being fingers.* Similarly, *mutatis mutandis*, for the case of doorknobs. Since there are minds, the ontological conditions which the mind-dependence of *doorknobhood* imposes on there being doorknobs are *ipso facto* satisfied. The mind-dependence of *doorknobhood* is not an argument for there not being doorknobs.

I wouldn't be going on about this so, except that it appears to have occasioned much confusion, and some inadvertent comedy, in the cognitive science community. (And in ever so many Departments of English Literature. And in France.) Here, for one example among multitudes, is George Lakoff getting himself into a thorough muddle about Tuesdays:

> If . . . symbols get their meaning only by being associated with things in the world, then weeks must be things in the world. But weeks do not exist in nature . . . Does 'Tuesday' refer to an aspect of 'external reality' reality external to human beings? Obviously not. That reality is constituted by the minds of human beings collectively . . . it is not an 'external' reality. [The word] 'Tuesday' cannot get it

I'm unclear exactly what work Lakoff thinks "external" is doing in this passage, and his persistently putting it into shudder-quotes suggests that he is too. But notice the repeated contrast of "constituted by human minds" and the like with (externally) "real" and the like. The inference that we're being offered is apparently: constituted by minds and so *not* (externally) real.

Now, it's true, of course, that Tuesdays are mind-dependent in at least the following pretty straightforward sense: whether today is Tuesday depends on what conventions people adhere to; and that people adhere to the conventions that they do, or to any conventions at all, depends on their having minds. So: no minds, no Tuesdays. But *it does not follow that there are no Tuesdays*; the minor premiss is missing. Nor does it follow that there is no fact of the matter about whether today is Tuesday (or about whether it is true that today is Tuesday). Nor does it follow that Tuesdays aren't real. Nor does it follow that 'Tuesday' doesn't really refer to Tuesday. As for whether it follows that Tuesdays aren't "'externally'" real, or that 'Tuesday' doesn't refer to an "'external'" reality, that depends a lot on what "'external'" means. Search me. I would have thought that minds don't have outsides for much the same sorts of reasons that they don't have insides. If that's right, then the question doesn't arise.

Likewise, there are many properties that are untendentiously mind-dependent though plausibly *not* conventional; *being red* or *being audible* for one kind of example; or *being a convincing argument*, for another kind; or *being an aspirated consonant*, for a third kind; or *being a doorknob*, if I am right about what doorknobs are. It does not follow that there are no doorknobs, or that no arguments are convincing, or that nothing is audible, or that the initial consonant in 'Patrick' is anything other than aspirated.[2] All that follows is that whether something is audible, convincing, aspirated, or a doorknob depends, *inter alia*, on how it affects minds like ours. Nor does it follow that doorknobs aren't "in the world". Doorknobs are constituted by their effects on our minds, and *our minds are in the world*. Where on earth else could they be?

[2] Compare Jackendoff: "Look at the representations of, say, generative phonology . . . It is strange to say that English speakers know the proposition, *true in the world independent of speakers* [*sic*], that syllable-initial voiceless consonants aspirate before stress . . . In generative phonology . . . this rule of aspiration is regarded as a principle of internal computation, not a fact about the world. Such semantical concepts as implication, confirmation, and logical consequence seem curiously irrelevant" (1992: 29). Note that, though they are confounded in his text, the contrast that Jackendoff is insisting on isn't between *propositions* and *rules/principles* of computation; it's between *phenomena of the kind that generative phonology studies* and *facts about the world*. But that 'p' is aspirated in 'Patrick' is a fact about the world. That is to say: it's a fact. And of course the usual logico-

I'm considering (and endorsing) reasons why no sort of Idealism is implied by the view that the relation between *being a doorknob* and *falling under a concept that minds like ours typically acquire from stereotypic doorknob-experiences* is metaphysical and constitutive. I've been arguing that *not even Idealism about doorknobs* follows; doorknobs are real but mind-dependent, according to the story I've been telling.

But I think there's another, and considerably deeper, point to make along these lines: I haven't suggested, and I don't for a moment suppose, that *all* our concepts express properties that are mind-dependent. For example, we have the concept WATER, which expresses the property of *being water*, viz. the property of *being* H_2O. We also have the concept H_2O, which expresses the property of *being* H_2O, viz. the property of *being water*. (What distinguishes these concepts, according to me, is that the possession conditions for H_2O, but not for WATER, include the possession conditions for H, 2, and O. See Chapters 1 and 2.) Assuming informational semantics, having these concepts is being locked to the property of *being water*; and being water is a property which is, of course, *not* mind-dependent. It is not a property things have in virtue of their relations to minds, ours or any others.

I suppose that natural kind predicates *just are* the ones that figure in laws; a fortiori, since water is a natural kind, there isn't a problem about how there could be laws about the property that the concept WATER expresses. But if water isn't mind-dependent, where do concepts like WATER come from? How do you lock a mental representation to a property which, presumably, things have in virtue of their hidden essences? And what, beside hypothesis testing, could explain why you generally get WATER from experience with water and not, as it might be, from experience with giraffes? What, in short, should an enthusiast for informational theories of content say about concepts that express natural kinds?

All in due time. For now, I propose to tell you a fairy tale. It's a fairy tale about how things were back in the Garden, before the Fall; and about what the Snake in the Garden said; and about how, having started out by being Innocents, we've ended up by being scientists.

Concepts of Natural Kinds

How Things Were, Back in the Garden

Once upon a time, back in the Garden, all our concepts expressed (viz.

expresses the property that things have when they seem to us to be of the same kind as instances of the doorknob stereotype; and we had the concept HOLE, which expresses the property that things have when they seem to us to be of the same kind as instances of the hole stereotype; and we had the concept A NICE DAY, which expresses the property that things have when they seem to us to be of the same kind as instances of the nice day stereotype . . . etc. (Also, I suppose we had logico-mathematical concepts; about which, however, the present work has nothing to say.)

Because the concepts we had back in the Garden were all concepts of mind-dependent properties, there was, back then, a kind of appearance/reality distinction that we never had to draw. We never had to worry about whether there might be kinds of things which, though they satisfy the DOORKNOB stereotype, nevertheless are not doorknobs. We never had to worry that there might be something which, as it might be, had all the attributes of a doorknob but was, in its essence, a Twin-doorknob. Or, who knows, a giraffe.[3]

But also, because we were Innocent, we didn't have the concept WATER, or the concept CONSONANT, or the concept LEVER, or the concept STAR. Perhaps we had concepts that were (extensionally) sort of like these; perhaps we used to wonder who waters the plants. But, if so, these concepts were importantly different from the homophonic counterparts that we have now. For it's compatible with the real concept WATER that there should be stuff that strikes us as being of the very same kind as instances of the water stereotype but that isn't water because it has the wrong kind of hidden essence (XYZ, perhaps). And it's compatible with the real concept STAR that there should be things that strike us as very different from paradigm stars, but which do have the right kind of hidden essences and are therefore stars after all (a black dwarf, perhaps; or the Sun). And it's compatible with the concept CONSONANT that we have now that there should be sorts of things that strike us as neither clearly consonants nor clearly not consonants but which, because they have the right kinds of hidden essences, really are consonants whether or not we think they are (*l*s and *r*s, perhaps).

I'll presently have much more to say about what concept of water we could have had in the Garden; and about how it would have been different

[3] There were, to be sure, *faux* doorknobs, fake doorknobs, *trompe l'œil* doorknobs, and the like; these were particulars which looked, at first glance, to satisfy the doorknob stereotype but, on closer examination, turned out not to do so. *Doorknob* vs. *trompe l'œil doorknob* is a distinction *within* mind-dependent properties; hence quite different from the difference between *doorknob* and, as it might be, *water*. (In consequence, drawing an

from the concept of water that we have now. And about how to square that difference with what an atomistic and informational semantics says about the individuation of concepts. But this will do to be getting on with: back in the Garden, when we were Innocent, we never thought about kinds of things which, though they are much the same in their effects on us, are *not* much the same in their effects on one another. Or about kinds of things which, though they are much the same in their effects on one another, are strikingly different in their effects on us. Back in the Garden, when we were Innocent, we took it for granted that there isn't any difference between similarity *for us* and similarity *sans phrase*; between the way we carve the world up and the way that God does.

Then came the Snake.

What the Snake Said

'I have here,' the Snake said, 'some stuff that will no doubt strike you, in your Innocence, as a sample of bona fide, original, straight off the shelf, *X*-ness. But come a little closer—come close enough to see how the stuff is put together—and you'll see that it isn't *X* after all. In fact, it's some kind of *Y*.'

'Sucks to how it's put together,' we replied, in our Innocence. 'For a thing to strike us as of a kind with paradigm *X*s *just is* for that thing to be an *X*. *X*-ness *just is* the property of being the kind of thing to which we do (or would) extrapolate from appropriate experience with typical *X*s. Man is the measure; *vide* doorknobs.'

'That,' the Snake replied, 'depends. Since we're assuming from the start that *X*s and *Y*s are, for practical purposes, indistinguishable in their effects on you, it follows that thinking of both *X*s and *Y*s as *X*s will do you no practical harm. For example, for purposes of longevity, reproductive efficiency, and the like, it's all one whether you ingest only *X*s under the description '*X*' or you ingest both *X*s and *Y*s under that description. But that is ingest; I am in earnest. If you want to carve Nature at the joints, if you want to know *how the world seems to God*, you will have to learn sometimes to distinguish between *X*s and *Y*s even though they taste (and feel, and look, and sound, and quite generally strike you as) much the same. It's entirely up to you of course. Far be it from me to twist your arm. (Sign here, please. In blood.)'

We fell for that, and it was, on balance, a fortunate Fall. The trouble with being Innocent is that, although how God made things *sometimes* shows up in broad similarities and differences in the way that they strike us in very highly contrived, quite unnatural environments; experimental environments, as it might be. For it's sometimes only in terms of a taxonomy that classifies things by similarities and differences among the ways that they do (or would) behave in *those* sorts of environments, that we can specify the deep generalizations that the world obeys. We are, after all, peculiar and complicated sorts of objects. There is no obvious reason why similarity in respect of the way that things affect us should, in general, predict similarity in the way that they affect objects that are less peculiar than us, or less complicated than us, or that are peculiar and complicated in different ways than us.[4]

Unless, however, we contrive, with malice aforethought, that things should strike us as alike only if they are alike in respect of the deep sources of their causal powers: that they should strike us as alike only if they share their hidden essences. So, for example, we can set things up so that the chemicals in the bottles will both turn the paper red (and thereby strike us as similar) if, but only if, they are both acids. Or, we can set things up so that both meters will register the same (and thereby strike us as similar) if, but only if, there's the same amount of current in both the circuits; and so on. The moral is that whereas you lock to *doorknobhood* via a metaphysical necessity, if you want to lock to a natural kind property, you have actually to *do the science*.

So much for the fairy tale. It's intuitively plausible, phylogenetically, ontogenetically, and even just historically, to think of natural kind concepts as late sophistications that are somehow constructed on a prior cognitive capacity for concepts of mind-dependent properties. But intuitively plausible is one thing, true is another. So, *is* it true? And, what does "doing the science" amount to? How, having started out as Innocents with no concepts of natural kinds, could we have got to where we are, with natural kind concepts like WATER? I turn to these questions in, more or less, that order.

[4] In just this spirit, Keith Campbell remarks about colours that if they are "integrated reflectances across three overlapping segments clustered in the middle of the total electromagnetic spectrum, then they are, from the inanimate point of view, such highly arbitrary and idiosyncratic properties that it is no wonder the particular colors we are familiar with are manifest only in transactions with humans, rhesus monkeys, and machines especially built to replicate just their particular mode of sensitivity to photons" (1990: 572–3). (The force of this observation is all the greater if, as seems likely, even the reflectance theory underestimates the complexity of colour psychophysics.)

See also J. J. C. Smart who, it seems to me, got more of this right than he is these days given credit for: "This account of secondary qualities explains their unimportance in

Natural Kinds Come Late

I think natural kind concepts have been getting more of a press than they deserve of late. It's past time to put them in their place; and their place is that of self-conscious and cultivated intellectual achievements. Much of what is currently being written about concepts—by philosophers, but also, increasingly, by psychologists suggests that natural kind concepts are the paradigms on which we should model our accounts of concept acquisition and concept possession at large. This is, I think, hopeless on the face of it. For one thing, as Putnam in particular has argued, natural kind concepts thrive best · ·maybe only in an environment where conventions of deference to experts are in place. But, patently, only creatures with an *antecedently* complex mental life could make a policy of adherence to such conventions. Adherence to conventions of deference couldn't be a precondition for conceptual content in general, if only because deference has to stop somewhere; if *my* ELM concept is deferential, that's because *the botanist's* isn't. Anyhow, it seems just obvious that concepts like STAR in, as one says, the 'technical sense'—the concept of stars that is prepared to defer about the Sun and black dwarfs on the one hand and meteors and comets on the other—come after, and sometimes come to replace, their colloquial counterparts.

As I say, this view flies in the face of the current fashions in developmental cognitive psychology, which stress how early, and how universally, natural kind concepts are available to children. But I find that I'm not much convinced. There is, to be sure, getting to be a lot of evidence (contra Piaget) that young children are deeply into appearance/reality distinctions: they're clear that you can't make a horse into a zebra just by painting on stripes (Keil 1989); and they're clear that, for some categories (animals but not vases, for example), what's on the inside matters to what kind a thing belongs to (Carey 1985). It's usual to summarize such findings as showing that young children are 'essentialists', and if you like to talk that way, so be it.[5] My point, however, is that being an essentialist in this sense clearly does *not* imply having natural kind concepts; not even if a cognitivist picture of concept possession is assumed for sake of the argument. What's further required, at a minimum, is the idea that what's 'inside' (or otherwise hidden) somehow is causally responsible for how

[5] I don't, myself, advise it. Grant that children think that properties that don't appear can matter to whether a thing is a horse. It isn't implied that they think what a bona fide essentialist should: that there are properties (other than being a horse) that *necessitate* a

things that belong to the kind appear; for their 'superficial signs'. It is, of course, an empirical issue, but I don't know of any evidence that children think that sort of thing.

If it's easy to miss the extent to which natural kind concepts are sophisticated achievements, that's perhaps because of a nasty ambiguity in the term. (One that we've already encountered, in fact; it's why I had to pussyfoot about whether they had WATER in the Garden). Consider this dialectic:

> *Did Homer have natural kind concepts?*
> Sure, he had the concept WATER (and the like), and water is a natural kind.

But also:

> *--Did Homer have natural kind concepts?*
> Of course not. He had no disposition to defer to experts about water (and the like); I expect the notion of an expert about water would have struck him as bizarre. And, *of course* Homer had no notion that water has a hidden essence, or a characteristic microstructure (or that anything else does); a fortiori, he had no notion that the hidden essence of water is causally responsible for its phenomenal properties.

A 'natural kind concept' can be the concept of a natural kind; or it can be the concept of a natural kind *as such* (i.e. the concept of a natural kind *as* a natural kind). It's perfectly consistent to claim that Homer had plenty of the first but none of the second. In fact, I think that's pretty clearly true. So the suggestion is that, in the history of science, and in ontogeny, and, for all I know, in phylogeny too, concepts of natural kinds as such only come late. Homer, and children, and animals, have few of them or none. Somehow, concepts of natural kinds as such emerge from a background of concepts of mind-dependent properties, and of concepts of natural kinds that *aren't* concepts of natural kinds as such. Presumably it's because they do somehow emerge from a background of other kinds of concepts that concepts of natural kind as such don't have to be innate.

Fine. So now all I owe you is a story about what "emerging" comes to: and I have to tell this story in a way that an informational semantics can tolerate, viz. without assuming that there is more to concept possession than locking even in the case of bona fide, full-blown, natural kind concepts as such. *Then* I get to go sailing.

I'll start with natural kind concepts and informational semantic and

Natural Kinds and Informational Semantics

We've just distinguished between *merely* having a natural kind concept and having a natural kind concept *as such*. What I'm asking now is *whether an atomistic informational semantics* can honour that distinction. And I'm inviting you to share my concern that, prima facie, it cannot. Prima facie an informational semantics has to say that whether you have the concept WATER is a matter of whether you are locked to water; if you are then you do, and if you aren't then you don't. Whereas (still prima facie) having WATER as a full-blown natural kind concept requires also having, for example, concepts like MICROSTRUCTURE and HIDDEN ESSENCE and NATURAL KIND. Atomism and informational semantics are natural allies, and it's been my strategy throughout to enlist each in the other's service. But maybe we've come to where their joint resources run out. If the possession conditions for full-blown natural-kind-as-such concepts invoke the possession conditions of concepts like NATURAL KIND, then they aren't atomistic.

So, the issue is how an informational semantics should treat full-blown natural kind terms. That's a large topic, and I wish I didn't have to think about it. For what it's worth, however, here's a sketch of a story: whether Homer had the (our) concept WATER doesn't depend on what other concepts he had (on whether he had HIDDEN ESSENCE and MICROSTRUCTURE, for example). Rather, it depends on whether he was locked to water *as such*; or was merely locked to water *in any reasonably nearby world*.

Homer had (and children and animals have) a concept that is locked to water via its familiar phenomenological properties; via its 'superficial signs'. So the locking Homer had was reliable only in worlds where water *has* the familiar phenomenological properties; which is to say only in nomologically possible worlds near ours. That is, I suppose, the usual, pretheoretic way of having a natural kind concept. The kind-constituting property is a hidden essence and you get locked to it via phenomenological properties the having of which is (roughly) nomologically necessary and sufficient for something to instantiate the kind. This explains, by the way, why concepts like WATER exhibit the d/D effect. WATER, like DOORKNOB, is typically learned from its instances; but that's not, of course, because *being water* is mind-dependent. Rather, it's because you typically lock to *being water* via its superficial signs; and, in point of nomological necessity, water samples are the only things around in which those superficial signs inhere.

theoretic. For us (but not for Homer), WATER is a concept whose locking to water is mediated by our adherence to a theory about what water is. Since, by assumption, this theory that we adhere to is true, the locking depends on a property that water has in every metaphysically possible world, not just in nomologically possible worlds that are near here. *We're locked to water via a theory that specifies its essence, so we're locked to water in every metaphysically possible world.* That, I'm suggesting, is what an informational semanticist should say that it *is* to have a concept of a natural kind *as* a natural kind: it's for the mechanism that effects the locking not to depend on the superficial signs of the kind, and hence to hold (*ceteris paribus* of course) even in possible worlds where members of the kind lacks those signs.

So, does this, or doesn't it, amount to Homer's having had the same concept of water that we do? Did they or didn't they have the concept WATER back in the Garden?

Actually, I don't much care which you say, so long as you like the general picture. Suffice it that it's quite in the spirit of informational semantics to decide to talk like this: Homer did have the concept WATER (he had a concept that is nomologically linked to *being water*) and, of course, *being water* isn't a mind-dependent property. So Homer had a concept of a natural kind. But WATER wasn't, for Homer, a concept of a natural kind *as such*; and for us it is. We're locked to *being water* via a chemical-cum-metaphysical theory, that specifies its essence, and that is quite a different mechanism of semantic access from the ones that Homer relied on. In particular, the two ways of locking to water support quite different counterfactuals. This shows up (*inter alia*) in the notorious thought experiments about Twin-Earth: we think that XYZ wouldn't be water; Homer wouldn't have understood the question.

But an entirely informational and atomistic semantics can also do justice to the intuition that Homer had the same WATER concept as ours. All the metaphysics of concept possession requires, of our concept WATER or Homer's, is being locked to water. If you are locked to water our way, you have the concept WATER as a natural kind concept; if you are locked to concept WATER Homer's way, you have the concept WATER, but not *as* a natural kind concept. But, on a perfectly natural way of counting, if you are locked to water either way, you have the concept WATER. (I suppose that God is locked to *being water* in still a third way; one that holds in every metaphysically possible world but *isn't* theory-mediated. That's OK with informational semantics; God can have the concept WATER too. He can't, however, have the *pretheoretic* concept

If you're lucky, you can have concepts of natural kinds on the cheap. Homer maybe didn't need much to get WATER locked to *water*; maybe all he needed was innate detectors for the phenomenological properties which, in point of nomological necessity, water has in all the worlds near to him (and us). But, of course, you only get what you pay for: Homer didn't have the concept of water *as* a natural kind concept. To have that, he would need to have been locked to the essence of water *via* the essence of water; that is, in a way that doesn't depend on water's superficial signs. Probably, de facto, all such lockings (except God's) are theory-mediated; indeed, they are perhaps all *meta*theory-mediated; they may well depend, de facto, on having not just concepts of natural kinds, but also the concepts NATURAL KIND and HIDDEN ESSENCE. Which *nobody* did until quite recently.

But I want to emphasize what I take to be a main moral of the discussion: the 'de facto' matters. Just as IA says there are *no* concepts the possession of which is metaphysically necessary for having WATER (except WATER), so I'd like it to say that there are no concepts the possession of which is metaphysically necessary for having WATER *as a natural kind concept* (except WATER); *all* that's required is being locked to water in a way that doesn't depend on its superficial signs. But, of course, metaphysically necessary is one thing, on the cards is quite another. I'm quite prepared to believe that, de facto, until we had (indeed, had more or less self-consciously), the concepts that cluster around NATURAL KIND, there was probably no way that we *could* link to WATER except the sort of way that Homer did and children and animals do; viz. via water's metaphysically accidental but nomologically necessary properties. But now we have a theory that tells us what water is, and we are linked to water via our acceptance of that theory. Science discovers essences, and doing science thereby links us to natural kinds *as such*.

I think, by the way, that the ethological analogies play out quite nicely on this sort of analysis. It's natural and handy and, for most purposes harmless, to say that ducklings have the concept MOTHER DUCK innately; that male sticklebacks have the concept CONSPECIFIC RIVAL innately, and so on. And it's quite true in such cases that, given normal experience, the creatures end up locked to the properties that these concepts express. So, as far as informational semantics is concerned, they therefore end up having concepts that have these properties as their contents. But, in fact, the innate endowment that they exploit in doing so is quite rudimentary. Male sticklebacks get locked to *conspecific rivalhood* via not much more than an innate ability to detect red spots. To do so,

This is nomologically necessary (anyhow, it's counterfactual supporting) in the stickleback's ecology, and nomological necessity is transitive. So sticklebacks end up locked to *conspecific rivalhood* via *one of its reliable appearances*.

To repeat: informational semantics suggests that, so far as the requisite innate endowment is concerned, if the world co-operates you can get concepts of natural kinds very cheap. That's what the sticklebacks do; it's what Homer did; it's what children do; it's what all of us grown-ups do too, most of the time. By contrast, for you to have a natural kinds concept *as such* is for your link to the essence of the kind *not* to depend on its inessential properties. This is a late and sophisticated achievement, historically, ontogenetically, and phylogenetically, and there is no reason to take it as a paradigm for concept possession at large. I suppose you start to get natural kind concepts in this strong sense only when it occurs to you that, if generality and explanatory power are to be achieved, similarity and difference in respect of how things affect minds like ours has sometimes got to be ignored in deciding what kinds of things they are; perhaps, de facto, this happens only in the context of the scientific enterprise.

Well, what about the 'technical' concept WATER? Does that have to be innate if it's primitive?

Of course not. For one thing, on the present view, there really is no 'technical concept water'; there's just, as it were, the technical way of having the concept WATER. Once you've got a concept that's locked to water via its (locally reliable) phenomenological properties, you can, if you wish, make a project of getting locked to water in a way that doesn't depend on its superficial signs. The easy way to do this is to get some expert to teach you a theory that expresses the essence of the kind. To be sure, however, that will only work if the natural kind concept that you're wanting to acquire is one which somebody else has acquired already. Things get a deal more difficult if you're starting *ab initio*; i.e. without any concepts which express natural kinds as such. It's time for me to tell my story about how concepts of natural kinds might "emerge" in a mind that is antecedently well stocked with concepts of other kinds. Actually, it's a perfectly familiar story and not at all surprising.

'Emerging'

which, however, aren't concepts of natural kinds as such.[6] Then what you need to do to acquire a natural kind concept as a natural kind concept *ab initio* is: (i) construct a true theory of the hidden essence of the kind; and (ii) convince yourself of the truth of the theory. If the theory is true, then it will say of a thing that it is such-and-such when and only when the thing is such-and-such; and if you are convinced of the truth of the theory, then you will make it a policy to consider that a thing is such-and-such when and only when the theory says that it is. So your believing the theory locks you to such-and-suches via a property that they have in every metaphysically possibly world; namely, the property of being such-and-suches; the property that makes the theory true. The upshot is that, if the moon is blue, and everything goes as planned, you will end up with a full-blown natural kind concept; the concept of *such-and-suches as such*.

Aha, but how do you go about constructing a true theory of the essence of such-and-suches and convincing yourself that it is true? How do you do it in, say, the case of being water*?*

Oh, well, you know: you have to think up a theory of what water is that both explains why the superficial signs of *being water* are reliable and has the usual theoretical virtues: generality, systematicity, coherence with your other theories, and so forth. You undertake to revise the theory when what it says about water isn't independently plausible (e.g. independently plausible in light of experimental outcomes); and you undertake to revise your estimates of what's independently plausible (e.g. your estimates of the construct validity of your experimental paradigms) when they conflict with what the theory says about water. And so on, round and round the Duhemian circle.

In short, you do the science. I suppose the Duhemian process of scientific theory construction is possible only for a kind of creature that antecedently has a lot of concepts of properties that are mind-dependent, and a lot of natural kind concepts that aren't concepts of natural kinds as such. And it's also only possible for a kind of creature that is able to pursue policies with respect to the properties that it locks its concepts to. Probably, we're the only kind of creature there is that meets these conditions. Which explains, I suppose, why we're so lonely.

[6] As I remarked in Chapter 6, I rather suspect that these, together with the concepts of natural kinds as such, exhaust the sorts of concepts that there are; but I don't know how to argue that they do.

Notice, in any case, that this is a mixed taxonomy. The distinction between concepts of mind-dependent properties and the rest is ontological; mind-dependence is a property of the property that a concept expresses. By contrast, the distinction between natural-kind-as-

A natural kind enters into lots of nomic connections to things other than our minds. We can validate a theory of the kind with respect to those connections because the theory is required to predict and explain them. You can't follow this Duhemian path in the case of DOORKNOB, of course, because there is nothing to validate a theory of doorknobs against except how things strike us. In effect, what strikes us as independently plausibly a doorknob *is* a doorknob; the mind-dependence of *doorknob-hood* is tantamount to that. The more we learn about what water is, the more we learn about the world; the more we learn about what doorknobs are, the more we learn about ourselves. The present treatment implies this and, I think, intuition agrees with it. At least, Realist intuition does.

We do science when we want to lock our concepts to properties that *aren't* constituted by similarities in how things strike us. We do science when we want to reveal the ways that things would be similar *even if we weren't there*. Idealists to the contrary not withstanding, there's no paradox in this. We can, often enough, control for the effects of our presence on the scene in much the same ways that we control for the effects of other possibly confounding variables. To be sure, here as elsewhere, the design of well-confirmed theories is hard work and often expensive. And the only recompense is likely to be the cool pleasure of seeing things objectively; seeing them as they are when you're not looking. Objectivity is an educated taste, much like Cubism. Maybe it's worth what it costs and maybe it's not. It's entirely your choice, of course. Far be it from me to twist your arm.

So much, then, for how we got from the Garden to the laboratory. It is, as I say, quite a familiar story.

Short Summary

You aren't actually required to believe any of what's in this chapter or the last; I have mostly just been exploring the geography that reveals itself if conceptual atomism is taken seriously. Still, I do think our cognitive science is in crisis, and that we're long overdue to face the dilemma that confronts it. On one hand, everybody knows, deep down, that Inferential Role Semantics makes the problem of concept individuation intractable. And, on the other hand, everybody gags on Informational Atomism. (Well, *practically* everybody does.) And nobody seems to be able to think of any other alternatives. Probably that's because those are all the alternatives that there are.

sweeten the pill. It seemed to me, for a long while, that a cost of atomism would be failing to honour the distinction between theoretical concepts and the rest. For, surely, theoretical concepts are ones that you have to believe a theory in order to have? And, according to conceptual atomism, there are *no* concepts that you have to believe a theory in order to have. But it doesn't seem to me that way now. A theoretical concept isn't a concept that's *defined* by a theory; it's just a concept that is, de facto, locked to a property via a theory. Informational Atomism doesn't mind that *at all*, so long as you keep the "de facto" in mind.

Likewise, it used to seem to me that atomism about concepts means that DOORKNOB is innate. But now I think that you can trade a certain amount of innateness for a certain amount of mind-dependence. *Being a doorknob* is just: striking our kinds of minds the way that doorknobs do. So, what you need to acquire the concept DOORKNOB "from experience" is just: the kind of mind that experience causes to be struck that way by doorknobs. The price of making this trade of innateness for mind-dependence is, however, a touch of Wotan's problem. It turns out that much of what we find in the world is indeed "only ourselves". It turns out, in lots of cases, that we *make things be of a kind* by being disposed to *take them to be of a kind*.

But not in every case; not, in particular, in the case of kinds of things that are alike in respect of the hidden sources of their causal powers, regardless of their likeness in respect of their effects on us. To describe it in terms of *those* sorts of similarities is to describe the world the way that God takes it to be. Doing science is how we contrive to cause ourselves to have the concepts that such descriptions are couched in. Not philosophy but science is the way to get Wotan out of his fly-bottle. That story seems to me plausibly true; and it is, as we've seen, compatible with an informational and atomistic account of the individuation of concepts. But dear me, speaking of fly-bottles, how Wittgenstein would have loathed it; and Wagner and Virginia Woolf too, for that matter. Well, you can't please everyone; I'll bet it would have been all right with Plato.

Short Conclusion: A Consolation for Philosophers

That's really the end of my story; but a word about what I think of as the Luddite objection to conceptual atomism is perhaps in order.

It's natural, *pace* Appendix 5A, to suppose that conceptual atomism means that there are no conceptual truths, hence that there are no analytic

analyticity' entailed not just 'no analyses' but 'no analytic philosophy' as well. Technological unemployment would then begin to threaten.

But I guess I'm not inclined to take that prospect very seriously; certainly I'm not one of those end-of-philosophy philosophers. If, there aren't any conceptual analyses, the moral isn't that we should stop doing philosophy, or even that we should start doing philosophy in some quite different way. The moral is just that we should stop saying that conceptual analysis is what philosophers do. If analytic philosophers haven't been analysing concepts after all, at least that explains why there are so few concepts that analytic philosophers have analysed.

I guess what I really think is that philosophy is just: whatever strikes minds like ours as being of the same kind as the prototypical examples. But maybe that's wrong; and, if it is, then maybe if we were to stop saying that philosophy is conceptual analysis, that would leave philosophers without a defensible metatheory. Well, if so, so be it. We wouldn't be worse off in that respect than doctors, lawyers, dentists, artists, physicists, chicken sexers, psychologists, driving instructors, or the practitioners of any other respectable discipline that I can think of. Why should philosophers be exempt from this practically universal predicament? There are many classes of performances in which intelligence is displayed, but the rules or criteria of which are unformulated. Efficient practice precedes the theory of it; methodologies presuppose the application of the methods, of the critical investigation of which they are the products . . . It is therefore possible for people intelligently to perform some sorts of operations when they are not yet able to consider any propositions enjoining how they should be performed.

But, bless me, it seems that I am quoting from *The Concept of Mind*.[1] I'm sure that means that it's time for me to stop.

APPENDIX 7A
Round Squares

I want briefly to consider an ontological worry about IA that's relatively independent of the main issues that this chapter is concerned with.

It seems pretty clear that IA is going to have to say that it's *metaphysically* impossible for there to be a primitive concept of a self-contradictory property; e.g. a primitive concept ROUND SQUARE. (Remember that "ROUND SQUARE" is a name, not a structural description. The notation leaves it open whether the corresponding

concept is atomic.) How the argument goes will depend on the details of IA's formulation. But, roughly: IA says that concepts have to be locked to properties. Maybe it's OK for a concept to lock to a property that exists but happens not to be instantiated (like *being a gold mountain*), but presumably there isn't any property of *being a round square* for the *necessarily* uninstantiated concept ROUND SQUARE to lock to.

That's all right if ROUND SQUARE is assumed to be complex; it's pretty plausible that there really isn't anything to having ROUND SQUARE beyond the inferential dispositions that its compositional semantics bestows (viz. the disposition to infer ROUND and SQUARE). But the corresponding *primitive* concept would have *neither* content (there's no property for it to lock to) *nor* compositional structure (it has no constituents), so there could be nothing to having it at all. The objection is that it's not obvious that it's metaphysically necessary that ROUND SQUARE couldn't be primitive.

A possible reply is that it's also not obvious that it could, so all you get is a hung jury. But I think maybe we can do a little better. Consider a *non*-self-contradictory property like *being a red square*. It's common ground for any RTM that there is a *complex* concept of this property (constructed from the concepts RED and SQUARE). But it's built into *informational* versions of RTM that it also allows there to be a *simple* concept of this property; viz. a *primitive* mental representation REDSQUARE (*sic*; this *is* intended to be a structural description) that is locked to *being red and square*. Presumably, one could acquire REDSQUARE ostensively. That is, one could get locked to *being red and square* (not by first getting locked to *being red* and *being square*, but) by learning that redsquares (*sic*) are the things that look like *those*. So Informational Atomism acknowledges the metaphysical possibility of having the concept of a red square without having either the concept RED or the concept SQUARE. (You won't, of course, admit that RED SQUARE could be, in this sense, primitive if you boggle at concepts without conceptual roles. But if you boggle at concepts without conceptual roles you can't accept a pure informational semantics *at all*, so why should you care what a pure informational semantics says about concepts of self-contradictory properties?)

If, on the other hand, you find it intuitively plausible that there *are* two ways of having a concept of a red square (viz. RED SQUARE, which you can't have unless you've got RED and SQUARE, and REDSQUARE, which you can because it's primitive) then everything is OK about IA's treatment of the concept ROUND SQUARE. For the (anyhow, my)

intuition is very strong that there is only *one* way to have *that* concept. In particular, that there is no concept of a round square that one could have without also having ROUND and SQUARE. If you share the intuition that there is this asymmetry, between RED SQUARE and ROUND SQUARE, then you should be very happy with IA. IA *explains* the asymmetry because it entails that there can be no *primitive* concept without a corresponding property for it to lock to.