2

The Quinean Backdrop

In this chapter I will discuss the doctrines required to support the Quine/Putnam indispensability argument. Although a great deal of Quine's philosophy is interconnected, making the isolation of particular doctrines very difficult, I will argue that the two essential theses for our purposes—confirmational holism and naturalism—can be disentangled from the rest of the Quinean web. What is more, I will show that these two theses do not depend in any significant way on some of Quine's more controversial views about the philosophy of language.

2.1 Introducing Naturalism

Naturalism, in its most general form, is the doctrine that we ought to seek accounts of the nature of reality that are not "other-worldly" or "unscientific," but to be more precise than this is to immediately encounter trouble. For instance, David Papineau points out that "nearly everybody nowdays wants to be a 'naturalist', but the aspirants to the term nevertheless disagree widely on substantial questions of philosophical doctrine" (Papineau, 1993, p. 1). In one way this is not at all surprising, for, after all, there is no compulsion for all naturalists to agree on other philosophical stances, distinct from naturalism, and such stances, when combined with naturalism, presumably yield widely different results. It all depends on what you mix your naturalism with.

There is, however, another reason for disagreement among naturalistic philosophers: Different philosophers use the word 'naturalism' to mean different things. Naturalism involves a certain respect for the scientific enterprise—that much is common ground—but exactly how this is cashed

out is a matter of considerable debate. For instance, for David Armstrong naturalism is the doctrine that "nothing but Nature, the single, all-embracing spatio-temporal system exists" (Armstrong, 1978, Vol. 1, p. 138),[1] whereas, for Quine, naturalism is the "abandonment of the goal of a first philosophy" (Quine, 1981a, p. 72).[2]

One issue on which naturalistic philosophers disagree widely, and which is of fundamental importance for our purposes, is the ontological status of mathematical entities. We saw in the last chapter how the Quine/Putnam indispensability argument legitimates belief in mind-independent mathematical objects, and that this argument depends on naturalism. On the other hand, philosophers such as David Armstrong cite naturalism as grounds for rejecting belief in any such mind-independent abstract objects.[3]

While there is no way of preventing philosophers from mixing their naturalism with other philosophical doctrines (so long as the mix is consistent), there is good reason for requiring that the various, often contradictory, positions that fly under the banner of naturalism be disentangled, from one another. This is an extremely complicated matter, but in the next three sections I make some progress toward this goal. At the very least I hope to identify the difference between the varieties of naturalism that may be used to undermine mathematical realism and the Quinean variety.

2.2 Quinean Naturalism

Quine's aphoristic characterisations of naturalism are well known. In "Five Milestones of Empiricism" he tells us that naturalism is the

> abandonment of the goal of a first philosophy. It sees natural science as an inquiry into reality, fallible and corrigible but not answerable to any supra-scientific tribunal, and not in need of any justification beyond observation and the hypothetico-deductive method. (Quine, 1981a, p. 72)

And that:

[1] This conception of naturalism has resonances in D. C. Williams's (1944) materialism.

[2] Quine's conception of naturalism also has important predecessors. For example, Bertrand Russell identifies the new philosophical movement he calls 'realism' with the abandonment of "the claim to a special philosophic method or a peculiar brand of knowledge to be obtained by its means" (Russell, 1924, p. 69). He goes on to say that this new philosophy "regards philosophy as essentially one with science" and that "[i]t conceives that all knowledge is scientific knowledge" (Russell, 1924, pp. 69-70).

[3] I will look at Armstrong's rejection of abstract objects later in this chapter (section 2.4), and in considerably more detail in the next chapter.

> [t]he naturalistic philosopher begins his reasoning within the inherited world theory as a going concern. He tentatively believes all of it, but believes also that some unidentified portions are wrong. He tries to improve, clarify, and understand the system from within. He is the busy sailor adrift on Neurath's boat. (Quine, 1981a, p. 72)

The aphorisms are useful, but they also mask a great deal of the subtlety and complexity of Quinean naturalism. Indeed, the subtleties and complexities of naturalism are far greater than one would expect for such a widely held and intuitively plausible doctrine. We would thus be well served spending a little time coming to better understand Quinean naturalism.

As I see it, there are two strands to Quinean naturalism. The first is a normative thesis concerning how philosophy ought to approach certain fundamental questions about our knowledge of the world. The advice here is clear: look to science (and nowhere else) for the answers. Science, although incomplete and fallible, is taken to be the best guide to answering all such questions. In particular, "first philosophy" is rejected. That is, Quine rejects the view that philosophy precedes science or oversees science. This thesis has implications for the way we should answer metaphysical questions: We should determine our ontological commitments by looking to see which entities our best scientific theories are committed to. Thus, I take it that naturalism tells us (1) we ought to grant real status only to the entities of our best scientific theories and (2) we ought to (provisionally) grant real status to all the entities of our best scientific theories. For future reference I'll call this first strand of Quinean naturalism the *no-first-philosophy thesis* (or NFPT for short) and its application to metaphysics the *Quinean ontic thesis*.

It is worth pointing out that the Quinean ontic thesis is distinct from a thesis about how we determine the ontological commitments of *theories*. According to this latter thesis, the ontological commitments of theories are determined on the basis of the domain of quantification of the theory in question.[4] Call this thesis *the ontological commitments of theories thesis*. One could quite reasonably believe the ontological commitments of theories thesis without accepting the Quinean ontic thesis. For instance, I take it that Bas van Fraassen (1980) is such a person. He accepts that our current physics is committed to entities such as electrons and the like, but it does not follow that he believes that it is rational to believe in these entities in order to believe the theory. The ontological commitments of theories thesis is purely descriptive, whereas the Quinean ontic thesis is, as I've pointed out, normative. From here on I shall be concerned only with the Quinean

[4] See Quine (1948, pp. 12–13) for details.

ontic thesis, but it is worth bearing in mind the difference, because I don't think that the ontological commitments of theories thesis rightfully belongs to the doctrine of naturalism. It is *an* answer to the question of how we determine the ontological commitments of theories, but it is not the only naturalistic way such questions can be answered.

The second strand of Quinean naturalism is a descriptive thesis concerning the subject matter and methodology of philosophy and science. Here naturalism tells us that philosophy is continuous with science and that together they aim to investigate and explain the world around us. What is more, it is supposed that this science–philosophy coalition is up to the task. That is, all phenomena are in principle explicable by science. For future reference I'll call this strand the *continuity thesis*.

Although it is instructive to distinguish the two strands of Quinean naturalism in this way, it is also important to see how intimately intertwined they are. First, there is the intriguing interplay between the two strands. NFPT tells us that we ought to believe our best scientific theories and yet, according to the continuity thesis, philosophy is part of these theories. This raises a question about priority: In the case of a conflict between philosophy and science, which gets priority? I'll have more to say on this issue in section 5.2, but for now it will suffice to say that philosophy does not occupy a privileged position. That much is clear. But it also appears, from the fact that philosophy is seen as part of the scientific enterprise, that science (in the narrow sense—i.e., excluding philosophy) occupies no privileged position either.

The second important connection between the two strands is the way in which the continuity thesis lends support to NFPT. The traditional way in which first philosophy is conceived is as an enterprise that is prior and distinct from science. Philosophical methods are seen to be a priori while those of science are a posteriori. But accepting the continuity thesis rules out such a view of the relationship between philosophy and empirical science. Once philosophy is located within the scientific enterprise, it is more difficult to endorse the view that philosophy oversees science. I'm not claiming that the continuity thesis entails NFPT, just that it gives it a certain plausibility.[5]

In the discussion so far I've glossed over the question of what constitutes our best scientific theories.[6] I'll have something to say about how we answer this question in section 4.3. What is important in the meantime is that there

[5] Indeed, the continuity thesis cannot entail NFPT since the former is descriptive and the latter normative.

[6] There is also the question of what constitutes a *scientific* theory as opposed to a *non-scientific* theory. I won't enter into that debate here: I'll assume that we have at least an intuitive idea of what a scientific theory is.

is room for disagreement about what our best scientific theories are.

Now let me say something about why one ought to embrace naturalism. I won't embark on a general defence of naturalism—that would be far too ambitious. Instead, I shall ask that you accept, for the sake of the discussion, some suitably broad sense of this doctrine.[7] I suggest that we accept something like the conception of naturalism that I opened section 2.1 with:

> We ought to seek only scientific accounts of reality.

(Where "scientific" is construed fairly broadly.) What I need to do now is convince you of *Quinean* naturalism. In order to do this, I'll need to identify what is distinctive about Quinean naturalism. In particular, I need to identify and justify those features of the Quinean conception of naturalism that support the crucial first premise of argument 2 on page 11.

First let me mark out the common ground. Naturalists of all ilks agree that we should look only to science when answering questions about the nature of reality. What is more, they all agree that there is at least prima facie reason to accept all the entities of our best scientific theories. That is, they all agree that there is a metaphysical component to naturalism. So they are inclined to accept the first part of the Quinean ontic thesis (the "only" part) and are inclined to, at least provisionally, accept the second part (the "all" part). (Most naturalists believe that naturalism entails scientific realism but they are inclined to be a little reluctant to embrace *all* the entities of our best scientific theories.)[8] What I take to be the distinctive feature of Quinean naturalism is the view that our best scientific theories are continuous with philosophy and are not to be overturned by first philosophy. It is this feature that blocks any first-philosophy critique of the ontological commitments of science. Consequently, it is this feature of Quinean naturalism that is of fundamental importance to the indispensability argument.

Now defences of such fundamental doctrines as naturalism are hard to come by. Typically such doctrines are justified by their fruits. So in order to defend Quinean naturalism over other versions (I've already mentioned that I won't be providing a general defence of naturalistic philosophy), I'll examine some of the consequences of the Quinean position. This examination will occupy us on and off for much of the remainder of this book (especially in the next chapter). Let me begin here by showing how Quinean

[7] It is worth bearing in mind that the primary targets of the indispensability argument are scientific realists disinclined to believe in mathematical entities. These scientific realists typically subscribe to some form of naturalism, so my acceptance, without argument, of a broadly naturalistic perspective is not as serious an assumption as it may first seem.

[8] For example Keith Campbell (1994) advocates "selective realism" and Quine restricts commitment to indispensable entities.

naturalism provides a nice defence against scepticism.

The sceptic might ask what justification we have for postulating physical objects from the meagre input of certain two-dimensional electromagnetic irradiation patterns on our retinas. Our theories clearly outstrip our evidence for them, so it is the business of epistemology to give a factual account of the relation between the two.[9] Quine points out, however, that "the skeptical challenge springs from science itself, and that in coping with it we are free to use scientific knowledge" (Quine, 1974, p. 3). After all, the idea of being deceived about physical objects, for instance, depends on science in two ways: (1) the deception consists in believing something other than the scientific picture of the world and (2) it is science itself that informs us that our data about the world is both incomplete and fallible. If scepticism originates within science, it is only reasonable that epistemologists are justified in using whatever portion of science they require to combat scepticism. From the point of view of Quine's naturalised epistemology, there is no more secure vantage point than the vantage point of our best scientific theories. Thus, the naturalised epistemologist "no longer dreams of a first philosophy, firmer than science, on which science can be based; he is out to defend science from within, against its self doubts" (Quine, 1974, p. 3).

2.3 The Methodologies of Philosophy and Science

In this section I continue my defence of Quinean naturalism by defusing an objection that it faces. The objection is that Quinean naturalism (in particular the continuity thesis) fails to acknowledge an important methodological difference between science and philosophy. After all, so the objection goes, it is clear that philosophy proceeds by a priori methods, such as thought experiments and deduction, whereas science proper proceeds by a posteriori methods, the celebrated scientific method of hypotheses and observation, typically involving *real* experiments and *in*duction. In reply I will rehearse Quine's argument against the possibility of a priori knowledge. In addition to this, I wish to cast doubt on the legitimacy of characterising philosophy as proceeding by a priori methods and science as proceeding by a posteriori methods. Moreover, the discussion of this objection will help elucidate certain features of Quinean naturalism—namely, its connection to confirmational holism.

[9] Quine calls this relation "[t]he relation between the meager input and the torrential output" (Quine, 1969a, p. 83).

2.3.1 The "Two Dogmas" Argument

In his famous article "Two Dogmas of Empiricism," Quine launches a two-fold attack on the analytic/synthetic distinction. Since his arguments here are well known, I won't spend too much time on them; I'll just refresh our memories with a quick review. The first part of this argument is to show that there is no non-circular definition of "analytic." For instance, he argued that we cannot define analyticity by way of the notion of synonymy. As Graham Priest (1979) points out, however, this circularity argument is not so much a condemnation of analyticity, since many important concepts can be defined only in circular terms.[10] The point here is that the circularity argument prevents a defence of analyticity by appeal to synonymy when analyticity comes under fire, because synonymy, being part of the circle, is just as much under fire as analyticity. Quine's attack proper, then, is an argument from the history of science that *no* belief can be held onto no matter what. I can do no better than to quote Quine here:

> Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system. Even a statement very close to the periphery can be held true in the face of recalcitrant experience by pleading hallucination or by amending certain statements of the kind called logical laws. Conversely, by the same token, no statement is immune to revision. Revision even of the logical law of excluded middle has been proposed as a means of simplifying quantum mechanics; and what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle? (Quine, 1951, p. 43)

The main point is that the history of science has taught us that what were once considered analytic truths, such as that Pythagoras's theorem holds in our world[11] or that any massive body can be accelerated without bound, have been given up in order to cohere with new and better scientific theories. Thus, by an inductive argument from such examples, we conclude that there are no analytic truths.[12]

Putnam (1976) points out that the idea of an analytic truth as one that is confirmed no matter what is quite different from the original Kantian

[10] Indeed, if Quine is correct about language, this is true of *all* concepts!

[11] By this I simply mean to rule out the claim that Pythagoras's theorem holds in abstract Euclidean spaces.

[12] Perhaps the conclusion here seems a bit stronger than the argument will support, but at the very least we can conclude that even if there are analytic truths, we are notoriously bad at recognising them. We should thus give them no privileged place in science. This weaker conclusion is all that Quine really requires, and certainly all that I require in this discussion.

notion of "the predicate contained in the concept"; it is much more like the traditional notion of apriority. So if we take Quine's argument from the history of science on its own, we have an argument against the a priori/a posteriori distinction. Furthermore, the argument, thus construed, *does* stand on its own; it does not depend on the circularity argument, which has attracted some criticism since the publication of "Two Dogmas of Empiricism."[13] We see then that if the a priori/a posteriori distinction cannot be maintained, the claim that Quinean naturalism fails to recognise this distinction is misguided. Note, however, that for Quine some beliefs may be more central than others, and so may have an a priori "feel" to them but strictly speaking they are not a priori.

It is also worth noting explicitly that in the foregoing defence of the continuity thesis we have seen an important consequence of Quinean naturalism: some sort of holism about our scientific theories. The fact that we cannot distinguish between a priori and a posteriori portions of our theory and also the fact that it seems that isolated hypotheses do not enjoy empirical confirmation or disconfirmation—only bodies of hypotheses may be said to be confirmed or disconfirmed—suggests confirmational holism. While it's clear that there is a close relationship between naturalism and holism, I do not wish to take too much for granted here. I do not wish to presuppose that the confirmational holism required for the success of the indispensability argument is written into naturalism. It would be nice if it were, but I think it's safer to suppose that it is not. I will thus argue for confirmational holism separately in section 2.5. For now I'll continue my defence of Quinean naturalism against the charge that it fails to respect the methodological differences between philosophy and science.

2.3.2 How Useful Is the Distinction?

It seems that even if one accepts the argument from the preceding section, it might be objected that Quinean naturalism fails to distinguish between the quite different methodologies of philosophy and science: The former proceeds by pseudo a priori methods; that is, by relying on beliefs central to science, such as logic, whereas empirical science proceeds by clearly a posteriori methods. This objection, though, is again misguided. The history of science is littered with examples of science proceeding by apparently a priori methods. This is the domain of the theoretical scientist, who must tease out consequences of theories using deduction and thought experiments. To give but one example, Galileo's famous law that all bodies

[13] See, for instance, Priest (1979), Grice and Strawson (1956), and Putnam (1976) for some of these criticisms.

fall to earth with the same velocity regardless of their mass was derived, it seems, by a thought experiment, which showed that the received view from Aristotelian physics (heavier objects fall more quickly than lighter ones) was inconsistent.[14] Furthermore, this law flew in the face of empirical evidence, as stones were repeatedly observed to fall faster than leaves. The success of Galileo's law over the received view demonstrates that internal consistency is quite rightly seen as a more important feature of a scientific theory than empirical adequacy, and since determining internal consistency is the business of the theoretical scientist, pseudo a priori methods are not only part of science, they are an *important* part of science.

The converse, however, does not seem true; philosophers do not engage in empirical inquiry. I have two things to say in relation to this. The first is that this in itself doesn't seem enough to exclude philosophy from the domain of science. After all, as we have already seen, theoretical physicists don't perform experiments either, and yet no one wishes to bar them from the science club! The other relevant point here is that while philosophers don't engage in empirical inquiry themselves, they are certainly not unaware of developments in modern experimental science. For example, those philosophers working in philosophy of mind must keep a close eye on developments in experimental psychology, computer science, and neurophysiology.

Perhaps, though, there is some fundamental difference between the philosopher's thought experiments and the scientist's. Frank Jackson suggests one way someone might wish to distinguish between the two is that scientific thought experiments tell us something about the world, whereas philosophical thought experiments tell us something about the way we use our language (Jackson, 1998, p. 78). For instance, Galileo's thought experiment tells us that all objects fall with the same velocity irrespective of mass, whereas Hilary Putnam's famous "twin earth" thought experiment tells us something about the way we use the term "water."[15] This is not

[14] Consider two bodies, $\mathcal{B}_1$ and $\mathcal{B}_2$, where $\mathcal{B}_1$ is heavier than $\mathcal{B}_2$. If we were to tie these two bodies together with a piece of twine of sufficient strength and allow them to fall freely from the same height, Aristotelian physics would tell us: Since $\mathcal{B}_1$ is heavier than $\mathcal{B}_2$ we should expect $\mathcal{B}_1$ to fall faster until the piece of twine becomes tight, and then $\mathcal{B}_2$ will retard $\mathcal{B}_1$'s motion so the velocity of the system will be *slower* than if $\mathcal{B}_1$ had been dropped on its own. On the other hand, the mass of the system is greater than the mass of $\mathcal{B}_1$'s alone and so the system's velocity should be *faster* than if $\mathcal{B}_1$ had been dropped on its own.

[15] Putnam (1973) invites us to consider earth and twin earth. These two worlds are exactly alike except for the following: In the former world the occupants use the word "water" to refer to H_2O and in the latter the occupants use it to refer to some other substance XYZ, which is otherwise similar to H_2O. Putnam's (and many others') intuition is that we, occupants of earth, would not call the twin-earth substance, XYZ,

at all clear to me though. Surely all Galileo's thought experiment shows is that there is an inconsistency in Aristotelian physics, which was rectified by dropping the proposition that heavier objects fall faster than lighter ones. The contradiction *might* have been resolved by altering the way we use some of the crucial words. Indeed, it is true that Galileo's law that all bodies fall with the same velocity is false unless what we mean by "fall" is "fall *in a vacuum*." So to some extent, at least, Galileo's thought experiment does help to clarify our use of the relevant bits of language.

Perhaps a better example is Einstein's special relativity thought experiment. Much of Einstein's work in the special theory of relativity was concerned with clarifying what is *meant* by length, time, velocity, and simultaneity rather than telling us how the world is (although it clearly does the latter as well). Indeed, Einstein himself takes this view of at least part of his work in the 1905 essay on special relativity, as illustrated by the following passage:

> Thus with the help of certain imaginary physical experiments we have settled what is to be understood by synchronous stationary clocks located at different places, and have evidently obtained a definition of "simultaneous," or "synchronous," and of "time." (Einstein, 1905, p. 40)

Take for example the classical (Galilean) addition of velocities formula:

$$V_{\mathcal{G}} = v_a + v_b \tag{2.1}$$

and compare it with the relativistic formula

$$V_{\mathcal{R}} = \frac{v_a + v_b}{1 + \frac{v_a v_b}{c^2}} \tag{2.2}$$

where v_a is the velocity of some body a (moving with uniform velocity) relative to some inertial frame $\mathcal{F}$, v_b is the velocity of some other body b (also moving with uniform velocity) relative to a, $V_{\mathcal{G}}$ is the velocity of b relative to $\mathcal{F}$ in the Galilean/Newtonian theory, $V_{\mathcal{R}}$ is the velocity of b relative to $\mathcal{F}$ in the special theory of relativity, and c is the speed of light in a vacuum. It is clear that in general $V_{\mathcal{G}} \neq V_{\mathcal{R}}$ and it is well known that equation (2.2) has replaced (2.1) as the correct formula for evaluating addition of velocities.[16]

"water." Putnam's conclusion is that the meaning of "water" is not fixed entirely by the internal states of speakers.

[16] Of course (2.1) is still used for velocities small in relation to c, but it is understood that values so obtained are estimates.

It might well be argued that there is a change of meaning in some of the relevant terms (such as "velocity") for such a revision of the addition of velocities formula to occur. Similarly the replacement of the classical kinetic energy and momentum formulae in special relativity indicates changes of meanings of these terms in the new theory. The moral of all this is simply that it is not at all clear that scientific thought experiments are concerned *only* with the way the world is. They also shed light on the way we use bits of the language of the relevant theory and, in particular, which concepts are the key ones. For instance, special relativity tells us that $V_{\mathcal{R}}$ is the important concept when adding velocities, not $V_{\mathcal{G}}$, despite the latter's intuitive appeal.

As a final attempt to maintain some sort of distinction between philosophy and science, someone might argue that Jackson's distinction between scientific thought experiments and philosophical thought experiments *can* be upheld; it's just that Einstein, for instance, was indulging in *philosophical* analysis in his 1905 essay. But surely if the distinction in question cannot be maintained by appeal to scientific and philosophical practices, then one ought to wonder whether the distinction is really picking out anything more significant than the distinction between, say, physics and chemistry. At the very least, it is clear that the boundary between philosophy and science is vague and this is enough to support the acceptance of the continuity thesis.

To sum up, then. First, if Quine is right and there really is no a priori knowledge,[17] then there can be little substance to the thought that philosophy is a priori and science a posteriori. Second, even if you disagree with Quine on this, it is not at all clear that the cleavage between the a priori and the a posteriori corresponds to any significant cleavage at all, let alone to the cleavage between philosophy and science. The fact that Quinean naturalism fails to respect such a cleavage is far from a deficiency of the position; it is one of its great strengths.

So, having dealt with what I think are the more obvious objections to Quinean naturalism, I now turn to the task of discussing its main rival.

[17] This may not be strictly correct. Hilary Putnam (1979b) suggests that there is at least one a priori truth, but this is of little interest since it is not the sort of proposition that is likely to play a significant role in any scientific theory. For the record, Putnam claims that the statement "Not every statement is both true and false" is true a priori since "to deny that statement would be to forfeit rationality itself" (p. 129).

2.4 The Causal Version of Naturalism

One interesting way in which someone can give an alternative account of naturalism is by giving an alternative account of what our best scientific theories are.[18] This can be done by placing some restriction on current scientific theories, such as by believing some portion of these theories and remaining instrumental about the rest.[19] This style of account is considered very seductive by many, and in this section I will consider one such account advanced by David Armstrong.

Armstrong defines naturalism "as the doctrine that reality consists of nothing but a single all-embracing spatio-temporal system" (Armstrong, 1980a, p. 149). This conception of naturalism has an important consequence (at least according to Armstrong): We should believe in only causally active entities (or perhaps, more generously, *potentially* causally active entities). I shall refer to this as the *Eleatic Principle*,[20] or causal requirement. According to Armstrong, non-spatio-temporally located entities would be incapable of acting on particulars, and so can play no explanatory role in science. He concludes that we have no rational reason to postulate them. Elsewhere he admits that the latter argument is less than conclusive, but nonetheless "gives us good reason for denying the existence of such entities" (Armstrong, 1989, p. 7). I will not discuss Armstrong's argument in any detail here (I'll have more to say about it in chapter 3, along with other justifications of the Eleatic Principle).

I've already noted the disagreement between Quine and Armstrong on the ontological status of mathematical entities. For Armstrong, nominalism follows fairly directly from the Eleatic Principle. This is in stark contrast with Quinean naturalism, which, as we've seen already, supports Platonism. Clearly the Eleatic Principle is the crucial difference. It is interesting, however, to ask after the status of this principle. If Armstrong is using the Eleatic Principle to overrule science on ontological matters—science being

[18] Clearly the alternative account of what our best scientific theories are had better not be too radical lest it cease to qualify as naturalistic. Construing astrology and literal readings of the Bible as our best scientific theories would not do.

[19] Indeed, it seems scientists themselves are instrumental about the more fanciful areas of science; that is, those which are not firmly supported by theory and experimental evidence. This instrumentalism seems to shift to realism as the area in question becomes better understood. However, I take such instrumentalism to be simply an indication that the area of theory in question *isn't* (yet) part of the best theory.

[20] It is called "the Eleatic Principle" after a passage from Plato's *Sophist* in which the Eleatic stranger suggests that causal power is the mark of being (Plato, 1935, pp. 247d–e). David Armstrong cites this passage (1978, Vol. 2, pp. 45–46) and Graham Oddie (1982) coined the phrase "the Eleatic Principle."

committed to mathematical entities and Armstrong ruling against them—then he is guilty of practising first philosophy. Perhaps, more plausibly, Armstrong sees the Eleatic Principle as part of our best scientific theory.

Whether the Eleatic Principle is considered part of a first philosophy or part of our best science doesn't matter too much; it's clear that Armstrong's naturalism is committed to a causal test of this kind. Either way, this test will require justification and I devote the whole of the next chapter to the important task of examining the arguments put forward for it. For now let me join David Lewis in expressing general concerns about philosophers who wish to interfere with science or mathematics because of commitment to principles such as the Eleatic Principle. (Lewis is specifically speaking of those who wish to interfere with mathematics based on the view that there are no classes, but the general lesson is clear.) He says:

> I am moved to laughter at the thought of how *presumptuous* it would be to reject mathematics for philosophical reasons. How would *you* like the job of telling the mathematicians that they must change their ways, and abjure countless errors, now that *philosophy* has discovered that there are no classes? Can you tell them, with a straight face, to follow philosophical argument wherever it may lead? If they challenge your credentials, will you boast of philosophy's other great discoveries: that motion is impossible, that a Being than which no greater can be conceived cannot be conceived not to exist, that it is unthinkable that anything exists outside the mind, that time is unreal, that no theory has ever been made at all probable by evidence (but on the other hand that an empirically ideal theory cannot possibly be false), that it is a wide-open scientific question whether anyone has ever believed anything, and so on, and on, *ad nauseam*?
>
> Not me! (Lewis, 1991, p. 59)

Nor me!

2.5 Holism

Holism comes in many forms. Even in Quine's philosophy there are at least two different holist theses. The first is what is usually called *semantic holism* (although Quine calls it *moderate holism* (1981a, p. 71)) and is usually stated, somewhat metaphorically, as the thesis that the unit of meaning is the whole of the language. As Quine puts it:

> The idea of defining a symbol in use was ... an advance over the impossible term-by-term empiricism of Locke and Hume. The statement, rather than the term, came with Bentham[21] to be recognized as the unit accountable to an empiricist critique. But what I am now urging is that even in taking the statement as unit we have drawn our grid too finely. The unit of empirical significance is the whole of science. (Quine, 1951, p. 42)

Semantic holism is closely related to Quine's denial of the analytic/synthetic distinction and his thesis of indeterminacy of translation. He argues for the former in a few places, but most notably in "Two Dogmas of Empiricism" (1951), while the latter is presented in *Word and Object* (1960).

The other holist thesis found in Quine's writings is *confirmational holism* (also commonly referred to as the Quine/Duhem thesis). As Fodor and Lepore point out (1992, pp. 39–40), the Quine/Duhem thesis receives many different formulations by Quine and it is not clear that all these formulations are equivalent. For example, in *Pursuit of Truth* Quine writes:[22]

> [T]he falsity of the observation categorical[23] does not conclusively refute the hypothesis. What it refutes is the conjunction of sentences that was needed to imply the observation categorical. In order to retract that conjunction we do not have to retract the hypothesis in question; we could retract some other sentence of the conjunction instead. This is the important insight called *holism*. (Quine, 1992, pp. 13–14)

And in "Two Dogmas of Empiricism," in a much quoted passage, he suggests that "our statements about the external world face the tribunal of sense experience not individually but only as a corporate body" (Quine, 1951, p. 41). In a similar vein in "On Mental Entities" he tells us:

> As Pierre Duhem urged, it is the system as a whole that is keyed to experience. It is taught by exploitation of its heterogeneous and sporadic links with experience, and it stands or falls, is retained or modified, according as it continues to serve us well or ill in the face of continuing experience. (Quine, 1953, p. 222)

[21] Interestingly, in the original version of the paper in *Philosophical Review* the reference here is to Russell. This was exchanged for Frege in the first edition of *From a Logical Point of View* and, finally, for Bentham in the second edition of *From a Logical Point of View* (Fodor and Lepore, 1992, p. 216).

[22] Cf. Duhem (1906, p. 187) for a similar statement of the thesis.

[23] By "observation categorical" Quine simply means a statement of the form "whenever P, then Q." For example, "where there's smoke, there's fire."

In the last two of these three passages Quine emphasizes the *confirmational* aspects of holism—it's the whole body of theory that is tested, not isolated hypotheses. In the first passage he emphasizes *disconfirmational* aspects of holism—when our theory conflicts with observation, any number of alterations to the theory can be made to resolve the conflict. Despite the difference in emphasis, I take it that these theses are equivalent (or near enough). Moreover, I take it that they are all true, modulo some quibbles about how much theory is required to face the tribunal at any time.

It's somewhat ironic that Quine argues for confirmational holism (which in some form or another is a relatively uncontroversial thesis)[24] from his semantic holism, which is one of the most controversial parts of Quine's philosophy. The debate about the rejection of the analytic/synthetic distinction, in particular, is still raging half a century after the publication of "Two Dogmas of Empiricism." I don't deny that confirmational holism follows from semantic holism; it's just that there are easier, less controversial roads to confirmational holism. Since it's only confirmational holism that we require for the indispensability argument, I intend to explore these other roads and thus avoid the semantic holism debate.

Both Duhem (1906) and Lakatos (1970) have argued for confirmational holism without any (obvious) recourse to semantic considerations. They emphasize the simple yet undeniable point that there is more than one way in which a theory, faced with recalcitrant data, can be modified to conform with that data. Consequently, certain core doctrines of a theory may be held onto in the face of recalcitrant data by making suitable alterations to auxiliary hypotheses. This point is driven home by appeal to case studies from the (actual and imagined) history of science. Indeed, such is the influence of Duhem, Lakatos, and Quine on this point that now few require convincing. I will, however, present one (fictional) example, from Lakatos, to illustrate the point:

> The story is about an imaginary case of planetary misbehaviour. A physicist of the pre Einsteinian era takes Newton's mechanics and his law of gravitation, N, the accepted initial conditions, I, and calculates, with their help, the path of a newly discovered small planet, p. But the planet deviates from the calculated path. Does our Newtonian physicist consider that the the deviation was forbidden by Newton's theory and therefore that, once established, it refutes the theory N? No. He suggests that there must be a hitherto unknown planet p', which perturbs the path of p. He calculates the mass, orbit, etc. of this hypothetical planet and then asks an experimental

[24] Of course it is not completely uncontroversial—very little in philosophy is.

> astronomer to test his hypothesis. The planet p' is so small that even the biggest available telescopes cannot possibly observe it; the experimental astronomer applies for a research grant to build yet a bigger one. In three years time, the new telescope is ready. Were the unknown planet p' to be discovered, it would be hailed as new victory of Newtonian science. But it is not. Does our scientist abandon Newton's theory and his idea of the perturbing planet? No. He suggests that a cloud of cosmic dust hides the planet from us. He calculates the location and properties of this cloud and asks for a research grant to send up a satellite to test his calculations. Were the satellite's instruments (possibly new ones, based on a little-tested theory) to record the existence of the conjectural cloud, the result would be hailed as an outstanding victory for Newtonian science. But the cloud is not found. Does our scientist abandon Newton's theory, together with the idea of the perturbing planet and the idea of the cloud which hides it? No. He suggests that there is some magnetic field in that region of the universe which disturbs the instruments of the satellite. A new satellite is sent up. Were the magnetic field to be found, Newtonians would celebrate a sensational victory. But it is not. Is this regarded as a refutation of Newtonian science? No. Either yet another ingenious auxiliary hypothesis is proposed or ... the whole story is buried in a dusty volume of periodicals and the whole story never mentioned again. (Lakatos, 1970, pp. 100–101)

The story, although humourous in tone, does illustrate how much theory is employed for what might seem like straightforward scientific observations. Thus, the unit of confirmation or disconfirmation is not the single hypothesis but, rather, some substantial body of hypotheses. Moreover, a fairly good case can be made for thinking that in some instances this larger body of hypotheses is the whole of science.

Before leaving the doctrine of holism, I wish to consider one last question: Might one accept confirmational holism as stated, but reject the claim that mathematical propositions are one with the rest of science? That is, might it not be possible to pinpoint some semantic difference between the mathematical propositions employed by science and the rest, with empirical confirmation and disconfirmation reserved for the latter? Carnap (1937), with his appeal to "truth by convention," suggested precisely this. Quine, of course, denies that this can be done (1936; 1951; 1963), but the reasons for his denial would take us deep into semantic holism. For our purposes, it will suffice to note that there is no obvious way of disentangling the purely mathematical propositions from the main body of science. Our empirical theories have the so-called empirical parts intimately intertwined with the

mathematical. A cursory glance at any physics book will confirm this, where one is likely to find mixed statements such as: "planets travel in elliptical orbits"; "the curvature of space-time is not zero"; "the work done by the force on the particle is given by $W = \int_a^b \mathbf{F} \cdot d\mathbf{r}$."

Thus, even if you reject Quine's semantic holism and you think that mathematical and logical language is different in kind from empirical language, you need not reject confirmational holism. In order to reject confirmational holism, you would need (at the very least) to separate the mathematical vocabulary from the empirical in all of our best scientific theories. Clearly this task is not trivial.[25] If you still feel some qualms about confirmational holism, though, you may rest assured—this doctrine will be called into question in later chapters when I discuss some of the objections to the indispensability argument. For the moment, at least, I invite you to join me in accepting confirmational holism.

2.6 The First Premise Revisited

Let me close this chapter with a summary of how confirmational holism and Quinean naturalism combine to yield the first premise of argument 2 on page 11. First, you might wonder whether holism is required for the argument. After all, (Quinean) naturalism alone delivers something very close to the crucial first premise. (More specifically, the Quinean ontic thesis is very suggestive of the required premise.) As a matter of fact, I think that the argument can be made to stand without confirmational holism: It's just that it is more secure *with* holism. The problem is that naturalism is somewhat vague about ontological commitment to the entities of our best scientific theories. It quite clearly rules out entities *not* in our best scientific theories, but there seems room for dispute about commitment to some of the entities that *are* in these theories. Holism helps to block such a move since, according to holism, it is the whole theory that is granted empirical support.

So, naturalism tells us to look to our best scientific theories for our ontological commitments. We thus have provisional support for all the entities in these theories and no support for entities not in these theories. For reasons of parsimony, however, we may wish to grant real status to only those entities that are *indispensable* to these theories. However, we are unable to pare down our ontological commitments further by appealing to some distinction based on empirical support because, according to holism,

[25] As we shall see, in chapter 4, Hartry Field undertakes this task for reasons not unrelated to those I've aired here.

all the entities in a confirmed theory receive such support. In short, holism blocks the withdrawal of the provisional support supplied by naturalism. And that gives us the first premise of the Quine/Putnam indispensability argument.

In the next chapter I will continue the defence of Quinean naturalism. In particular, I will demonstrate some of the serious problems that face the Eleatic Principle. (As we have seen in this chapter, this principle is the central tenet of Quinean naturalism's main rival: the causal version of naturalism.) There are three aspects to this strategy; two relevant to the task of motivating Quinean naturalism and one that's more general. First, discrediting the Eleatic Principle undermines much of the plausibility of the causal version of naturalism, leaving Quinean naturalism looking all the more attractive. Second, I will show that when one considers how the Eleatic Principle might be defended against some of the objections I raise for it, one is drawn toward a position not unlike Quinean naturalism. This again helps fortify the Quinean position. Finally, if the Eleatic Principle were tenable, it would present a quite general problem for mathematical realists, independent of the brand of naturalism subscribed to. Undermining the Eleatic Principle thus clears the way for Platonist philosophies of mathematics.

3

The Eleatic Principle

In the last chapter I identified the crucial difference between Quinean naturalism and Armstrong's causal version of naturalism. The difference, I argued, is that the latter subscribes to the Eleatic Principle, while the former does not. In this chapter I will continue the defence of Quinean naturalism. In particular, I will highlight some problems faced by the Eleatic Principle and suggest that Quinean naturalism is better equipped to deal with these difficulties.

The Eleatic Principle or causal criterion, you will recall, is a causal test that entities must pass in order to gain admission to some philosophers' ontologies. This principle justifies belief in only those entities to which causal power can be attributed, that is, to those entities that can bring about changes in the world. The idea of such a test is rather important in modern ontology, since it is neither without intuitive appeal nor without influential supporters. Its sympathisers (if not supporters) have included David Armstrong (1978, Vol. 2, p. 5), Brian Ellis (1990, p. 22), and Hartry Field (1989, p. 68), to name but a few.

Clearly this principle requires some justification. In this chapter I will look at the arguments that have been put forward for such a principle and suggest some problems for each of these. Of course in such a survey I cannot provide a decisive refutation of the principle, but I do hope to show that, despite its intuitive appeal, the Eleatic Principle's main justifications either look ad hoc or do not justify a version of the principle that delivers the intuitively correct results about some fairly uncontroversial cases. This is not an entirely negative result though. Once we look at the shortcomings of the motivations for the Eleatic Principle, a more general principle suggests itself. This more general principle looks very much like Quine's