

Contents

1. Introduction	1
Part I: Naturalism	
2. Pillars of the Tradition	21
3. Naturalism Characterized	50
Part II: Ontology	
4. The Discovery Problem	77
5. Proper Function	108
6. Pragmatic Arguments	128
7. What Price Antirealism?	158
Part III: Alternatives	
8. Intuitionism	173
9. Supernaturalism	212
References	227
Index	241

1

Introduction

PHILOSOPHICAL naturalism has dominated the Western academy for well over a century. It is not just fashionable nowadays; it enjoys the lofty status of academic orthodoxy. However, there is an important sense in which naturalism's status as orthodoxy is without rational foundation. Furthermore, the costs of embracing it are surprisingly high. The goal of this book is to defend these two claims.

In the present chapter I will introduce several of the concepts and assumptions that will occupy center stage in the book's main argument. In Section 1, I will introduce the notion of a *research program*. I will also defend the conclusion that it is impossible to adopt a research program on the basis of evidence. This will constitute my argument for the conditional claim that *if* naturalism is a research program, its status as orthodoxy is without rational foundation. In later chapters I will argue that naturalism is indeed a research program. In Section 2, I will introduce the central thesis of the book and explain in some detail the concepts involved in that thesis. Finally, in Section 3, I will provide a brief outline of what is to come.

1. Research Programs

Inquiry is a process in which we try to revise our beliefs in some way—by acquiring new ones, discarding old ones, or both. This is so whether we are trying to answer scientific questions like 'What causes thunder?', philosophical questions like 'Is capital punishment immoral?', or more mundane questions like 'Where are my keys?', 'Why does my stomach hurt?', or 'What's in that dark room?' But not just any attempt at belief revision counts as a process of inquiry. A severe blow to the head might bring about revisions in your belief

structure; but hitting yourself on the head with a baseball bat is not, all by itself, an investigative process. True inquiry is a process in which we try to revise our beliefs on the basis of what we take to be evidence.

But this means that, in order to inquire into anything, we must *already* be disposed to take some things as evidence. In order even to begin an inquiry, we must already have various dispositions to trust at least some of our cognitive faculties as sources of evidence and to take certain kinds of experiences and arguments to be evidence. Such dispositions (let's call them *methodological dispositions*) may be reflectively and deliberately acquired. But for most people most of the time, they are probably acquired unreflectively and even unconsciously. We tend automatically to trust our senses, our reasoning abilities, and our memories. The fact that mathematical and logical propositions seem obviously to be true we readily and unreflectively take as strong and sufficient evidence in their favor. But we cast a cool eye of skepticism upon our untested beliefs on other topics. In matters of health, auto mechanics, gardening, and the like, we demand empirical evidence or testimony from an expert. Hunches about the stock market we might trust if we are expert investors (perhaps assuming that good reasoning lurks tacitly behind them). But, in our sober moments, hunches about lottery tickets and games of Russian roulette are steadfastly ignored. We find deductive arguments compelling, inductive arguments less so, and many other sorts of arguments laughable. In short, we are disposed to trust certain ways of acquiring information with respect to various topics and to distrust others; and, though sometimes our being so disposed is a result of conscious and reflective activity on our part, quite often it is not.

We may also note that methodological dispositions differ from person to person. Some take hunches of various sorts or burnings in the breast very seriously as sources of evidence; others regard all such episodes as epistemically worthless. Some are inclined to trust just about anyone's testimony; others are recalcitrant skeptics. Furthermore, people differ with respect to what they take to be *basic* sources of evidence—sources that are to be trusted even in the absence of positive evidence in favor of their reliability. For some (maybe most), sense perception, reason, memory, rational intuition, and religious experience *all* count as basic sources. For others, at least some of these sources are derivative—they are to be trusted only after their reliability has been verified by evidence from the basic sources.

For purposes here, a *research program* is a set of methodological dispositions. *Individual* research programs are maximal sets of methodological dispositions, where a set of dispositions is maximal just in case it is possible to have all of the dispositions in the set but it is not possible to have all of them *and* to have other methodological dispositions as well. *Shared* research programs, on the other hand, are relevantly distinctive subsets of individual research programs. An example will help to clarify this. Consider the disposition to treat sense perception as a source of evidence under specific conditions $C_1 \dots C_n$. All by itself, this disposition does not count as an individual research program since one could have it while at the same time having a disposition to trust (say) religious experience under various circumstances. But that disposition in conjunction with a disposition to treat as evidence nothing other than sense perception as exercised in circumstances $C_1 \dots C_n$ would constitute an individual research program. Furthermore, the more general disposition to trust nothing other than sense perception, though compatible with a variety of different and more specific methodological dispositions, might nonetheless be characterized as a shared research program because (in some context of discussion) it might be seen as the relevantly distinctive core of a variety of different but closely related individual research programs.

I should mention in passing that the way in which I use the term ‘research program’ differs from, but is nevertheless related to, the way in which Imre Lakatos has famously used the term. For Lakatos, a (scientific) research program has three elements: a distinctive ‘hard core’ of theses that are, for all practical purposes, treated as immune to revision; a ‘protective belt’ of auxiliary theses that may be revised as needed to accommodate observations that would otherwise seem to threaten the hard core; and a set of problem-solving strategies.¹ The most obvious difference here is that, unlike research programs in my sense of the term, Lakatosian (scientific) research programs include theses, as well as methodological rules specifying which theses are potential candidates for rejection. Interestingly enough, however, Lakatos does say that the whole of science itself can be regarded as a kind of research program; and what constitutes it as such is not any particular set of hypotheses or rules aimed at protecting particular hypotheses but *simply* those problem-solving strategies that are

¹ See esp. Lakatos (1970: 4, 47 ff.).

distinctive of science in general (1970: 47). Thus, Lakatos seems to have, in addition to his more narrow notion of a scientific research program, a somewhat broader notion according to which some research programs may be nothing more than sets of problem-solving strategies. This broader notion is not far off from the notion I aim to characterize.

Despite the academic connotations normally associated with the term, I take it that every inquirer has a research program, and this regardless of whether the inquiry is conducted in a university research lab, an auto shop, or the person's own backyard. This is because I think that every inquirer must have some methodological dispositions, and that anyone who has methodological dispositions must have a maximal set thereof. It seems just obvious that, for any kind of experience or argument, anyone who is disposed to treat *some* experiences or arguments as evidence will either be disposed to treat *that* kind of experience or argument as evidence or be disposed not to treat it as evidence. But if that is right, then the person's total set of methodological dispositions will preclude all others. (For example, if such a person lacked a disposition to trust tabloids, she would not be able to acquire that disposition without losing a disposition not to trust them. Hence, she could not *retain* her original set of methodological dispositions while adding the new one.) But, of course, to say this is just to say that every such person has a maximal set of methodological dispositions. Thus, every inquirer has a research program.

Like methodological dispositions generally, our research program is something we *bring* to the table of inquiry. Method is therefore prior to theory; and so, at least in the first instance, our research program is not something that we intentionally adopt as a result of inquiry. Thus, in the first instance, a research program is not adopted on the basis of evidence. But must it always be this way? Can't we discard a research program and adopt a new one as a result of inquiry? Can't we theorize about our methods and acquire evidence that some methodological dispositions ought to be cultivated and others put aside? And then can't we respond to that evidence, and so embrace a research program for *reasons* rather than by happenstance, habit, or arbitrary preference? Initially, it might seem that we can. For example, we can acquire evidence that our vision is unreliable, that our memory is untrustworthy, or that we are unusually gullible; and so we can make a conscious effort to put aside our dispositions to trust

our faculties in these domains. If we succeed, we will have acquired a new research program. And it appears that we will have done so on the basis of evidence.

But the appearances are misleading. Research programs can be discarded on the basis of evidence, but they cannot be adopted on that basis. To see why, let us consider a concrete example. Suppose you start with a research program that includes a disposition to distrust sources of information that you take to be unreliable *and* a disposition to treat all testimony as evidence. (Very small children sometimes seem to have this sort of research program.) But suppose that, over time, you acquire what you take to be good evidence that the testimony of slick Uncle Bill is terribly unreliable. Then one of three things might happen. (a) You might acquire a disposition to distrust Uncle Bill, thereby losing your disposition toward universal gullibility. (b) You might discount the evidence in favor of Uncle Bill's unreliability, thereby losing some of the dispositions that initially led you to take that evidence seriously. Or (c) you might lose your disposition to distrust sources that you believe are unreliable. Given that at least *one* of these three things must happen, we have here a situation where evidence forces you to discard your research program. But evidence does *not* determine which of these three responses is the correct one. And so your adoption (say) of a research program that retains the disposition to distrust sources that you believe to be unreliable is not something that is done on the basis of evidence.

In making this point, I assume that one cannot have contradictory dispositions (i.e., one cannot at the same time be disposed both to do A and not to do A). But one might object that this very assumption should lead one to believe that the scenario just described is impossible. After all, if one *really* has a general disposition to distrust apparently unreliable sources, doesn't that imply that one's disposition to trust testimony is really only a disposition to trust apparently reliable testimony? Or, alternatively, if one really has a disposition to trust all testimony, doesn't that imply that one's general disposition to distrust apparently unreliable sources does not apply to testimony? Perhaps; and if so, then it seems that the right conclusion to draw here is that (a) research programs cannot defeat themselves in the way described above, and so (b) just as research programs cannot be adopted on the basis of evidence, so too they cannot be discarded on that basis either. This is a consequence I could be content with. It would force some revisions in what I say later on. (Most importantly:

I will say later on that a certain shared research program—*intuitionism*—is self-defeating under certain circumstances. But if I were to accept (a), I would recast the same argument as an argument for the conclusion that it is impossible for someone to be an intuitionist in those circumstances and so, for those who find themselves in those circumstances, intuitionism is not an alternative to naturalism.) But, as a matter of fact, I do not accept either the objection or its consequence. The reason is that, as I see it, the disposition toward universal gullibility and the disposition to distrust apparently unreliable sources do not come into conflict until the evidence of Uncle Bill's unreliability is uncovered. Thus, though the research program is ultimately self-defeating, it does not in fact defeat itself until the right sort of investigation has been undertaken. And I think that the same holds true for any other self-defeating research program.

In general, the reason why research programs cannot be adopted on the basis of evidence is that evidence can only be recognized as such from within a research program. As we have just illustrated, a research program might generate evidence that prescribes its own rejection. Furthermore, a research program might also generate evidence that some other research program is to be accepted. Still, it cannot be on the basis of this evidence that one accepts the favored program. For once the old program is rejected, the evidence arising out of the old program in favor of the new cannot be recognized as evidence. Or, at any rate, it cannot be recognized as evidence until a new program is in place which sanctions the sources that generated it. But even if this new program happens to be the program favored by the old, still it will not have been on the basis of the evidence generated by the old program that the new one was accepted. For the new program had to be accepted before that evidence could again be recognized as such. Thus, from an evidential point of view, the best we can say on behalf of a research program is that it is *self-supporting* (where a research program is self-supporting if and only if it does not prove self-defeating). We cannot say that it is supported by evidence that is somehow generated and recognizable as such independently of the program.

So when it comes to rejecting one program in favor of another, the decision to adopt the favored program must be made on pragmatic grounds, broadly speaking, rather than evidential grounds. In such cases, one chooses the program whose consequences are most attract-

ive, or whose canons are most convenient to adopt, or whose adoption will most irritate one's enemies, or whatever.² Furthermore, even if it happens to be true that (say) one rationally ought to adopt the program whose consequences are more attractive rather than the program whose adoption will most irritate one's enemies, there are no discernible grounds for asserting this truth. For, again, one could discern the grounds for asserting it *as grounds* for asserting it only within the context of some research program.

Understanding all of this is crucial to understanding the proper place of philosophical naturalism on the methodological landscape. As I said earlier, philosophical naturalism now enjoys the status of orthodoxy (much like a certain religious program enjoyed the status of methodological orthodoxy during the medieval years). In light of the foregoing, however, it should be clear why, if naturalism is a research program, its status as orthodoxy is without rational foundation. As we have already seen, there is no method-neutral basis on which to assess the decision to adopt a particular research program. Hence, we can have no rational grounds for declaring any such decision to be either categorically rational or categorically irrational. Thus, if naturalism is indeed a research program, as I will later argue that it is, there is no basis at all for saying that it is the sort of program that everybody, or every intelligent or right-thinking person, ought to adopt. This, of course, is no special deficiency of naturalism. If it is a deficiency at all, it is one shared by every other research program. But the point is important nonetheless, for it goes some distance toward disarming the current presumption in favor of naturalism. It also suggests that, on the assumption that naturalism is not self-defeating, the most appropriate tactic for trying to persuade others to reject it is to provide pragmatic considerations against it—highlighting its unattractive consequences, showing that other research programs do not suffer from the same problems, and the like. Since I am not prepared to defend the conclusion that naturalism is self-defeating, I will adopt the latter strategy in my own attack on naturalism.³

² This is not to say that a research program can be adopted at will. Perhaps we have no direct voluntary control over what research program we adopt. Still, this is no reason to doubt that our adopting a particular research program is ultimately guided by pragmatic considerations.

³ But see Ch. 8 for a limited defense of Alvin Plantinga's widely discussed 'evolutionary argument against naturalism'. Strictly speaking, Plantinga's argument is not directed at naturalism as I understand it; however, as I will note in Ch. 8, there is good reason to think that the argument nevertheless has some bearing on naturalism.

2. The Central Thesis

The central thesis of this book is that naturalists cannot, by their own lights, be justified in accepting two metaphysical views that many philosophers—naturalists in particular—very much want to accept. Those views are *realism about material objects* (RMO) and *materialism*. I will also argue that, on the assumption that standard naturalistic arguments against mind–body dualism are successful, naturalists must give up a third thesis: *realism about other minds* (ROM). Realism about other minds is, strictly speaking, a different thesis for every person. For me, it is the thesis that minds other than my own exist; for you it is the thesis that minds other than your own exist; and so on. Materialism and realism about material objects deserve more extended comment.

I will define materialism as the view that nothing exists except for spacetime, material objects and events in spacetime, and the properties exemplified by spacetime and the objects and events therein. The categories of material object and event are meant to be collectively broad enough to include things like electrons and fields; and they are meant to exclude such things as God (as traditionally conceived), Cartesian souls, angels, and the like. I admit that this definition is imprecise; but I am not aware of any definition in the literature that manages to be more precise while at the same time expressing a view that all materialists are plausibly thought to hold in common. Some, of course, will think that my definition is overly broad. For example, David Chalmers (1996) would surely admit that all mental properties are properties of material objects in spacetime; but he explicitly denies being a materialist. But that is no problem for present purposes. Materialism as I characterize it here is entailed by most of the theses standardly identified with materialism (and also physicalism). So in showing that naturalists must give up materialism as I have defined it, I will thereby also have shown that naturalists must give up those other doctrines as well.

Realism about material objects is, very roughly, the thesis that material objects exist independently of minds or mental activity. Notoriously, however, this very rough expression is not at all satisfactory. For one thing, many philosophers who count themselves realists about material objects are not at all opposed to the view that no material thing could possibly exist apart from God's mental activity; and those realists who do oppose this view do not usually take it as

part of their realism to oppose it. The role, if any, that God's mental activity plays in the world is not what is at issue in the realism/antirealism debate. Thus, realism about material objects should be understood not just as the view that mind-independent material objects exist, but as the view that material objects exist independently of human or other creaturely minds.

A more significant problem with the rough-and-ready formulation lies in the fact that the bare thesis that mind-independent material objects exist is ambiguous between the weak thesis that *some* material thing exists independently of the mental and the stronger thesis that *familiar* material things, or tokens of most current common-sense and scientific types, exist independently of the mental. For convenience, let's call these versions of realism *weak realism* about material objects and *strong realism* about material objects, respectively. (We will also do well to note that one can be a realist about the *external world* without being a realist about material objects. That is, one might hold that *something* exists independently of the mental without holding that some *material object* exists independently of the mental.⁴) I suspect that most people who profess to be realists about material objects take the word 'realism' to refer only to weak realism, even if they happen in fact to embrace strong realism about material objects. The reason I say this is that philosophers like Peter Unger (1979) and Peter van Inwagen (1990), who reject strong realism about material objects because of familiar puzzles about composition, are never numbered among the 'antirealist' crowd. Accordingly, for present purposes, I will take RMO to be equivalent to weak realism about material objects.

Still, we do not yet have a fully satisfactory expression of RMO. There is the further problem of figuring out what it means for some material thing to exist independently of the mental. What sort of dependence is being denied? Most realists will want to say that the dependence relation must be understood in such a way that minds themselves exist independently of the mental.⁵ Thus, we cannot say that the relevant dependence relation is bare logical dependence upon minds or mental activity; for minds are logically dependent upon the existence of minds and mental activity. Furthermore, realists also typically want to say that the dependence relation must be under-

⁴ Cf. the distinctions drawn between various kinds of realism in Devitt (1991).

⁵ A notable exception, however, is Michael Devitt. See Devitt (1991: 16).

stood in such a way that (for example) artifacts do not turn out to be mind-dependent simply by virtue of the fact that their existence is caused in part by the activity of a mind. If my house is mind-dependent, it is not because the builder's mental activity is among the causes of its existence. Some sort of mental activity is among the causes of the existence of most human beings; but, again, we do not think that that suffices to make *us* mind-dependent. Thus, the relevant dependence relation is not causal dependence either. But what other kinds of dependence are there?

One kind is the sort of dependence that thoughts bear to minds. Thoughts are essentially such as to be in a mind. Not only can they not exist in worlds without minds or mental activity, but they cannot exist in worlds with minds or mental activity unless they are among the contents of one of the minds. Clearly realists will want to say that neither minds nor material objects depend in this way upon the existence of minds. To say otherwise is to embrace idealism, the most extreme version of antirealism about material objects. But most participants in the realism/antirealism debate do not seem to think that idealism is the only version of antirealism about material objects. As most have it, antirealism about material objects comes in a variety that is compatible with realism about the external world. But if that is right, then there must be yet another kind of dependence relation that objects might bear to minds.

Often one hears antirealism about material objects characterized as the view that material objects are 'conceptual constructs', or as the view that we somehow 'carve up reality' according to our 'conceptual schemes', or as the view that by virtue of being concept users, we are in some sense 'worldmakers'. Making sense out of this sort of talk without construing antirealism about material objects as equivalent to idealism is extraordinarily difficult. After all, there is no obvious sense at all in which our minds or conceptual schemes function like carving knives upon a cosmic turkey. If the carving metaphor is not vacuous, at the very least it is obscure. And describing material objects as conceptual constructs, or talking about ourselves as if we make worlds by using our concepts, seems to do no better. Mythical creatures are, perhaps, conceptual constructs in some meaningful sense. After all, they exist only in our minds, and our ideas of them do indeed seem to be constructed out of other ideas or concepts. And mythical worlds (like Tolkien's Middle Earth) may in some sense be made by us. But understanding antirealism along these lines—i.e.

taking ourselves to make our world and its inhabitants in the same way that Tolkien made hobbits, orcs, and the rest of Middle Earth—takes us back to understanding antirealism as equivalent to idealism.

Nevertheless, I think that *some* sense can be made out of these metaphors. The way to do it is to understand the non-idealist variety of antirealism (let's call it 'constructivism') as nothing more than the thesis that none of the properties that appear to be sortal properties of nonabstract, nonmental objects are intrinsic to anything. Apparent sortal properties are properties that appear to correspond to sorts, or kinds, of objects. Examples include properties like *being an electron*, *being a horse*, *being a lump of cells*, *being a statue*, and so on. Intrinsic properties are properties that can be had by something regardless of whether it is accompanied or unaccompanied by any other contingent being.⁶ Intrinsic properties are opposed to extrinsic properties—properties that are exemplified by a thing only by virtue of relations that the thing bears to other contingently existing objects.

If sortal properties are in fact extrinsic properties, it is a bit difficult to say exactly what relations they might involve. Though other options are available, everyone who in fact believes that sortal properties are extrinsic believes that all apparent sortal properties, and thus all apparent essential properties, involve relations between minds and something else. What is difficult is specifying the something else.

It is tempting to say that, on this view, sortal properties involve relations between minds and the objects the properties sort. So, for example, an object counts as a horse (on this way of understanding the view) by virtue of the fact that we conceive of it as a horse. The sortal property *being a horse* thus involves relations between our minds and the horse itself. The trouble with this view, however, is that saying this makes it sound like the object's existence is somehow prior to its having any sortal properties. And this is something that just about everybody (whether realist or antirealist) wants to deny.

Another alternative is to say that sortal properties involve relations between objects and the *stuff* that composes them. On this view, the

⁶ This is a rough approximation of the definition defended in Langton and Lewis (1998). One untoward consequence of the definition is that the property *having been created by God* turns out to be intrinsic if God is a necessary being. I suspect the definition could be modified so as to avoid this consequence without substantially affecting the discussion that follows. But since the consequence does not matter for present purposes, I will not attempt to make such modifications.

word 'stuff' is supposed to refer to whatever it is that exists independently of our minds and is such that our conceiving of it in various ways is what makes it the case that the world contains the various material objects that we think it contains. A pure stuff ontology says that no objects exist independently of human mental activity, but only stuff. Unfortunately, however, it is not at all clear that a pure stuff ontology is coherent.⁷ Furthermore, even if it were coherent, it seems that the only plausible explanation for why we conceive of some stuff as a horse and other stuff as a table is that the equine stuff is intrinsically propertied differently than the tabular stuff. But it is not clear how pure stuff can have any properties. Properties seem to be the sorts of things that have to be had by an *object*.

For this reason, my preference is to say that, if sortal properties are extrinsic, they involve relations between human minds and a mind-independent world which consists of one or more mind-independent objects of an unidentifiable sort. The objects must be of an *unidentifiable* sort because the constructivist thesis is that all of the sortal properties we are familiar with are extrinsic. If we can identify an object's sort, its sort is not intrinsic, and so the object is not mind-independent. Furthermore, constructivism also implies that we cannot know how many mind-independent objects there are. For, obviously enough, if we don't know *what* mind-independent objects exist, we can't count them. On this view, then, apart from the mind-independent world, whose ultimate composition is inscrutable, everything that exists does so only by virtue of our conceiving of the world in the ways that we do. (Whose concepts count? That is a difficult question with no clear answer, so we'll leave it aside.) Note too that there is still room for talk of stuff on this view. We can use the word 'stuff' to refer to whatever mind-independent reality makes up the material objects of everyday discourse. We *should* use this word because, if constructivism is true, we have absolutely no idea whether it is one mind-independent object or many that makes up the world of material objects. Thus, we have no guarantee that individual material objects will be made up of any mind-independent objects at all. (They all might be made of nonobjectual portions of a single mind-independent thing.) But using the word 'stuff' this way, we do not encounter the problems associated with a pure stuff ontology; for, on the present interpretation of constructivism, we have mind-

⁷ On this, see Blackson (1992) and Carter and Bahde (1998).

independent objects in our ontology to serve as property bearers. This interpretation of constructivism therefore seems to be coherent. However, I should note that nothing in what follows depends on taking this interpretation rather than the one that presupposes a pure stuff ontology. My aim in offering it is just to show that even if, as I think, a pure stuff ontology is incoherent, still the view that sortal properties involve relations between minds and something else can be interpreted coherently.

A point worth highlighting is that, insofar as constructivism is distinct from idealism, it involves the view that some relationship to human minds is a necessary condition for an object's exemplifying any of the sortal properties that are apparently exemplified by material objects. It is perhaps also important to note that the view is not quite as crazy as it initially appears. Even realists will acknowledge that *some* apparent sortal properties are extrinsic. For example, many philosophers think that something can be a statue only if it is a work of art, and that something counts as a work of art only if it is the product of intelligent design. But if this is right, then the property *being a statue* is not intrinsic to anything. Constructivism just goes a (big) step further and says that *all* apparent sortal properties are this way. But this obviously does not imply that we 'make the world' by using our concepts in any radical sense. Indeed, on my interpretation, the view presupposes that the world exists and has intrinsic properties wholly independently of human beings and their mental activity. It's just that, according to constructivists, what makes it the case that (say) stuff in the world arranged horse-wise constitutes a horse rather than some other kind of object with a wholly different set of persistence conditions is the fact that the world which includes that stuff stands in a certain relation to us and our mental activity. Had we conceived of the world differently, different sortal properties would have been exemplified throughout, and thus different objects would have existed. Hence the carving metaphor; hence the metaphor of conceptual constructing; and hence the world-making metaphor.

I have made a point of distinguishing constructivism from idealism. But the two views are related. Though constructivism does not entail idealism, idealism does entail constructivism. The reason is that if apparent sortal properties are intrinsic to anything, they are intrinsic to the objects that they appear to sort; but, on idealism, those objects do not exist. So idealism trivially entails the thesis that apparent sortal properties are not intrinsic to anything, but not the

other way around. Furthermore, as we have already seen, constructivism is compatible with what I have called realism about the external world. What it rules out is simply the claim that *apparent* sortal properties—i.e. the properties that appear to us to be sortal properties—are intrinsic. It therefore also rules out the claim that the property *being a material object* is intrinsic, since that is an apparent sortal property. (Does it rule out the claim that the property *being an object* is intrinsic? I say no. ‘Object’ is not a sortal term; rather, objects are just the things that sortal terms sort. ‘Material object’, on the other hand, is a sortal term—at least if the term ‘material’ is not devoid of content, as I will assume it is not.)

I suspect that some will be inclined to think that if constructivism is true, then strictly speaking there are no material objects at all. On this view, then, constructivism does not really deny mind-independence. Rather, it denies the existence of material objects altogether.⁸ One reason for taking constructivism this way is that it might seem to be part of the very concept of a material object that material objects have their sortal properties intrinsically. Thus, the thesis that apparent sortal properties are not intrinsic to anything might seem to be logically equivalent to the thesis that there are no material objects. I am very sympathetic with this view about constructivism; but for present purposes I will presuppose that it is false. Thus, I will continue to assume that the dispute between constructivists and their opponents is over the question whether material objects (e.g. trees, human bodies, rocks, etc.) are mind-dependent rather than over the question whether material objects exist. Nothing substantial depends on this presupposition; but making it will help the discussion that follows to move along more smoothly.

Many readers will by now realize that the brand of antirealism I have just been attempting to characterize is the sort commonly attributed to philosophers like Immanuel Kant, John Dewey, Michael Dummett, Richard Rorty, Nelson Goodman, and Hilary Putnam. No doubt such readers also know that the view that I have just described (as well as, perhaps, any other reasonably well-defined view that might appear to be worthy of the label antirealism) cannot *uncontroversially* be attributed to any of these philosophers. So let me emphasize that, though I do claim to have made *some* sense out of the typical antirealist metaphors, I make no pretense to having made complete

⁸ Cf. the understanding of global anti-realism recommended in Cortens (1999).

sense out of them, or to having characterized antirealism in a way that most antirealists would accept, or even to having established that any nonidealist variety of antirealism really is coherent. However, I do believe that most contemporary philosophers who are typically seen as falling into the antirealist camp take themselves to be defending something *other* than realism or idealism; and I also believe that if there is any coherent characterization of their views that does not present them either as realists or idealists, that characterization must lie somewhat in the direction of the view described above.

Summing up, I see two ways of rejecting RMO (which, again, is equivalent to what I called *weak* realism about material objects). One way is to hold that everything that exists either is a mind or is among the contents of a mind. The other way is to endorse constructivism, the view that none of the properties that appear to be sortal properties of nonmental, nonabstract objects are intrinsic. Constructivism is compatible with realism about the external world; idealism is not. My main thesis, again, is that naturalists are committed to *rejecting* RMO and materialism (and, given certain assumptions, ROM as well). I will also argue that naturalists are committed to accepting constructivism, but not necessarily to accepting idealism.

3. What is to Come

Toward establishing the central thesis of the book, I will begin in Part I by defending my characterization of naturalism as a research program. One striking point of contrast between the naturalistic orthodoxy of the present century and the religious orthodoxy of earlier centuries is that the canons and consequences of the current orthodoxy—the doctrines that one is committed to by virtue of being orthodox—are not at all well defined. Much is done in the name of naturalism, but few have bothered to spell out in clear and precise terms what exactly it means to *be* a naturalist. For this reason, it is less than clear what (if any) philosophical costs or benefits attach to embracing naturalism. In Chapters 2 and 3, I argue that the characterization of naturalism that is most faithful to the tradition, and the one that best explains both the similarities and the differences one finds among contemporary naturalists, is one which takes naturalism to be not a view but a research program—strictly, a *shared* research program; but in what follows I will omit the qualifier. Chapter 2

provides a brief discussion of the prehistory of naturalism, together with a more extended discussion of the relevant views of naturalism's two main spokesmen in the twentieth century, John Dewey and W. V. Quine. The purpose of this chapter is to give the reader a clear perspective on the core dispositions underlying the naturalistic tradition. Chapter 3 then argues that characterizing naturalism as a view rather than a research program inevitably portrays naturalism either as a self-defeating thesis or as a view commitment to which would be inconsistent with the core dispositions of the tradition. Thus, I argue, the fairest and most plausible characterization of naturalism treats it as a research program—in particular, a research program wherein one treats the methods of science and those methods alone as basic sources of evidence.

In Part II, I argue that commitment to the naturalistic research program precludes one from accepting RMO and materialism. The argument turns on the prospects (or lack thereof) that naturalists have for solving what I call the Discovery Problem. Very roughly, the Discovery Problem is just the fact that intrinsic modal properties seem not to be discoverable by the methods of science. I describe this problem in Chapter 4, and I argue that *if* there is good reason to think that the problem cannot be solved, then naturalists cannot be justified in accepting RMO. In Chapters 5 and 6, I go on to argue that, indeed, there is good reason for thinking that the problem cannot be solved. In Chapter 7, I argue that, having been forced to give up RMO, naturalists are committed to accepting constructivism (but not necessarily idealism). I also argue that, in accepting constructivism, naturalists must give up materialism. Finally, I show that, once materialism has been given up, standard arguments against mind-body dualism turn their teeth toward ROM.

Finally, in Part III, I consider alternatives to naturalism. The alternatives that I consider are what I call *intuitionism*, a (shared) research program which takes the methods of natural science and rational intuition, but nothing else, as basic sources of evidence, and *supernaturalism*, a (shared) research program wherein one takes at least the methods of the natural sciences and religious experience as basic sources.⁹ Though I think that the consequences of naturalism described in Part II are sufficiently unattractive to motivate one to

⁹ I thus leave open the question whether supernaturalists will treat rational intuition, or anything else, as basic. As will become clear in Ch. 8, some versions of supernaturalism which treat rational intuition as a basic source will be viable, others will not be.

look for an alternative, not just any alternative will do. I argue in Chapter 8 that, unless one has intuitions that support the view that our world is the product of intelligent design (or unless one has very unusual beliefs or intuitions), intuitionism is self-defeating. I also argue that, though there might be empirical reason for thinking that intuition is reliable in some domains, this fact does not save intuitionism from self-defeat, nor is it of any use to a naturalist in trying to avoid the ontological consequences described in Part II. In Chapter 9, I argue that embracing supernaturalism offers our best hope of avoiding those consequences. However, I also argue that not just any version of supernaturalism will do the job, but only those that give rise to evidence for the conclusion that our world and, in particular, our cognitive faculties are the products of intelligent design.

The upshot, then, is that the most viable research programs that provide no evidence against the conclusion that our world is a world without design fare worse by certain widely accepted pragmatic standards than those that do. By the lights of any of the research programs discussed in this book, unless they somehow support the conclusion that our world is the product of design, both naturalism and intuitionism give rise to an ontology radically at odds with common belief and, surprisingly enough, also at odds with philosophical theses that many proponents of these research programs are most interested in accepting. Again, showing this is not sufficient to show that every right-thinking person *ought* to embrace supernaturalism or a design-friendly version of intuitionism (though, for all I know, this stronger claim may in fact be true). But the conclusion is significant nonetheless. For as a matter of fact the most effective way of persuading reflective people to reject one research program in favor of another is to make the rejected alternative appear less attractive than the former. I hope that I will be successful in persuading some naturalists to jump ship; but I make no pretense to having shown that it is irrational for them to do otherwise.