

Innateness and Ontology, Part I: The Standard Argument¹

I find only myself, every time, in everything I create.

Wotan in *Die Walküre*, Act II

Are you also puzzled, Socrates, about cases that might be thought absurd, such as hair or mud or dirt or any other trivial and undignified objects. Are you doubtful whether or not to assert that each of these has a separate form? . . . Not at all, said Socrates. In these cases, the things are just the things we see; it would surely be too absurd to suppose that they have a form.

Plato, *Parmenides*

Virginia Woolf has summed up this state of things with perfect vividness and conciseness in the words, 'Tuesday follows Monday'.

E. M. W. Tillyard, *Shakespeare's Last Plays*

Introduction

RTM requires there to be infinitely many concepts that are complex and finitely many that are primitive. RTM also requires concepts to have their contents essentially. The versions of RTM that are currently standard in philosophy and in cognitive science, however, want still more: most lexical concepts should *not* be primitive, and the content of concepts should be determined, at least *inter alia*, by their inferential-cum-causal relations to one another. I think, however, that the evidence is getting pretty solid that the last two conditions can't be met; lexical concepts typically don't act as though they were internally structured by either psychological or linguistic

¹ This chapter reconsiders some issues about the nativistic commitments of RTMs that I first raised in Fodor 1975 and then discussed extensively in 1981a. Casual familiarity with the latter paper is recommended as a prolegomenon to this discussion.

test. And the question *which* aspects of a concept's inferential role are the ones that determine its meaning appears to be hopeless. Thus far has the World Spirit progressed.

I propose, therefore, that we scrap the standard versions of RTM and consider, in their place, a doctrine that I'll call Informational Atomism. (IA for short.) IA has an informational part and it has an atomistic part. To wit:

Informational semantics: content is constituted by some sort of nomic, mind-world relation. Correspondingly, having a concept (concept possession) is constituted, at least in part, by *being in* some sort of nomic, mind-world relation.

Conceptual atomism: most lexical concepts have no internal structure.

As far as I can tell, nobody but me thinks that IA has a prayer of being true; not even people who are quite sympathetic to RTM. Now, why is that, do you suppose?

I can imagine three objections to IA (however, see Appendix 7A). The first of these I'm prepared not to take very seriously, but the second two need some discussion. Most of this chapter and the next one are devoted to them. I should say at the outset that I regard what follows as very tentative indeed. Though the standard versions of RTM have been explored practically to death, IA is virgin territory. The best I hope for is a rough sketch of the geography.

First objection: If atomism is true and most lexical concepts have no internal structure, then there is no such thing as the analysis of most of the concepts that philosophers care about. That BROWN COW has a philosophical analysis (into BROWN and COW) isn't much consolation.

Reply: Strictly speaking, you can have conceptual analysis without structured concepts since, strictly speaking, you can have analyticity without structured concepts (see Appendix 5A). You do, however, have to live with the failure of attempts to reduce analyticity to conceptual containment. And you have to live with the general lack of empirical sanction for claims that satisfying the possession conditions for some concept *A* requires satisfying the possession conditions for some other concept *B*. As far as I can tell, there is little or no evidence for such claims except brute appeals to intuition; and, as we saw in Chapter 4, a case can be made that the intuitions thus appealed to are corrupt.

On the other hand, who cares about conceptual analysis? It's a commonplace that its successes have been, to put it mildly, very sparse. Indeed, viewed from the cognitive psychologist's perspective, the main

quotidian concepts, its answers to 'What is their content?' and to 'How do you acquire them?' are, respectively, 'It has none' and 'You don't'. It's worth bearing in mind that analytic philosophy, from Hume to Carnap inclusive, was a *critical* programme. For the Empiricists, the idea was to constrain the conditions for concept possession a priori, by constraining the acceptable relations between concepts and percepts. It would then turn out that you really don't have many of the concepts that you think you have; you don't have GOD, CAUSE, or TRIANGLE at all, and though perhaps you do have DOG, it's not the sort of concept that you had supposed it to be. "When we run over the libraries, persuaded of these principles, what havoc must we make?" (Hume 1955: 3.) Post-Positivist philosophical analysis has wavered between reconstruction and deconstruction, succeeding in neither. Most practitioners now hold that we do have DOG, CAUSE, and TRIANGLE after all; maybe even GOD. But they none the less insist that there are substantive, a priori, epistemological constraints on concept possession. These, in the fullness of time, analysis will reveal; to the confusion of Sceptics, Metaphysical Realists, Mentalists, Cartesians, and the like. Probably of Cognitive Scientists too.

But, between friends: nothing of the sort is going to happen. In which case, what's left to a notion of conceptual analysis that's detached from its traditional polemical context? And what on earth are conceptual analyses for?

Second objection: The informational part of IA says that content is constituted by nomic symbol-world connections. If that is true, then there must be laws about everything that we have concepts of. Now, it may be there are laws about *some* of the things that we have concepts of (fish, stars, grandmothers(?)). But how could there be laws about, as it might be, *doorknobs*?² Notice that it's only in conjunction with conceptual atomism that informational semantics incurs this objection. Suppose the concept DOORKNOB is definitionally equivalent to the complex concept . . . *ABC* . . . Then we can think the former concept if there are laws about each of the constituents of the latter. In effect, all informational semantics per se requires for its account of conceptual content is that there be laws about the properties expressed by our *primitive* concepts. However, IA says that practically every (lexical) concept is primitive. So, presumably, it says that DOORKNOB is primitive.³ So there must be laws about doorknobs

² For discussions that turn on this issue, see Fodor 1986; Antony and Levine 1991; Fodor 1991.

³ Actually, of course, DOORKNOB isn't a very good example, since it's plausible a

qua doorknobs, as it were, not qua *ABCs*. But how could there be laws about doorknobs? Doorknobs, of all things!

Third objection: If most lexical concepts have no internal structure, then most lexical concepts must be primitive. But primitive concepts are, *ipso facto*, unlearned; and if a concept is unlearned, then it must be innate. But how could DOORKNOB be innate? *DOORKNOB*, of all things!! Prima facie, this objection holds against (not just IA but) any version of RTM that is not heavily into conceptual reduction; that is, against any theory that says that the primitive conceptual basis is large. In particular, it holds prima facie against any atomistic version of RTM, whether or not it is informational.

Objections two and three both turn on the peculiarly central roles that primitive concepts play in RTMs. Primitive concepts are supposed to be the special cases that problems about conceptual content and concept acquisition reduce to. But if not just RTM but also conceptual atomism is assumed, then the special case becomes alarmingly general. If, for example, DOORKNOB is primitive, then whatever metaphysical story we tell about the *content* of primitive concepts has to work for DOORKNOB. And so must whatever psychological story we tell about the *acquisition* of primitive concepts. And the metaphysical story has to work in light of the acquisition story, and the acquisition story has to work in light of the metaphysical story. Hume wouldn't have liked this at all; he wanted the primitives to be just the sensory concepts, and he wanted them to be acquired by the stimulation of an innate sensorium. Pretty clearly, he gets neither if DOORKNOB is among the primitives.

I propose, in this chapter, to explore some of the ways that these issues play out in IA versions of RTM. We'll consider how, because of the way it construes conceptual content, IA is maybe able to avoid some extremes of conceptual nativism to which other atomistic versions of RTM are prone. (Though at a price, to be sure. No free lunches here either.) In Chapter 7, I'll take up the question about laws.

The Standard Argument

There is a plausible argument which says that informational atomism implies radical conceptual nativism; I'll call it the 'Standard Argument' (SA). Here, in very rough form, is how the Standard Argument is supposed to go.

SA begins by assuming that learning a concept is an *inductive* process;

relatively unproblematic when the concept to be acquired is a definition. If the concept BACHELOR is the concept UNMARRIED MALE, you can learn BACHELOR by learning that things fall under it in virtue of *being male* and *being unmarried*. But, on pain of circularity, the (absolutely) primitive concepts can't themselves be learned this way. Suppose the concept RED is primitive. Then to learn RED inductively you'd have to devise and confirm the hypothesis that things fall under RED *in virtue of being red*. But you couldn't devise or confirm that hypothesis unless you already had the concept RED, *since the concept RED is invoked in the formulation of the hypothesis*. So you can't have learned the concept RED (or, *mutatis mutandis*, any other primitive concept) inductively, by hypothesis testing and confirmation. But SA assumes that induction is the only sort of concept learning that there is. So it follows that you can't have learned your primitive concepts at all. But if you have a concept that you can't have learned, then you must have it innately. So the Standard Argument says. What, if anything, is wrong with this?

To begin with, it might be replied that the inductive account of concept acquisition is plausible only assuming a cognitivist account of concept possession; an account of concept possession according to which having a concept is *knowing something*. This assumption is natural enough if you are thinking of concepts on the model of definitions (*/stereotypes/theories*): having a concept is knowing what its definition (*/stereotype/theory*) is. By contrast, IA is explicitly *non-cognitivist* about concept possession; it says that having a concept is (not knowing something but) being in a certain nomic mind-world relation; specifically, it's being in that mind-world relation in virtue of which the concept has the content that it does. This changes the geography in ways that may be germane to the present issues. Because it is non-cognitivist about concept possession, IA invites a correspondingly non-cognitivist account of how concepts are acquired. That might be just what you're looking for if you're looking for a way out of SA.

Avoiding nativism by endorsing a non-cognitivist view of concept possession is, of course, hardly a new idea. At least since Ryle (1949), a lot of philosophical ink has been invested in the thought that having a concept is knowing how, not knowing that. Correspondingly, concept acquisition is arguably *learning how*, rather than *learning that*, and it isn't obvious that learning how needs to be inductive. Maybe construing concept possession as know-how is all that avoiding SA requires. I think philosophers quite generally find this plausible.

appears that how-learning itself depends on that-learning.⁴ For example, my linguist friends tell me that learning how to talk a first language requires quite a lot of learning that the language has the grammar that it does. I tell my linguist friends that my philosophy friends tell me that it is a priori and necessary that this cannot be so. Then my linguist friends laugh at me. What am I to do?

And, for another thing, whatever the *general* story about knowing how and knowing that may be, the particular skills that concept possession is usually supposed to implicate are *perceptual* and *inferential*, and these look to be just saturated with knowing that. Surely, you can't identify a dog by its barking unless you know(/believe) that dogs bark. Surely, you won't infer from *dog* to *animal* unless you know(/believe) that dogs are animals. Indeed, in the second case, opposing knowing how to knowing that looks like insisting on a distinction without a difference.⁵

Where we've got to is: even if it's supposed that concepts are skills,⁶ very little follows that helps with avoiding SA. That's because to avoid SA you need a non-cognitivist view of concept possession. And supposing concepts to be skills doesn't guarantee a non-cognitivist view of concept possession, because it is perfectly possible to be a cognitivist *about the possession of skills*, if not in every case, then at least in the case of the skills that concept possession requires. The moral: *it's unclear that Ryle can deny SA the premiss that it centrally requires, viz. that concept acquisition is mediated by hypothesis formation and testing*.

But IA can. Let's see where this leads.

Following Loewer and Rey (1991a) (who are themselves following the usage of ethologists) I'll say that acquiring a concept is getting *nomologically locked* to the property that the concept expresses. So, then, consider a supplemented version of IA (I'll call it SIA) which says everything that IA does and also that concept possession is some kind of locking. The question before us is whether SIA requires radical nativism.

⁴ That learning how can't depend on learning that in *every* case is, I suppose, the moral of Lewis Carroll's story about Achilles and the tortoise: Carroll 1895/1995.

⁵ CogSci footnote: the present issue isn't whether inferential capacities are 'declarative' rather than 'procedural'; it's whether they are interestingly analogous to skills. A cognitive architecture (like SOAR, for example) that is heavily committed to procedural representations is not thereby required to suppose that drawing inferences has much in common with playing basketball or the piano. Say, if you like, that someone who accepts the inference from *P* to *Q* has the habit of accepting *Q* if he accepts *P*. But this sort of 'habit' involves a relation among one's propositional attitudes and, *prima facie*, being able to play the piano doesn't.

Notice that the question before us is *not* whether SIA *permits* radical nativism; it's patent that it does. According to SIA, having a concept is being locked to a property. Well, being locked to a property is having a disposition, and though perhaps there are some dispositions that must be acquired, hence can't be innate, nothing I've heard of argues that being locked to a property is one of them. If, in short, you require your metaphysical theory of concept possession to *entail* the denial of radical nativism, SIA won't fill your bill. (I don't see how any metaphysics could, short of question begging, since the status of radical nativism is surely an empirical issue. Radical nativism may be false, but I doubt that it is, in any essential way, *confused*.) But if, you're prepared to settle for a theory of concepts that is plausibly *compatible* with the denial of radical nativism, maybe we can do some business.

If you assume SIA, and hence the locking model of concept possession, you thereby deny that learning concepts necessarily involves acquiring beliefs. And if you deny that learning concepts necessarily involves acquiring beliefs, then you can't assume that hypothesis testing is an ingredient in concept acquisition. It is, as I keep pointing out, primarily cognitivism about the metaphysics of concept possession that motivates inductivism about the psychology of concept acquisition: hypothesis testing is the natural assumption about how *beliefs* are acquired from experience. But if it can't be assumed that concept acquisition is *ipso facto* belief acquisition, then it can't be assumed that locking DOORKNOB to *doorknobhood* requires a mediating hypothesis. And if it can't be assumed that locking DOORKNOB to *doorknobhood* requires a mediating hypothesis, then, a fortiori, it can't be assumed that it requires a mediating hypothesis in which the concept DOORKNOB is itself deployed. In which case, for all that the Standard Argument shows, DOORKNOB could be *both* primitive *and* not innate.

This maybe starts to sound a little hopeful; but not, I'm afraid, for very long. The discussion so far has underestimated the polemical resources that SA has available. In particular, there is an independent argument that seems to show that *concept acquisition has to be inductive, whether or not the metaphysics of concept possession is cognitivist*; so SA gets its inductivist premiss even if SIA is right that having a concept doesn't require having beliefs. The moral would then be that, though a non-cognitivist account of concept possession may be necessary for RTM to avoid radical nativism, it's a long way from being sufficient.

In short, Patient Reader, the Standard Argument's way of getting radical nativism goes like this:

hypothesis-testing) model of concept *learning* → (3) primitive concepts can't be learned.

SIA denies (1), thereby promising to block the standard argument. If, however, there's some *other* source for (2)—some plausible premiss to derive it from that doesn't assume a cognitivist metaphysics of concept possession—then the standard argument is back in business.

And there is. Here's a narrowly based argument for the hypothesis-testing model of concept acquisition; one that presupposes neither a cognitivist account of concept possession nor even any general inductivist thesis about the role of hypothesis testing in the acquisition of empirical beliefs.

Nobody, radical nativists included, doubts that what leads to acquiring a concept is typically *having the right kinds of experiences*. That experience is *somehow* essentially implicated in concept acquisition is common ground to both Nativists and Empiricists; their argument is over whether concepts are abstracted from, or merely occasioned by, the experiences that acquiring them requires. That this is indeed the polemical situation has been clear to everybody concerned (except the Empiricists) at least since Descartes. In short, SIA, like everybody else, has to live with the fact that it's typically acquaintance with doorknobs that leads to getting locked to doorknobhood. So, like everybody else, SIA has to explain why it's those experiences, and not others, that eventuate in locking to that property. *But that's enough, all by itself, to make the search for a non-inductivist account of concept acquisition look pretty hopeless*. For, even if a cognitivist model of concept possession is *not* assumed, the hypothesis-testing story has the virtue of solving what I'll call the doorknob/DOORKNOB problem:⁷ why is it so often experiences of doorknobs, and so rarely experience with whipped cream or giraffes, that leads one to lock to *doorknobhood*?

According to the hypothesis-testing model, the relation between the content of the concepts one acquires and the content of the experiences that eventuate in one's acquiring them is *evidential*; in particular, it's mediated by content relations between a hypothesis and the experiences that serve to confirm it. You acquire DOORKNOB from experience with doorknobs because you use the experiences to confirm a hypothesis about the nature of *doorknobhood*; and doorknobs, unlike giraffes or whipped cream, are *ceteris paribus* a good source of *evidence* about the nature of doorknobs. Come to think of it, one typically gets DOORKNOB from

⁷ I had thought at first that I would call this the fire hydrant/FIRE HYDRANT

experience with *good* or *typical* examples of doorknobs, and good or typical doorknobs are a *very good* source of evidence about doorknobs. I'll return to this presently.

If, by contrast, you assume that, in the course of concept acquisition, the relation between the eliciting experience and the concept acquired is *not* typically evidential—if, for example, it's just 'brute causal' (for this terminology, see Fodor 1981a)—then why *shouldn't* it be experience with giraffes that typically eventuates in locking to doorknobhood? Or vice versa? Or both? It appears there's more to be said for the hypothesis-testing model of concept acquisition than even SA had supposed.⁸ Compare a proposal that Jerry Samet once made for avoiding the assumption that hypothesis testing mediates concept acquisition (and hence for avoiding the Standard Argument): perhaps concepts are not learned but 'caught', sort of like the flu (Samet 1986). No doubt this suggestion is a bit underspecified; the 'sort of' does all the work. But there's also a deeper complaint: it's left wide open why you generally catch DOORKNOB from doorknobs and not, as it might be, from using public telephones (again sort of like the flu).

UnDarwinian Digression

At this point in the dialectic, there's a strong temptation to dump the load on Darwin; a standard tactic, these days, when a philosopher gets in over his head. Suppose that the mechanism of concept acquisition is indeed non-cognitivist; suppose, for example, that it's some sort of triggering. Still, wouldn't a mechanism that triggers the concept *X* consequent upon experience with *X*s be more of a help with surviving (or getting reproduced, or whatever) than, say, a mechanism that triggers the concept *X* consequent upon encounters with things that aren't *X*s? If so, then maybe SIA together with not-more-than-the-usual-amount of handwaving about Darwin might after all explain why the relation between the content of experiences and the content of the concepts they eventuate in locking to is so rarely arbitrary.

⁸ *Linguistic footnote:* as far as I can tell, linguists just take it for granted that the data that set a parameter in the course of language learning should generally bear some natural, unarbitrary relation to the value of the parameter that they set. It's hearing *sentences without subjects* that sets the null subject parameter (maybe); what could be more reasonable? But, on second thought, the notion of triggering as such, unlike the notion of hypothesis testing as such, requires no particular relation between the state that's acquired and the experience that occasions its acquisition. In principle *any* trigger could set *any* parameter. So, *prima facie*, it is an embarrassment for the triggering theory if the grammar that the child acquires

Well, maybe. But, of course, that's cold comfort if what you want is a *non-nativist* version of SIA. You can only trigger a concept that's there, genetically specified, waiting to be triggered. So the Darwinian/ethological story about concept acquisition does no better than the old-fashioned hypothesis-testing story at making DOORKNOB not be innate. Out of one frying pan but into another; ethologists are nativists by definition.

And, anyhow, even if the doorknob/DOORKNOB relation is selected for by evolution, what, if not inductive learning, could be the mechanism by which it is implemented? If concept acquisition isn't inductive, then *how* does Mother Nature contrive to insure that it *is* instances of *F-ness* (and not of *G-ness*) that trigger the concept *F* in the course of ontogeny? After all, if Mother N wants to select for the doorknob/DOORKNOB type of relation between concepts and their experiential causes, she has to do so by selecting *a mechanism that produces* that relation between one's concepts and their causes. This is a special case of the entirely general truth that whenever Mother N wants to select for *any* phenotypic property she has to do so by selecting a proximal mechanism that produces it. The obvious candidate to select if one wants to ensure that concept acquisition exhibits the d/D relation is inductive learning. But we have it on independent grounds that primitive concepts can't be learned inductively. There may be a way for a conceptual atomist to get out of this dilemma, but waving his hands about Darwin certainly isn't it.

The preliminary moral, anyhow, is that radical nativism is very hard for a conceptual atomist to avoid. If he starts out thinking about concept acquisition the way Empiricists do—as a kind of hypothesis testing—radical concept nativism follows; and if he starts out thinking about concept acquisition the way that ethologists do—as a kind of triggering—radical concept nativism still follows. It looks like a conceptual atomist ends up being a radical concept nativist pretty much however he starts out thinking about concept acquisition. So maybe conceptual atomism is just false.

Or maybe radical concept nativism is true, despite its wide unpopularity in the philosophical community. Speaking just as a private citizen, I've always sort of thought it wouldn't be all that surprising if radical concept nativism did turn out to be true. So it didn't much embarrass me that all the roads from concept atomism seemed to lead there. It is, after all, God and not philosophers who gets to decide what creatures have genotypically built in. That is surely *much* the best arrangement from the creature's point of view.

So, in any case, it seemed to me in 1975 or so. But maybe this relaxed

have an apparently respectable argument that they *must* be learned inductively: nothing else appears likely to account for the content relation between the concept that's acquired and the experience that mediates its acquisition. But look, it can't be that inductivism about the acquisition of primitive concepts is *both* circular *and* mandatory.

Please note that, though this is an embarrassment for those of us who are inclined towards atomism, it is also an embarrassment for those of you who aren't. For, whatever you may think about the *size* of the primitive conceptual basis—and, in particular, about whether DOORKNOB is in it—on any version of RTM *some* concepts are going to have to be primitive. And, on the one hand, SA does seem to show that primitive concepts can't be acquired inductively. And, on the other hand, whatever the primitive concepts are, their acquisition is pretty sure to exhibit the familiar d/D relation between the content of the concept and the content of the experience that occasions it. Of what concept does the acquisition not?⁹

In fact, it's the concepts that have traditionally been practically everybody's favourite candidates for being primitive that exhibit the doorknob/DOORKNOB effect most clearly. Like RED, for example. To be sure, philosophers of both the Cartesian and the Empiricist persuasion have often stressed the *arbitrariness* of the relation between the content of sensory concepts and the character of their causes. It's bumping into photons (or whatever) that causes RED; but RED and PHOTON couldn't be less alike in content. (According to Descartes, this shows that not even sensory concepts can come from experience. According to Locke, it shows that secondary qualities are mind-dependent.) Well, if the relation between sensory concepts and their causes really is arbitrary, then there can be no d/D problem about sensory concepts. In which case, if Empiricists are right and *only* sensory concepts are primitive, everything turns out OK. Sensory concepts don't have to be learned inductively, so they can be innate; just as the Standard Argument requires, and just as Empiricists and Rationalists have both always supposed them to be. Empiricism would be cheap at the price if it shows the way out of a foundational paradox about RTM.

But, on second thought, no such luck. The thing to keep your eye on, *pace* Locke and Descartes both, is that the relation between the content of

⁹ Well, maybe the acquisition of PROTON doesn't; it's plausible that PROTON is *not* typically acquired from its instances. So, as far as this part of the discussion is concerned, you are therefore free to take PROTON as a primitive concept if you want to. But I imagine

a sensory concept and the character of its cause is *not* arbitrary *when the cause is intentionally described*. The thing to keep your eye on is that we typically get the concept RED from (or, anyhow, on the occasion of) experiencing things *as red*.

There is, I think, more than a hint of a muddle about this in Fodor 1981a, where the following is a favourite line of argument: 'Look, *everybody*—Empiricists and Rationalists—agrees that there is at least one psychological mechanism which effects a *non-rational*, arbitrary relation between at least *some* primitive concepts and their distal causes. In particular, everybody agrees that the sensorium works that way.' "[E]ven the Empiricists hold that primitive concepts are merely triggered by [rather than learned from] experience . . . It is . . . just a fact about the way that we are put together than the sensory concepts we have are dependent in the ways they are upon the particular stimulations which occasion them" (ibid.: 275). On this account, Rationalism is simply the generalization of the Empiricist picture of the sensorium to cover whatever primitive concepts there turn out to be, sensory or otherwise: some kinds of arbitrary stimuli trigger (sensory) concepts like RED; other kinds of arbitrary stimuli trigger (non-sensory) concepts like DOORKNOB. What's the big sweat?

That I still like using the sensorium as a model of concept innateness at large will presently become clear. But, to repeat, *prima facie* it has a problem that needs to be taken seriously. The problem is that the triggering stimuli for RED *aren't* arbitrary when you take them under intentional (rather than psychophysical) description. If you take them under intentional description, the doorknob/DOORKNOB problem instantly emerges for sensory concepts too. It is encounters with doorknobs that typically occasion the acquisition of what Empiricists (and practically everybody else) have taken to be a *complex* concept like DOORKNOB; likewise it is typically encounters with red things (and not with green things, and not with square things, and not with elephants (unless they are red squares or red elephants)) that typically occasion the acquisition of what practically everybody takes to be a *primitive* concept like RED. Surely that's no accident in either case? And if it's not an accident, what else but an inductive model of concept acquisition could explain it?

This begins to seem a little worrying. It is perhaps tolerable that representational theories of mind should lead by plausible arguments to quite a radical nativism. But it is surely not tolerable that they should lead by plausible arguments to a contradiction. If the d/D effect shows that primitive concepts *must* be learned inductively, and SA shows that

concepts, then there aren't any concepts at all. And if there aren't any concepts all, RTM has gone West. Isn't it a bit late in the day (and late in the book) for me to take back RTM?

Help!

Ontology

This all started because we were in the market for some account of how DOORKNOB is acquired. The story couldn't be hypothesis testing because Conceptual Atomism was being assumed, so DOORKNOB was supposed to be primitive; and it's common ground that the mechanism for acquiring primitive concepts can't be any kind of induction. But, as it turned out, there is a further constraint that whatever theory of concepts we settle on should satisfy: it must explain why there is so generally a content relation between the experience that eventuates in concept attainment and the concept that the experience eventuates in attaining. At this point, the problem about DOORKNOB metastasized: assuming that primitive concepts are triggered, or that they're 'caught', won't account for their content relation to their causes; apparently only induction will. But primitive concepts can't be induced; to suppose that they are is circular. What started as a problem about DOORKNOB now looks like undermining all of RTM. This is not good. I was relying on RTM to support me in my old age.

But, on second thought, just *why* must one suppose that only a hypothesis-testing acquisition model can explain the doorknob/DOORKNOB relation? The argument for this is, I'm pleased to report, non-demonstrative. Let's go over it once more: the hypothesis-testing model takes the *content* relation between a concept and the experience it's acquired from to be a special case of the *evidential* relation between a generalization and its confirming instances (between, for example, the generalization that *F*s are *G*s and instances of things that are both *F* and *G*). You generally get DOG from (typical) dogs and not, as it might be, from ketchup. That's supposed to be because having DOG requires believing (as it might be) that typical dogs bark. (Note, once again, how cognitivism about *concept possession* and inductivism about *concept acquisition* take in one another's wash.) And, whereas encounters with typical dogs constitute evidence that dogs bark, encounters with ketchup do not (*ceteris paribus*). If the relation between concepts and experiences

That is what is called in the trade a 'what-else' argument. I have nothing against what-else arguments in philosophy; still less in cognitive science. Rational persuasion often invokes considerations that are convincing but not demonstrative, and what else but a what-else argument could a convincing but non-demonstrative argument be? On the other hand, it is in the nature of what-else arguments that '*Q* if not *P*' trumps 'What else, if not *P*?'; and, in the present case, I think there is a prima facie plausible ontological candidate for *Q*; that is, an explanation which makes the d/D effect the consequence of a metaphysical truth about how concepts are constituted, rather than an empirical truth about how concepts are acquired. In fact, I know of two such candidates, one of which might even work.

First Try at a Metaphysical Solution to the d/D Problem

If you assume a causal/historical (as opposed to a dispositional/counterfactual) construal of the locking relation, it might well turn out that there is a metaphysical connection between acquiring DOORKNOB and causally interacting with doorknobs. (Cf. the familiar story according to which it's because I have causally interacted with water and my Twin hasn't that I can think *water*-thoughts and he can't.) Actually, I don't much like causal/historical accounts of locking (see Fodor 1994: App. B), but we needn't argue about that here. For, even if causally interacting with doorknobs is metaphysically *necessary* for DOORKNOB-acquisition, it couldn't conceivably be metaphysically *sufficient*; just causally interacting with doorknobs doesn't guarantee you any concepts at all. That being so, explaining the doorknob/DOORKNOB effect requires postulating some (contingent, psychological) mechanism that reliably leads from having *F*-experiences to acquiring the concept of *being F*. It understates the case to say that no alternative to hypothesis testing suggests itself. So I don't think that a causal/historical account of the locking relation can explain why there is a d/D effect without invoking the very premiss which, according to SA, it can't have: viz. that primitive concepts are learned inductively.

Note the similarity of this objection to the one that rejected a Darwinian solution of the d/D problem: just as you can't satisfy the conditions for having the concept *F* just in virtue of having interacted with *F*s, so too you can't satisfy the conditions for having the concept *F* just in virtue of your grandmother's having interacted with *F*s. In both cases, *concept acquisition requires something to have gone on in your head in consequence of the interactions.* Given the ubiquity of the d/D phenom-

Second Try at a Metaphysical Solution to the d/D Problem

Maybe what it is to be a doorknob isn't *evidenced* by the kind of experience that leads to acquiring the concept DOORKNOB; maybe what it is to be a doorknob is constituted by the kind of experience that leads to acquiring the concept DOORKNOB. A Very Deep Thought, that; but one that requires some unpacking. I want to take a few steps back so as to get a running start.

Chapter 3 remarked that it's pretty clear that if we can't define "doorknob", that can't be because of some accidental limitation of the available metalinguistic apparatus; such a deficit could always be remedied by switching metalanguages. The claim, in short, was not that we can't define "doorknob" *in English*, but that we can't define it *at all*. The implied moral is interesting: if "doorknob" can't be defined, the reason that it can't is plausibly not *methodological* but *ontological*; it has something to do with what kind of property *being a doorknob* is. If you're inclined to doubt this, so be it; but I think that you should have your intuitions looked at.

Well, but *what* could it be about *being a doorknob* that makes 'doorknob' not definable? Could it be that doorknobs have a "hidden essence" (as water, for example, is supposed to do); one that has eluded our scrutiny so far? Perhaps some science, not yet in place, will do for doorknobs what molecular chemistry did for water and geometrical optics did for mirrors: make it clear to us what they *really* are? But *what* science, for heaven's sake? And what could there be for it to make clear? Mirrors are puzzling (it seems that they double things); and water is puzzling too (what could it be made of, there's so much of it around?). But doorknobs aren't *puzzling*; doorknobs are *boring*. Here, for once, "further research" appears not to be required.

It's sometimes said that doorknobs (and the like) have *functional* essences: what makes a thing a doorknob is what it is (or is intended to be) used for. So maybe the science of doorknobs is psychology? Or sociology? Or anthropology? Once again, believe it if you can. In fact, the intentional aetiology of doorknobs is utterly transparent: they're intended to be used as doorknobs. I don't at all doubt that's what makes them what they are, but that it is gets us nowhere. For, if DOORKNOB plausibly lacks a conceptual analysis, INTENDED TO BE USED AS A DOORKNOB does too, and for the same reasons. And surely, *surely*, that can't, in either case, be because there's something secret about *doorknobhood* that depth psychology is needed to reveal? No doubt, there is a lot that we don't know about intentions towards doorknobs *qua* intentions; but I can't believe

them doorknobs, and either it's something complex or it's something simple. If it's something complex, then 'doorknob' must have a definition, and its definition must be either "real" or "nominal" (or both). If 'doorknob' has a nominal definition, then it ought to be possible for a competent linguist or analytical philosopher to figure out what its nominal definition is. If 'doorknob' has a real definition, then it ought to be possible for a science of doorknobs to uncover it. But linguists and philosophers have had no luck defining 'doorknob' (or, as we've seen, anything much else). And there is nothing for a science of doorknobs to find out. The direction this is leading in is that if 'doorknob' is undefinable, that must be because *being a doorknob* is a *primitive property*. But, of course, that's crazy. If a thing has *doorknobhood*, it does so entirely in virtue of others of the properties it has. If doorknobs don't have hidden essences or real definitions, that can't possibly be because *being a doorknob* is one of those properties that things have simply because they have them; ultimates like *spin*, *charm*, *charge*, or the like, at which explanation ends.

So, here's the riddle. How could 'doorknob' be undefinable (contrast 'bachelor' =_{df} 'unmarried man') and lack a hidden essence (contrast water = H₂O) without being metaphysically primitive (contrast spin, charm, and charge)?

The answer (I think) is that 'doorknob' works like 'red'.

Now I suppose you want to know how 'red' works.

Well, 'red' hasn't got a nominal definition, and *redness* doesn't have a *real* essence (ask any psychophysicist), and, of course, *redness* isn't metaphysically ultimate. This is all OK because *redness* is an *appearance property*, and the point about appearance properties is that they don't raise the question that definitions, real and nominal, propose to answer: viz. 'What is it that the things we take to be Xs have in common, over and above our taking them to be Xs?' This is, to put it mildly, not a particularly original thing to say about red. All that's new is the proposal to extend this sort of analysis to doorknobs and the like; the proposal is that there are lots of appearance concepts that aren't sensory concepts.¹⁰ That this should be so is, perhaps, unsurprising on reflection. There is no obvious reason why

¹⁰ So, then, which appearance properties are sensory properties? Here's a line that one might consider: *S* is a sensory property only if it is possible to have an experience of which *S*-ness is the intentional object (e.g. an experience (as) of red) *even though one hasn't got the concept S*. Here the test of having the concept *S* would be something like being able to think thoughts whose truth conditions include ... *S* ... (e.g. thoughts like *that's red*). I think this must be the notion of 'sensory property' that underlies the Empiricist idea that RED and the like are learned 'by abstraction' from experience, a doctrine which presupposes that a

a property that is constituted by the *mental* states that things that have it evoke in us must *ipso facto* be constituted by the *sensory* states that things that have it evoke in us.

All right, all right; you can't believe that something's being a doorknob is "about us" in anything like the way that maybe something's being red is. Surely 'doorknob' expresses a property that a thing either has or doesn't, regardless of our views; as it were, a property of things in themselves? So be it, but *which* property? Consider the alternatives (here we go again): is it that 'doorknob' is definable? If so, what's the definition? (And, even if 'doorknob' is definable, some concepts have to be primitive, so the present sorts of issues will eventually have to be faced about *them*.) Is it that doorknobs qua doorknobs have a hidden essence? Hidden where, do you suppose? And who is in charge of finding it? Is it that being a doorknob is ontologically ultimate? You've got to be kidding.¹¹

If you take it seriously that DOORKNOB hasn't got a conceptual analysis, and that doorknobs don't have hidden essences, all that's left to make something a doorknob (anyhow, all that's left that I can think of) is *how it strikes us*. But if being a doorknob is a property that's constituted by how things strike us, then the intrinsic connection between the content of DOORKNOB and the content of our doorknob-experiences is metaphysically necessary, hence not a fact that a cognitivist theory of concept acquisition is required in order to explain.

To be sure, there remains something about the acquisition of DOORKNOB that does want explaining: viz. why it is the property that these guys (several doorknobs) share, and not the property that those guys (several cows) share, that we lock to from experience of good (e.g. stereotypic) examples of doorknobs. And, equally certainly, it's got to be something about our kinds of minds that this explanation adverts to. But, I'm supposing, such an explanation is *cognitivist* only if it turns on the *evidential* relation between *having the stereotypic doorknob properties* and *being a doorknob*. (So, for example, triggering explanations aren't

have an experience (as) of doorknobs, I suppose only a mind that has the concept DOORKNOB can do so.

'But *how* could one have an experience (as) of red if one hasn't got the concept RED?' It's easy: in the case of *redness*, but not of *doorknobhood*, one is equipped with sensory organs which produce such experiences when they are appropriately stimulated. *Redness* can be *sensed*, whereas the perceptual detection of *doorknobhood* is always *inferential*. Just as sensible psychologists have always supposed.

¹¹ The present discussion parallels what I regard as a very deep passage in Schiffer 1987 about *being a dog*. Schiffer takes for granted that 'dog' doesn't name a species, and (hence?) that dogs as such don't have a hidden essence. His conclusion is that there just isn't (except

cognitivist according to this criterion, *and wouldn't be even if (by accident) the concept DOORKNOB happened to be triggered by doorknobs*.) Well, by this criterion, my story isn't cognitivist either. My story says that what doorknobs have in common qua doorknobs is *being the kind of thing that our kind of minds (do or would) lock to from experience with instances of the doorknob stereotype*. (Cf. to be red *just is* to have that property that minds like ours (do or would) lock to in virtue of experiences of typical instances of redness.) Why isn't that OK?¹²

If you put that account of the metaphysics of doorknobhood together with the metaphysical account of concept possession that informational semantics proposes - having a concept is something like "resonating to" the property that the concept expresses - then you get: being a doorknob is having that property that minds like ours come to resonate to in consequence of relevant experience with stereotypic doorknobs. That, and not being learned inductively, is what explains the content relation between DOORKNOB and the kinds of experience that typically mediates its acquisition. It also explains how *doorknobhood* could seem to be undefinable and unanalysable without being metaphysically ultimate. And it also explains how DOORKNOB could be both psychologically primitive and not innate, the Standard Argument to the contrary notwithstanding.

Several points in a spirit of expatiation:

The basic idea is that what makes something a doorknob is just: being the kind of thing from experience with which our kind of mind readily acquires the concept DOORKNOB. And, conversely, what makes something the concept DOORKNOB is just: expressing the property that our kinds of minds lock to from experience with good examples of instantiated *doorknobhood*. But this way of putting the suggestion is too weak since experience with stereotypic doorknobs might cause one to lock to any of a whole lot of properties (or to none), depending on what else is going on at the time. (In some contexts it might cause one to lock to the property *belongs to Jones*.) Whereas, what I want to say is that *doorknobhood* is the property that one gets locked to when experience with typical doorknobs causes the locking and does so *in virtue of the properties they have qua typical doorknobs*. We have the kinds of minds that often

¹² Modal footnote (NB): Here as elsewhere through the present discussion, 'minds like ours' and 'the (stereo)typical properties of doorknobs' are to be read *rigidly*, viz. as denoting the properties that instances of *stereotypic doorknobs* and *typical minds* have in *this world*.

acquire the concept *X* from experiences whose intentional objects are properties belonging to the *X*-stereotype.¹³

Notice that this is not a truism, and that it's not circular: it's contingently true if it's true at all. What makes it contingent is that being a doorknob is neither necessary nor sufficient for something to have the stereotypic doorknob properties (not even in 'normal circumstances' in any sense of "normal circumstances" I can think of that doesn't beg the question). *Stereotype* is a statistical notion. The only theoretically interesting connection between *being a doorknob* and *satisfying the doorknob stereotype* is that, contingently, things that do either often do both.

In fact, since the relation between instantiating the doorknob stereotype and being a doorknob is patently contingent, you might want to buy into the present account of DOORKNOB even if you don't like the Lockean story about RED. The classical problem with the latter is that it takes for granted an unexplicated notion of 'looks red' ('red experience', 'red sense datum', or whatever) and is thus in some danger of circularity since "the expression 'looks red' is not semantically unstructured. Its sense is determined by that of its constituents. If one does not understand those constituents, one does not fully understand the compound" (Peacocke 1992: 408). Well, maybe this kind of objection shows that an account of *being red* mustn't presuppose the property of *looking red* (though Peacocke doubts that it shows that, and so do I). In any event, no parallel argument could show that an account of *being a doorknob* mustn't presuppose the property of *satisfying the doorknob stereotype*. The conditions for satisfying the latter are patently specifiable without reference to the former, viz. by enumerating the shapes, colours, functions, and the like that doorknobs typically have.

It's actually sort of remarkable that all of this is so. *Pace* Chapter 5, concepts really *ought* to be stereotypes. Not only because there's so much evidence that having a concept and having its stereotype are reliably closely correlated (and what better explanation of reliable close correlation could there be than identity?) but also because it is, as previously noted, generally *stereotypic* examples of *X-ness* that one learns *X* from. Whereas, what you'd *expect* people reliably to learn from stereotypic examples of *X* isn't

¹³ How much such experience? And under what conditions of acquisition? I assume that there are (lots of) empirical parameters that a formulation of the laws of concept acquisition would have to fill in. Doing so would be the proprietary goal of a serious

the concept *X* but the *X* stereotype.¹⁴ A stereotypic *X* is *always* a better instance of the *X* stereotype than it is of *X*; that is a truism.¹⁵

Interesting Digression

The classic example of this sort of worry is the puzzle in psycholinguistics about 'Motherese'. It appears that mothers go out of their way to talk to children in stereotypic sentences of their native language; in the case of English, relatively short sentences with NVN structure (and/or Agent Action Object structure; see Chapter 3). The child is thereby provided with a good sample of stereotypic English sentences, from which, however, he extracts not (anyhow, not only) the concept STEREOTYPIC ENGLISH SENTENCE, but the concept ENGLISH SENTENCE *TOUT COURT*. But why on Earth does he do that? Why doesn't he instead come to believe that the grammar of English is $S \rightarrow NVN$, or some fairly simple elaboration thereof, taking such apparent counter-examples as he may encounter as not well-formed? Remember, on the one hand, that Mother is following a strategy of screening him from utterances of unstereotypic sentences; and, on the other hand, that he'll hear lots of counter-examples to *whatever* grammar he tries out, since people say lots of ungrammatical things. I think the answer *must* be that it's a law about our kinds of minds that they are set up to make inductions from samples consisting largely of stereotypic English sentences to the concept ENGLISH SENTENCE (viz. the concept sentences satisfy in virtue of being well-formed relative to the grammar of English) and not from samples consisting largely of stereotypic English sentences to the concept STEREOTYPIC ENGLISH SENTENCE (viz. the concept sentences satisfy in virtue of being NVN).

In short, I do think there's good reason for cognitive scientists to be unhappy about the current status of theorizing about stereotypes. The kinds of worries about compositionality that Chapter 5 reviewed show that the relation a stereotype bears to the corresponding concept *can't* be constitutive. The standard alternative proposal is that it is simply heuristic; e.g. that stereotypes are databases for fast recognition procedures. But this seems not to account for the ubiquity and robustness of stereotype phenomena; and, anyhow, it begs the sort of question that we just discussed: why is it the concept *X* rather than the concept STEREOTYPIC *X* that one normally gets from experience with stereotypic *X*s? (*Mutatis mutandis*, if the way perception works is that you subsume things under

¹⁴ Reminder: 'the *X* stereotype' is rigid. See n. 12 above.

DOORKNOB by seeing that they are similar to stereotypic doorknobs, why is it that you generally see a doorknob as a doorknob, and not as *something that satisfies the doorknob stereotype*?) If our minds are, in effect, functions from stereotypes to concepts, that is a fact *about us*. Indeed, it is a *very deep* fact about us. My point in going on about this is to emphasize the *untriviality* of the consideration that we typically get a concept from instances that exemplify its stereotype.

That a concept has the stereotype that it does is never truistic; and that a stereotype belongs to the concept that it does is never truistic either. In particular, since the relation between a concept and its stereotype is always contingent, no circularity arises from defining 'the concept *X*' by reference to 'the stereotype of the concept *X*'. But, according to the present proposal, the relation between being a doorknob and instantiating the doorknob stereotype is, as it were, *almost* constitutive. Instantiating *doorknobhood* and instantiating the corresponding stereotype are logically, conceptually, and metaphysically independent in both directions.¹⁶ But the following is *metaphysically necessary*, according to the line I'm selling: being a doorknob is having the property to which minds like ours generalize from experiences (as of) the properties by which the doorknob stereotype is constituted. *That's what the mind-dependence of doorknobhood consists in.*

By way of a sort of summary, I want to rub in something that I said before: there is a sense, quite different from the one I've been discussing, in which it's pretty untendentious that *being a doorknob* is a mind-dependent property. Perhaps it's in the nature of doorknobs that they are artefacts. Perhaps, for example, nothing that *just grew* on a door could be a doorknob. Since it's in the nature of artefacts that have a certain kind of intentional history, it follows that there would be no doorknobs but that there are intentions with respect to doorknobs. A fortiori, there would be no doorknobs if there were no minds. Have this however you will; I raise the issue only to distinguish it from the one that I care about.

My line is that whether a thing is a doorknob is a matter of how it strikes us. By contrast, if being a doorknob is having the right sort of intentional history, then it's straightforwardly a matter of fact whether a

¹⁶ In principle, they are also *epistemically* independent in both directions. As things are now, we find out about the stereotype by doing tests on subjects who are independently identified as having the corresponding concept. But I assume that if we knew enough about the mind/brain, we could *predict* a concept from its stereotype and vice versa. In effect, given the infinite set of actual and possible doorknobs, we could predict the stereotype from which our sorts of minds would generalize to it: and given the doorknob stereotype we

thing *is* a doorknob. That's because what intentional history a thing has is metaphysically independent of what intentional history it strikes anyone as having.¹⁷ *Being married* is a matter of intentional history; one has to have said certain things, under certain conditions, with certain intentions, . . . etc. But whether Napoleon was married isn't up to us; nor, for that matter, is whether you are married up to you. Whether you are married is metaphysically independent of whether you wish or take yourself to be. It's too late to change your mind, and 'I forgot' does not defend against a charge of bigamy.

And, anyhow, I think the metaphysics of lots of concepts that do *not* subsume artefacts, and are patently not constituted by their intentional histories, works in much the same way that the metaphysics of DOORKNOB does. In fact, I rather think that is true of *all* concepts that aren't logico-mathematical and don't express natural kinds. More of this in Chapter 7.

I've been suggesting that whether a thing is a doorknob is maybe constituted by facts about whether we (do or would) take it to be a doorknob; just as whether something is red is maybe constituted by facts about whether it looks red to us. The metaphysical camel I'm trying to get you to swallow is, to repeat, an analogy between DOORKNOB and *appearance* concepts: with *doorknob* as with *red*, all there is to being it is how things tend to strike us. This account of the metaphysics of doorknobs would seem to explain why DOORKNOB exhibits the d/D effect without having to assume that DOORKNOB is learned inductively. So far, then, the present picture is compatible with the idea that DOORKNOB is primitive. So it's compatible with Semantic Atomism.

Suppose, if only for the sake of the discussion, that you're prepared to consider the ontology I've been trying to sell you. Then: *what's the bottom line about Innate Ideas?*

Innateness and Ontology

The natural, appalled, reaction to radical concept nativism is: 'But how *could* you have a concept like DOORKNOB innately?' To which the proper answer is: 'That depends a lot on what the concept DOORKNOB is and it depends a lot on what it is to have a concept.' According to the present proposal, to have a concept is to be locked to the corresponding property. But also, according to the present proposal, DOORKNOB is

an appearance concept; the property it expresses is constituted by the way that things that have it (do or would) strike us (if we have had or were to have appropriate experiences with stereotypic doorknobs). Well, if a property is *constituted* by the way that things that have it strike us (under certain circumstances), then being locked to the property requires only that things that have it do reliably strike us that way (under those circumstances).

The model, to repeat, is *being red*: all that's required for us to get locked to *redness* is that red things should reliably seem to us as they do, in fact, reliably seem to the visually unimpaired. Correspondingly, all that needs to be innate for RED to be acquired is whatever the mechanisms are that determine that red things strike us as they do; which is to say that all that needs to be innate is the sensorium. Ditto, *mutatis mutandis*, for DOORKNOB if *being a doorknob* is like *being red*: what has to be innately given to get us locked to doorknobhood is whatever mechanisms are required for doorknobs to come to strike us as such. Put slightly differently: if the locking story about concept possession and the mind-dependence story about the metaphysics of doorknobhood are both true, then the kind of nativism about DOORKNOB that an informational atomist has to put up with is perhaps not one of concepts but of mechanisms. That consequence may be some consolation to otherwise disconsolate Empiricists.

I suppose the philosophically interesting question about whether there are innate ideas is whether there are innate *ideas*. It is, after all, the thought that the 'initial state' from which concept acquisition proceeds must be specified in intentional terms (terms like 'content', 'belief', etc.) that connects the issues about concept innateness with the epistemological issues about a prioricity and the like. (By contrast, I suppose the *ethologically* interesting question is not whether what's innate is strictly speaking intentional, but whether it is domain specific and/or species specific. Perhaps you find the ethologically interesting question more interesting than the philosophically interesting question. And perhaps you're right to do so. Still, they are *different* questions.) Correspondingly, the 'innate sensorium' model suggests that the question how much is innate in concept acquisition can be quite generally dissociated from the question whether any *concepts* are innate. The sensorium is innate by assumption, and there would quite likely be no acquiring sensory concepts but that this is so. But, to repeat, the innateness of the sensorium isn't the innateness of anything that has intentional content. Since the sensorium isn't an idea, it is a fortiori not an *innate* idea. So, strictly speaking, the

To be sure, RED and DOORKNOB *could* both be innate for all I've said so far. But the main motivation for saying that they are is either that one finds inductivist theories of concept acquisition intrinsically attractive, or that noticing the d/D effect has convinced one that some such theory must be true whether or not it's attractive. Well, SA blocks the first motivation. And, as we've been seeing, it may be that the explanation of the d/D effect is metaphysical rather than psychological. In which case, unless I've missed something, there isn't *any* obvious reason why the initial state for DOORKNOB acquisition needs to be intentionally specified. A fortiori, there isn't any obvious reason why DOORKNOB needs to be innate. *NOT EVEN IF IT'S PRIMITIVE*. The moral of all this may be that though there has to be a story to tell about the structural requirements for acquiring DOORKNOB, intentional vocabulary isn't required to tell it. In which case, it isn't part of cognitive psychology.

Not even of "cognitive neuropsychology", if there is such a thing (which I doubt). Suppose we were able to specify, in neurological vocabulary, the initial state from which DOORKNOB acquisition proceeds. The question would then arise whether the neurological state so specified is intentional — whether it has conditions of semantic evaluation (and, if so, what they are). So far, we haven't found a reason for supposing that it does. To be sure, it is an innate, possibly quite complicated, state from which DOORKNOB may be acquired, given experience of e.g. doorknobs. But this is all neutral as to whether the initial state is an intentional state; it's all true *whether or not* the initial state is an intentional state. So it's all true whether or not the initial state for DOORKNOB acquisition is in the domain of *cognitive* neuropsychology (as opposed, as it were, to neuropsychology *tout court*).

None of this could be *much* comfort to a disconsolate Empiricist, since none of it is supposed to deny, even for a moment, that a lot of stuff that's domain specific or species specific or both has to be innate in order that we should come to have the concept DOORKNOB (or for that matter, the concept RED). But the issue isn't whether acquiring DOORKNOB requires a lot of innate stuff; anybody with any sense can see that it does. The issue is whether it requires a lot of innate *intentional* stuff, a lot of innate stuff that has content. All the arguments I know that say that innate intentional stuff has to mediate concept acquisition depend on assuming either that concept acquisition is inductive or that the explanation of the d/D effect is psychological or both. Well, where a primitive concept expresses a mind-dependent property, it is very unclear that either of these kinds of argument will work.

APPENDIX 6A

Similarity

'Hey, aren't you just saying that all that has to be innate in a DOORKNOB-acquisition device is the capacity to learn to respond selectively to things that are relevantly similar to doorknobs? And didn't Quine say that years ago?'

No, I'm not and no, he didn't. Not quite.

There are two ways to understand the claim that the process of acquiring DOORKNOB recruits an innate 'similarity metric'. One is platitudinous, the other is committed to innate ideas—in effect, to the innateness of the concept SIMILAR TO A DOORKNOB. The geography around here is pretty familiar, so we can settle for a quick tour.

On the first way of running it, the similarity story is just the remark that, given appropriate experience of doorknobs, creatures like us converge on a capacity to respond selectively to things that are like doorknobs *in respect of their doorknobhood*. This is perfectly self-evidently true; nobody reasonable could wish to deny it. It doesn't, however, *explain* the fact that we learn DOORKNOB from doorknobs; it just repeats the fact that we do. So construed, the similarity story is completely neutral on the issues this chapter is concerned with, viz. whether the structures in virtue of which we are able to converge on selective sensitivity to *doorknobhood* need to be innate, and whether they need to be intentional.

On the other, unplatitudinous, way of running the similarity theory, it is itself a version of concept nativism: it's the thesis that what's innate is the concept SIMILAR TO A DOORKNOB. There seems, to put it mildly, to be no reason to prefer that view to one that has DOORKNOB itself be innate. (Indeed, the first would seem to imply the second; since the concept SIMILAR TO A DOORKNOB is, on the face of it, a construct out of the concept DOORKNOB, it's hard to imagine how anyone could think the one concept unless he could also think the other.) None of this bothers Quine much, of course, because he pretty explicitly assumes the Empiricist principle that the innate dimensions of similarity, along which experience generalizes, are sensory. But Empiricism isn't true, and it is time to put away childish things.

Quine's story is that learning DOORKNOB is learning to respond selectively to things that are *similar to doorknobs*. What the story amounts to depends, in short, on how *being similar to doorknobs* is construed. Well, there's a dilemma: if *being similar to doorknobs* is elucidated by appeal to *doorknobhood*, then the story is patently empty; 'How is the concept that

doorknobs is spelled out by reference to properties other than *doorknobhood*, Quine has to say which properties these are, where the concepts of these properties come from, and how radical nativism with respect to them is to be avoided.

Like Quine, I've opted for the second horn of the dilemma. But, unlike Quine, I'm no Empiricist. Accordingly, I can appeal to the doorknob stereotype to say what 'similarity to doorknobs' comes to, and—since 'the doorknob stereotype' is independently defined—I can do so without invoking the concept DOORKNOB and thereby courting platitude.

So I'm not saying what Quine said; though it may well be what he should have said, and would have said but for his Empiricism. I often have the feeling that I'm just saying what Quine would have said but for his Empiricism.¹⁵

¹⁵ I am also, unlike Quine, not committed to construing locking in terms of a capacity for discriminated responding (or, indeed, of anything epistemological). Locking reduces to nomic connectedness. (I hope.) See Fodor 1990; Fodor forthcoming *b*.