The Poverty of Analysis

David Papineau

Introduction

Many different ideas parade under the banner of philosophical naturalism. One is a thesis about philosophical method. Philosophy investigates reality in the same way as science. Its methods are akin to scientific methods, and the knowledge it yields is akin to scientific knowledge. This 'methodological naturalism' is to be distinguished from 'ontological naturalism' understood as a general view about the contents of reality. Ontological naturalism maintains that reality involves nothing more than the entities studied in the natural sciences and contains no supernatural or transcendent realm. While both ontological and methodological naturalism claim a species of affinity between philosophy and science, the two doctrines are largely independent.

Part of the task in understanding these matters is to bring definition to this pair of naturalist doctrines. A surprisingly wide range of philosophers wish to style themselves as naturalists, and by no means all understand either the methodological or ontological commitments of naturalism in the same way. My focus in this paper will be on methodological naturalism. I shall aim to refine and defend methodological naturalism as a thesis about philosophical method. The ontological dimension of naturalism will not feature in what follows.

Methodological naturalism asserts that philosophical investigation is like scientific investigation. Clearly more needs to be said before we can subject this claim to serious assessment. Nobody can doubt that the two enterprises are similar in some respects (both aim for precision and truth, say) and different in other respects (philosophers don't use particle accelerators). If methodological naturalism is to have any significant content, it needs to be specified in what respects philosophical and scientific methods are supposed to be alike.

I am going to argue that philosophy is like science in three interesting and non-obvious ways. First, the claims made by philosophy are synthetic not analytic: philosophical claims, just like scientific claims, are not guaranteed by the structure of the concepts they involve. Second, philosophical knowledge is a posteriori not a priori: the claims established by philosophers depend on the same kind of empirical support as scientific theories. And finally, to complete the traditional trio, the central questions of philosophy concern actuality rather than necessity: philosophy is primarily aimed at understanding the actual world studied by science, not some further realm of metaphysical modality.

I do not intend these claims in a revisionary spirit. I am not recommending that philosophers start doing something different. Here I diverge from other philosophers in the methodologically naturalist camp who take their position to require a shift in philosophical method—philosophers should get out of their armchairs and become more involved with active scientific research. This is not my view. When I say that philosophical investigation is akin to scientific investigation, I am not urging philosophers to change their ways. I think that most philosophy is just fine as it is, including philosophy that sticks to traditional methods of abstract theorizing, argument, and reflection on possible cases. My aim is to show that philosophy of this kind is already akin to science, not that it needs reforming in order to become so.

In what follows I shall avoid offering any positive characterization of philosophy, and in particular of what makes it different from science. For what it is worth, I do have some views about this. If pressed, I would say that philosophy is characteristically concerned with theoretical tangles. It deals with issues where deep-seated assumptions pull us in opposite directions and it is difficult to see how to resolve the tension. Because of this, the gathering of new empirical data is often (though by no means always) of no help in resolving philosophical problems. The characteristic philosophical predicament is that we have all the data we could want, but still cannot see how to resolve our theoretical problems.

Still, as I said, I am not going to commit myself to any positive characterization of philosophy. My argument does not need one. My intended subject matter is philosophy as it actually is, not a hypothetical philosophy that fits some set of prior specifications. Of course, this sociological dimension means that my claims are strictly speaking hostage to the activities of any philosophical eccentrics or extremists who deviate from my account of philosophical practice. But I hope that readers will understand my claims sympathetically in this respect. I don't want to show that everybody who has ever called themselves a 'philosopher' vindicates my claims about the nature of philosophy. It will be quite enough if I can establish my theses for those kinds of philosophy that most of you regard as mainstream.

Before proceeding, I need to qualify my claims in another respect. They do not apply equally straightforwardly to all philosophical subject matters. The areas that fit my claims best are the 'theoretical' branches of philosophy, including metaphysics, philosophy of mind, philosophy of language, and epistemology. Things become more complicated when we are dealing with areas of philosophy that trade in normative claims, or mathematical claims, or logical or modal claims. Part of the difficulty here is that the contents of these claims are themselves matters of philosophical debate, and so any attempt to show that they fit my theses about the nature of philosophy will itself become embroiled in these debates. As it happens, I think that most of the spirit of my theses about the nature of philosophy applies to these claims too, give or take a bit. But to show this would require far more space than I have available here. For present purposes it will be enough if I can show that my theses apply to the more easily interpretable claims of theoretical philosophy.

In what follows, I shall devote most of my attention to my first thesis. The next four sections will be about the synthetic nature of philosophical claims. After that, I shall devote my final two sections to the issues of a posterioricity and modality.

Theories and Concepts I

It might seem that my account of philosophy falls at the first hurdle, at least in so far as it is intended as non-revisionary. What about the many philosophers who proclaim themselves to be concerned with the analysis or explication of concepts? A wide and varied range of contemporary philosophers describe their own philosophical practice as in large part concerned with the elaboration of conceptual truths. Does this not immediately belie my first thesis that philosophy as it is currently practised deals with synthetic rather than analytic claims?

I say that that these philosophers misdescribe their own practice. They may claim that they are concerned with conceptual truths, but they are wrong. When we look more carefully at what they actually do, we can see that they are in fact concerned with synthetic and not analytic matters. Indeed their claims about their practice are not even supported by everything they <u>say</u> they do. I shall show that when these philosophers go on to fill out their account of philosophy, their own characterization of their practice is perfectly consistent with my first thesis.

Anybody who thinks that there are conceptual truths to be uncovered must suppose that the relevant concepts have some kind of structure. They must be constitutively linked to other concepts in such a way as to place constraints on their proper application. The idea is then that this structure can be uncovered by reflection and analysis, perhaps including reflection on what we would say about a range of possible cases.

An initial question to ask about this kind of putative conceptual structure is how it relates to <u>theories</u> involving the relevant concepts. By 'theories' I mean sets of claims with synthetic consequences. A simple theory of pain in this sense would be constituted by the two claims that (a) bodily damage typically causes pains and (b) pains typically cause attempts to avoid further damage. For note that together these two claims have the manifestly synthetic consequence that bodily damage typically causes attempts to avoid further damage. We can take it that everyday thought endorses theories like this about a wide range of philosophically interesting topics, including not only mental kinds like pain, but also such categories as persons, free will, knowledge, names, and so on—after

all, this is simply to assume that everyday thought includes various synthetic assumptions about these kinds.

It is widely supposed that there is a close connection between everyday concepts and everyday theories. But there are different views about the nature of this connection. In this section and the next I shall distinguish 'verificationist' from 'descriptivist' accounts of the connection between concepts and theories. As we shall see, neither account lends any support to the thesis that philosophy is centrally concerned with analytic truths.

Let me start with the verificationist account. This assumes that possessing a concept is a matter of being disposed to use that concept in a certain way. In particular, it is a matter of applying the concept in response to perceptual experiences and other judgements, and of drawing further inferences in turn from judgements involving the concept.

Given this account of concepts, which concepts a thinker possesses will depend on what theories that thinker accepts. This is because accepting a theory affects your dispositions to apply the concepts it involves. For example, if you accept the phlogiston theory of chemistry, then you will hold that burning causes air to become saturated with phlogiston, that dephlogisticated air is easily breathable, and so on. Similarly, if you accept the baby theory of pain offered above, then you will be disposed to hold that those with bodily damage are in pain, and that those who are in pain will engage in avoidance behaviour. From the verificationist perspective, then, your commitment to these theories determines your concepts phlogiston and pain. Since the theories affect your dispositions to apply the concepts, they determine your concepts themselves.

Now, one issue which arises at this point is how much of accepted theory is supposed to make such a constitutive contribution to concepts? Do all accepted assumptions make a difference, or only some distinguished subset—and if the latter, what distinguishes this subset? However, we can by-pass these familiar questions here. The points I now want to make are quite orthogonal to this issue. They will apply to any view that takes the

acceptance of sets of synthetic claims to affect concepts, however those claims might be identified.

A more basic issue is whether it makes sense to suppose that the mere possession of a concept can require a thinker to embrace synthetic commitments. Some of you may suspect that there must be something amiss with an account of concepts which implies this. However, not all philosophers share this worry. Robert Brandom, for instance, does not. He is insistent that concept possession incurs synthetic commitments. For example, after discussing Michael Dummett's example of the concept Boche, Brandon says that this

'. . . shows how concepts can be criticized on the basis of substantive beliefs. If one does not believe that the inference from German nationality to cruelty is a good one, then one must eschew the concept Boche' (Brandom 1994, p 126).

Again, a page later, he explains

'The concept temperature was introduced with certain criteria or circumstances of appropriate application and with certain consequences of application. . . . The proper question to ask in evaluating the introduction and evolution of a concept is . . . whether the inference embodied . . . is one that ought to be endorsed' (Brandom 1994, p 127).

This account of concepts plays an important part in Brandom's understanding of the philosophical enterprise. Brandon takes philosophy to be centrally concerned with the explication of concepts. But for Brandom this is not a merely descriptive enterprise. Since concepts carry synthetic commitments, it is possible to criticize concepts on the grounds that these commitments are unwarranted. Brandom is quite explicit about this:

'I see the point of explicating concepts rather to be opening them up to rational criticism . . . Defective concepts distort our thought and constrain us by limiting the propositions

and plans we can entertain . . . Philosophy, in developing and applying tools for the rational criticism
of concepts, seeks to free us from these fetters, by bringing the distorting influences out into the light of conscious day, exposing the commitments implicit in our concepts as vulnerable to rational challenge and debate' (Brandom 2001, p 77).

The notion that concepts have synthetic implications and are therefore open to criticism is not peculiar to Brandom. It is a commonplace of much discussion of the role of concepts in philosophy. Thus in a recent discussion of philosophical intuitions Alvin Goldman asserts that

'A concept that embeds a bad theory is of dubious worth' (Goldman 2007, p 22).

Again, to take just one further example, in a recent paper on moral concepts we find Richard Joyce arguing that

'Sometimes discoveries lead us to decide that a concept (e.g., phlogiston or witch) is hopeless; sometimes we prefer to revise the concept, extirpate the problematic element, and carry on much as before' (Joyce 2006, p 142).

I alluded a moment ago to the oddity of a view of concepts on which the mere possession of a concept can incur synthetic commitments. In fact there are further aspects of the verificationist approach that should make us even more suspicious of its account of concepts. For a start, verificationism implies that theoretical change inevitably leads to conceptual change. If you alter your theoretical assumptions involving some concept, perhaps because empirical evidence has shown that these assumptions are mistaken, then you will change your dispositions to apply that concept—and so, according to verificationism, will end up with a new concept. 'Meaning incommensurability' then quickly follows: adherents of different theories must mean different things even when they use the same words, and so cannot communicate with each other in a common language. In the extreme case, this implies that those who reject the ontological

commitments of some theory cannot use the language of that theory to convey this. Since I do not accept the phlogiston theory, I cannot mean the same by 'phlogiston' as the theory's adherents, and so cannot communicate my disagreement to them by saying 'There is no phlogiston'.

For my money, these points are enough to discredit the verificationist account of the relation between concepts and theories. Still, I do not need to take a stand on the nature of concepts here. This is because I have no objection to what verificationists like Brandom say about philosophical practice itself, as opposed to their funny way of thinking about concepts. Brandom says that philosophy is concerned with concepts, and then explains that for him this means that philosophy should identify the synthetic assumptions that guide our use of concepts, and criticize these assumptions when necessary. This vision of philosophical practice is entirely in accord with my first thesis that philosophy is concerned with synthetic claims.

When philosophers like Brandom say that they are explicating concepts, an unwary audience might conclude that this means that that they are not concerned with synthetic matters. But by this conclusion is belied, not only by their philosophical practice, but also by their official explanation of this practice. If the possession of concepts requires commitment to synthetic claims, and explication of these concepts involves the assessment of these claims, then there is no difference between conceptual explication and ordinary synthetic theorizing.

Theories and Concepts II

Even if we reject verificationist thinking, there may still be a close connection between concepts and theories. Suppose that we dismiss the notion that concept possession hinges on dispositions to apply concepts. Then our concepts will not depend on which theories we <u>accept</u>. But they may still depend on which theories we <u>understand</u>.

To see how this might work, suppose that T(F) is some synthetic theory involving the concept F. Then it is open to us to regard the concept F as having its reference is fixed via the description 'the Φ such that T(Φ)'. That is, F can be understood as referring to the unique Φ that satisfies the assumptions in T, if there is such a thing, and to fail of reference otherwise. In this spirit, we might regard pain as referring to the mental state, if there is one such, that is typically caused by damage and gives rise to avoidance behaviour, and phlogiston as referring to the substance, if there is one such, that is emitted in combustion and absorbed during chemical reduction; and so on.

On this descriptivist account, there is still a close connection between concepts and theories. But your concepts no longer depend on which theories you accept. Which theories you accept will of course affect your dispositions to apply concepts. But for non-verificationists this won't make a difference to the concepts themselves. Even though I reject the phlogiston theory, and so apply the concept phlogiston quite differently from the eighteenth-century chemists who endorsed the theory, this doesn't stop me having the same concept as they had. For we can all understand the concept phlogiston as equivalent to the relevant description—the putative substance that is emitted during combustion and absorbed during reduction—independently of our divergent views as to whether this description is satisfied.

In line with this, note that on the descriptivist account of concepts no synthetic commitments are incurred by the mere possession of a concept. Somebody who possesses a concept F defied by some theory T will be committed to the 'Carnap sentence' of the theory—if (EΦ)(T(Φ)), then T(F))—but this claim will be analytic not synthetic. For example, if you have the concept phlogiston you will be committed to the relevant analytic claim, that if there is a substance emitted during combustion and absorbed during reduction, then it is phlogiston. But you needn't thereby be committed to the synthetic commitments of the phlogiston theory itself.

From the perspective of this approach to concepts, the original theory T(F) can be decomposed into the analytic Carnap sentence and the synthetic 'Ramsey sentence' of the

theory—(EΦ)(T(Φ)). The Ramsey sentence expressed the substantial commitments of the theory—there is an entity which . . .—while the Carnap sentence expresses the definitional commitment to dubbing that entity F. The original theory framed using the concept F is thus equivalent to the conjunction of the Ramsey and Carnap sentences.

This understanding of the relation between theories and concepts informs an influential contemporary vision of philosophical practice, inspired originally by the work of David Lewis and more recently codified by Frank Jackson (1998). As conceived by Jackson, philosophy proceeds in two stages. The first stage involves the identification and articulation of folk concepts. Here the aim is to figure out how everyday thought conceives of free will, mental states, persons, moral value, and other important philosophical categories. At this stage we will use traditional methods of conceptual analysis and reflection on possible cases. Then, once we have analysed such everyday concepts, we can turn to our most serious theories of the world to investigate what satisfies them. This second stage will involve synthetic claims about the underlying nature of reality—we will look to physics and any other basic sciences to inform us about possible candidates which might realize our everyday concepts. But while this second stage appeals to synthetic knowledge, it depends essentially on the first analytic stage, where the identification of everyday concepts plays an essential role in setting the agenda for further philosophical investigation.

Thus Jackson:

‘What then are the interesting philosophical questions that we are seeking to address when we debate the existence of free action and its compatibility with determinism, or about eliminativism concerning intentional psychology? What we are seeking to address is whether free action according to our ordinary conception, or something suitably close to our ordinary conception, exists and is compatible with determinism, and whether intentional states according to our ordinary conception, or something suitably close to it, will survive what cognitive science reveals about the operations of our brains’ (Jackson 1998, p 31, his italics).

One worry about this programme is whether the relevant concepts really have the requisite descriptive structure. Strong externalists about content will doubt that there are any analytic assumptions involving free will say, or person, that you must be committed to if you have these concepts, let alone assumptions that will uniquely identify the referents of these concepts. (Cf Williamson 2007, ch 4.)

Another worry, which arises even if we reject strong externalism, relates to the familiar question of which everyday assumptions play a definitional role. As before, are all assumptions to be included, or only some distinguished subset—and if the latter, what marks the distinction?

I think that these are serious worries, but I shall not press them here. This is because I think I can show that, even if there are analytic truths of just the kind that that Jackson supposes, they are of no significance to philosophy.

Jackson says that everyday concepts set the agenda for further metaphysical investigation. It is because everyday thought conceives of free action, and intentional states, and so on, in such-and-such ways that we philosophers are prompted to probe the nature of those things that fit these specifications.

But why think of the matter in this way? Doesn't it make far more sense to suppose that it is the synthetic theories implicit in everyday thought that raise the initial philosophical questions, not the mere analytic commitment to concepts. Even after we allow that everyday thought is indeed structured as Jackson supposes, the natural assumption is surely that it is the synthetic Ramsey sentences that matter to philosophy, not the analytic Carnap sentences. What makes philosophers interested in investigating further is the pre-theoretical supposition that there are entities fitting such-and-such specifications, not just the hypothetical specification that if there were such entities, then they would count as free actions, or intentional states, or whatever.

The point is most easily brought out by considering cases where current everyday thought endorses the definitional Carnap sentence involving some concept, but not the substantial Ramsey sentence. I think, and so do all of you, that if there is a category of women who ride on broomsticks, cast spells, and enter into pacts with the devil, then these women are witches. But of course none of us think that there is a real kind of this sort, and so have no inclination at all to conduct metaphysical investigations into its nature. Again, to take a somewhat more serious example, we can all agree, I take it, that if there are entities that are conscious, separable from bodies, and can survive death, then those things are souls. But only those few among us who think that there actually are souls will have any motive to probe their metaphysical nature further.

The point is that concepts themselves are ontologically non-committal. The mere possession of concepts carries no implications at all about the contents of reality, and so cannot point the way to further investigations, in the way that substantial synthetic claims can.

I am very much in favour of the idea that much philosophy involves subjecting everyday ideas to serious scrutiny. All of us, philosophers included, acquire much of our understanding of the world from the everyday culture in which we grow up. Some of this everyday lore is sound, and some is not. If we are serious about our understanding of the world, we need to examine the assumptions that we acquire from everyday thinking, and see how many of them stand up to serious examination. But none of this is anything to do with concepts. Since concepts on their own are non-committal about reality, they cannot lead us astray. But the synthetic commitments of everyday thought can, and so do need to be properly examined.

When Jackson and others who subscribe to his programme actually address serious metaphysical issues, they of course proceed in just the way I am advocating. That is, they take cases where everyday thinking commits us to substantial assumptions about the contents of reality, and ask whether these assumptions are sustainable. To this extent, I

would say that their official account of what they are doing is belied by their actual practice. Officially they say they start with concepts, but in fact they start with theories.

Moreover, even the official account of what they are doing is not always strictly maintained. The difference between concepts and theories is not always respected. So in a number of passages. Jackson talks about the initial exploration of folk ideas as a matter of identifying theories rather than concepts.

For example:

'. . . my intuitions reveal my theory of free action . . , your intuitions reveal your theory . . . to the extent that our intuitions coincide with those of the folk, they reveal the folk theory.' (Jackson 1998, p 32)

And later we find him saying that

'My intuitions about which possible cases to describe as cases of K-hood . . . reveal my theory of K-hood' (op cit, p 37).

As I have said, I am all in favour of beginning philosophical investigation with everyday theories. But this is not the same as beginning with mere concepts. Theories involve significantly more than concepts, as is shown by the cases of witches and souls, where we have the concepts but not the corresponding theories.

The Method of Possible Cases

My thesis that philosophy deals in synthetic claims might seem to be inconsistent with one salient feature of philosophical practice. Philosophers characteristically test philosophical claims by considering whether counterexamples are in some sense imaginable. At first pass, this certainly seems to support the view that philosophical claims are conceptual in nature. Imagination can plausibly show us whether or not

certain situations are conceptually possible, but presumably not whether they are actual. Correspondingly, it looks as if imagination can usefully test claims about what is conceptually required, but not about what actually occurs.

For example consider Gettier's demonstration that knowledge is not true justified belief. Gettier showed us how to construct possible cases in which people have true justified beliefs, but are not knowers (because, roughly speaking, the truth of their belief is accidental relative to their method of justification). Surely this shows that the philosophical claim being tested is that true justified belief conceptually requires knowledge. Otherwise how could the mere conceivability of counter-examples disprove it?

Again, consider Kripke's demolition of the descriptive theory of ordinary proper names. Kripke invited us to consider possible cases in which someone (Schmidt, say) satisfies all the descriptions associated with some name ('Gödel') yet is not the bearer of that name (because he is not the causal origin of its use). Here too it looks as if the mere conceivability of a counterexample is enough to discredit the thesis of interest, and thus that this thesis must be conceptual in nature.

One possible naturalist response would be to reject the method of reasoning by possible cases. Since philosophy is concerned with synthetic claims, just like the sciences, it can't possibly make progress just by reflecting on what is conceptually possible. Instead philosophers should get out of their armchairs and engage directly with experimental and observational findings.

This is not my view. I take it to be uncontentious that Gettier's and Kripke's thought-experiments led to genuine advances in philosophical knowledge. More generally, I regard reflection on possible cases as a highly fruitful mode of philosophical investigation. As I said at the beginning, I am not proposing any revisionary account of philosophical practice. From my point of view, the methods that philosophers use are just fine, including the method of reflection on possible cases. So instead of rejecting

armchair reflection, I am going to argue that armchair methods provide more than purely conceptual information and so can play a part in the assessment of synthetic claims.

The obvious comparison here is with thought-experiments in sciences. Many important advances in science have been prompted by pure reflection on possible cases. Famous examples include Archimedes on buoyancy, Galileo on falling bodies and the relativity of motion, Newton's bucket experiment, Maxwell's demon, and Einstein on quantum non-locality. Cases like these certainly suggest that armchair reflection can be relevant to establishing synthetic claims.

Scientific thought-experiments display a range of different structures. Let me focus on one of the simpler cases—Galileo's analysis of falling bodies. According to the Aristotelian orthodoxy of Galileo's time, heavier bodies fall faster than lighter ones. Galileo asks his readers to consider what will happen if a lighter body is tied to a heavier one by a piece of string (Galileo 1638). Since the Aristotelian theory says the lighter body will be inclined to fall more slowly than the heavier, it follows that the lighter should slow the heavier down when joined to it. But by the same coin the compound body consisting of the two tied together is heavier than the two individual bodies, and so should fall faster than both. The Aristotelian theory is thus shown to be inconsistent. Moreover, it looks as if the only consistent account will have the compound body falling at the same speed as the individual components, which implies that speed of fall is independent of weight.

In this kind of case it is clear that the relationship between weight and speed of fall is a synthetic matter. Concepts cannot guarantee anything this relationship. How then can armchair reflection show us what to think? The answer must be that armchair reflection is showing us more than that certain scenarios are <u>conceptually</u> possible. Of course, it can't show that there are any <u>actual</u> cases in which a compound body falls at the same speed as its components. Galileo didn't create a real case of two bodies tied together just by thinking about it. Still, Galileo didn't need an actual case to disprove the Aristotelian theory. If we construe that theory as saying that the faster fall of heavier bodies is

required by the laws of nature, it will be enough for Galileo to show that a case of a heavier body falling at the same speed as a lighter one is consistent with the laws of nature. And that is just what Galileo does. He asks us to consider a manifestly naturally possible scenario in which two bodies are tied together, and then judges that in such a case the laws of nature will lead the compound body to fall at the same speed as its components.

Obviously, the crucial step here is played by Galileo's intuition that a compound body will fall at the same speed as its components. And this is clearly a synthetic intuition, by no means guaranteed by the concepts it involves. That is why it can overturn the synthetic Aristotelian theory.

I want to suggest that philosophical thought experiments have the same structure. Explicit philosophical theories about the requirements for a thinker to know something, or for a thing to bear a name, (or for someone to have acted freely, or for one person to be the same as another, . . .) are synthetic claims about the relevant categories. Philosophers then test such synthetic proposals against their intuitions about possible scenarios. Thus Gettier appealed to the intuition that a belief whose truth is accidental relative to its method of justification is not knowledge; Kripke appealed to the intuition that something that is not the causal origin of a name is not its bearer; and so on. On my account, all these intuitions are synthetic claims about the relevant kind of scenario. This is why they have the power to discredit the initial philosophical theories.

From this perspective, there is nothing in the method of reasoning about possible cases to undermine the idea that philosophy is concerned with synthetic claims. It is simply a technique that enables us to counter the synthetic theories proposed by philosophers by the synthetic intuitions elicited by thought experiments.

There is one respect in which this account of thought experiments may be an oversimplification. I have suggested that thought experimental intuitions manifest certain general principles, such as that an accidentally true believer isn't a knower, or that the

causal origin of a name is its bearer, and so on. However Tim Williamson has pointed out that such general claims are arguably more than the thought experiments committed us to (2007, Ch 6). For example, in order to disprove the tripartite analysis of knowledge, Gettier only needed the particular counterfactual claim that, in the most obvious understanding of his scenario, the relevant thinker would not be a knower. There is no need to suppose that any thinker satisfying the explicit specifications of his scenario would fail to know, still less to suppose some still more general principle as that 'all accidentally true believer aren't knowers'. For Williamson, philosophical thought experiments thus appeal only to our ability to reason counterfactually, and do not demand any grasp of general principles.

I am happy to agree that counterfactual reasoning is enough for thought-experimental purposes, and correspondingly that is in by no means mandatory to suppose that general principles lie behind the relevant intuitions. Even so, I would like to continue working on the assumption that thought experiments display general principles. This may be an oversimplification, but I don't think it is too far from the truth. We may not fully understand counterfactual reasoning, but it is clear that it is strongly constrained by general claims about the working of the world. Williamson alludes to the role of imagination in counterfactual reasoning (2007, chs 5-6). But when I think about what would happen if I had dropped a vase, say, I do not imagine every outcome that is permitted by the concepts involved, such as that the vase floats gently onto the table. Rather I consider just those outcomes that are consistent with some such synthetic general claim as that heavy bodies fall rapidly when unsupported. Perhaps this general claim as just formulated is more precise than anything that governs our counterfactual thinking. Still, it seems clear that our counterfactual thinking must be informed by some such principle. In line with this, I shall continue to assume that the intuitions in philosophical thought experiments are informed by general principles. Attempts to state these principles explicitly may inevitably lead to oversimplification, but I propose to overlook this in the interests of facilitating investigation into their nature. (In what follows I shall use 'intuition' to refer to both the general principles informing our

counterfactual reasoning and the specific judgements about counterfactual situations that issue from them. When the distinction matters I shall draw it explicitly.)

The Encapsulation of Assumptions

There is an obvious objection to my proposed analogy between philosophical and scientific thought experiments. Consider Galileo's thought experiment again. The crucial intuition was that tying two bodies together won't make any difference to their speed of fall. Now, it is clear that this conjecture is hostage to further empirical investigation. It may strike us as obvious that Galileo is right, but even so empirical observation remains the ultimate test of his intuition. Galileo is in effect hazarding a guess—albeit a highly informed guess—as to the synthetic facts, and the final arbiter of this guess must be real observations. We need to find some actual bodies that are tied together see how they fall. Either they will conform to Galileo's intuition, or they won't. And both options are clearly left open by the terms in which the issue is posed.

Things seem rather different in philosophy. In the Gettier thought experiment, for example, the analogous intuition was that a belief isn't knowledge if its truth is an accident relative to its method of justification. But there seems no analogous room to check this intuition against real cases, by seeing whether or not actual thinkers with such accidentally true beliefs are knowers. For we already know what we will say about any such cases—namely, that these thinkers are certainly not knowers. The reflection involved in the philosophical thought experiment is itself enough to tell us what we will judge in any similar real situation, and thus to rule out any possibility of observing someone who is an accidentally true believer yet a knower. The same seems true of philosophical thought experiments in general. Take the Kripke case. We don't need to find any real cases of names whose original bearers don't fit the associated descriptions, in order to check whether or not the names really do name the original bearers. For again, we already know what we will say about any real such cases—namely, that the names apply to the original bearers even if they don't satisfy the descriptions. And this

again rules out any possibility of observing a name which turns out to refer to the satisfier of associated descriptions rather than the original bearer.

In short, the intuitions in play in the philosophical thought experiments don't seem to be falsifiable in the way they ought to be if they were synthetic claims. On the contrary, their inviolability to any observational falsification seems to argue strongly that they are analytic. And this would then imply that the philosophical thought experiments are serving to manifest the structure of our concepts, rather than to draw out our implicit empirical opinions.

However, this is not the only way of seeing the matter. An alternative is to hold that the relevant philosophical intuitions are synthetic, but encapsulated in the cognitive systems that make judgements about such categories as knowledge, names, free will, persons and so on. By way of analogy, consider the way that the human visual system detects the edges of physical objects by registering sharp changes in intensity in the visual field. We can think of the visual system as embodying the implicit 'assumption' that intensity changes are due to the edges of physical objects. This assumption is then 'encapsulated' in the sense that the visual system will continue to embrace it even in cases where we are personally aware that the intensity changes are due to something else, as when we are viewing the surface of a photograph.

Because of this, it is inevitable that we see intensity changes as edges. And this means that our visual system is never going to deliver intuitive particular 'judgements' that falsify the intuitive general 'assumption' that all intensity changes are due to object edges. There is no possibility of a visual observation of sharp intensity changes that are not seen as edges. Still, it is clear enough that the assumption that all intensity changes are due to object edges is akin to a synthetic rather than an analytic claim. Its approximate truth is not due to the structure of its content, but to the fact that most intensity changes in the actual world are due to the edges of physical objects.

I would like to say the same about the general intuitions which guide us in making particular judgements about knowledge, names, persons, free will and so on. The sub-personal cognitive mechanisms responsible for such judgements are not well understood, as is evidenced by the difficulty philosophers have in identifying the principles on which they operate. But it is clear enough how they must work: they take in information which do not presuppose the relevant categories, and use it to arrive at judgements about who knows what, and which words name which things, and when someone is the same person as someone else, and so on. I want to suggest that the particular intuitions displayed in philosophical thought experiments manifest the implicit general 'assumptions' on which such mechanisms hinge, in the way that the visual system hinges on the 'assumption' that intensity changes are due to object edges.

This is why there is no question of any direct judgements about particular cases falsifying such 'assumptions'. If my judgmental procedures decide who is a knower by assuming inter alia that accidentally true believers are not knowers, then clearly there isn't any question of my meeting up with a case where I judge such an accidentally true believer to be a knower after all. Again, if my judgemental procedures decide what things bears some name by noting the causal origin of the use of the name, them I'm not going to come across cases where I judge that some name is borne by something other than its causal origin. But this impossibility of direct falsification does not mean that the relevant general assumptions are analytic. They may yet have a substantial synthetic content, like the visual system's assumption that intensity changes are due to object edges.

Some readers may be feeling that I have not yet established a positive case for my first thesis that philosophy deals in synthetic matters. In this section and the last I may have succeeded in showing how the importance of thought-experiments can be made consistent with that thesis. But isn't it equally consistent with the contrary thesis that philosophy is centrally concerned with analytic matters? I may have been able to concoct a story which makes philosophical thought-experiments come out like scientific ones. But isn't the more natural account still that the point of these thought-experiments is to

articulate the structure of our concepts? Don't I owe some positive arguments against this natural account and in favour of the one I have contrived?

I have two responses to this line of thought. First, there are independent reasons for thinking that substantial synthetic assumptions are built into the automatic mechanisms that allow us to make particular judgements about philosophically salient categories like knowledge, names, persons, free will and so on. Judgements like these are important to us in our daily life, and it is therefore unsurprising that we should have unthinking mechanisms that allow us to form them quickly and efficiently. But it would be odd then to suppose that any inferential assumptions built into these mechanisms must be analyticities whose truth is guaranteed by the structure of their contents. The whole point of these mechanisms is to start with limited information and deliver further conclusions. It would run quite counter to this function if they were restricted to analytic inferences and precluded from engaging in ampliative ones.

My second response is that, if philosophical thought-experiments were concerned only with the elaboration of analyticities, they would be much less interesting than they are. They would tell us about the structure of our concepts, but they wouldn't help us to understand the rest of the world. (Cf Williamson 2007, 204-7.)

Recall that analytic knowledge comes in the form of conditional Carnap sentences. These simply explain that, <u>if</u> things satisfying certain requirements exist, <u>then</u> they count as such-and-suches, but analytic knowledge never deliver any categorical information about the contents of actuality. Correspondingly, the philosophical analysis of concepts may tell us that, <u>if</u> there is a propositional attitude that requires truth, justification, and so on, <u>then</u> it is knowledge—or again that, <u>if</u> words and things bear certain causal relations, <u>then</u> the words name the things.

But this seems far less than we actually get from the relevant thought-experiments. Thus I take Gettier to have shown not just that our <u>concept</u> of knowledge imposes a requirement of non-accidentally, but far more interestingly that this requirement is

satisfied by real knowledge—that is, the state that plays an important role in the world and is displayed in many paradigm cases. Similarly, I take Kripke to have shown that not just that we conceptualize names causally, but in addition that real name-bearer pairs—all those many instances we are familiar with—are causally related.

If the assumptions manifested in philosophical thought experiments really are synthetic, then of course their contents leave it open that they may turn out to be false. They may not be directly falsifiable via a simple contrary observation, for the reasons outlined above. Still, we can imagine how more sophisticated investigations may show them to be flawed. Compare the way in which, even though we never see sharp intensity changes as anything other than object edges, more elaborate investigation can show us that there are plenty of contrary cases. We can imagine reaching a similar conclusion about knowledge, say. We are notoriously unclear about the significance of knowledge. (Is it just that knowledge is an effective means to true belief, or because knowledge underpins a certain kind of robustness in action explanations, or because it is biologically more basic than true belief, or for some further reason?) Perhaps when we have a good answer to this question we will conclude that the principle that knowledge requires non-accidentality is a crude rule of thumb that works well enough in many cases, but on occasion leads us astray. (I doubt that it will actually turn out like this. My concern here is only to show how it is epistemologically possible.)

Philosophy is A Posteriori

Let me now turn to my second thesis that philosophical knowledge is a posteriori not a priori. It might seem that this will now follow quickly, given my first thesis that philosophical claims are synthetic. How can a synthetic claim possibly be known to be true independently of experience, given that its content alone leaves it open that it might be false? But of course this is too quick. Traditional theists and transcendental idealists both take themselves to have good answers to this question. And even if we reject these particular answers, there is room for other non-experiential accounts of synthetic knowledge.

On this topic, Timothy Williamson has argued that philosophical intuitions, though synthetic, should not be counted as a posteriori (2006, 165-9, 189-90). His reason is that experience does not play a normal evidential role in generating them. We can't point to past observations of supporting instances to support such claims as that accidentally true believers are not knowers, or that names refer to their causal origins. Our route to these claims is thus clearly unlike the normal justification of synthetic generalizations by inductive or abductive evidence. (Williamson does not conclude that such philosophical judgements are a priori—he thinks the traditional contrast is not useful here. But we need not pursue this point, given that I am about to argue contra Williamson that philosophical intuitions should definitely be counted as a posteriori.)

I am in broad agreement with Williamson's view of the provenance of philosophical intuitions. They are not products of normal inductions or abductions. This accords well with my suggestion that the underlying assumptions are 'encapsulated' in the cognitive mechanisms we use to decide on cases of knowledge, naming, and so on. Of course, there is plenty of room for debate about the means by which assumptions become encapsulated in this way. Strong nativists argue that all the relevant information is encoded in genes bequeathed to us by natural selection. Others hold that the relevant assumptions are laid down early in individual development, via the influence of surrounding culture and possibly also of acquaintance with particular paradigm cases. Still, whatever the precise truth on this matter, Williamson seems right to observe that the source of philosophical intuitions is not normal inductive or abductive evidence.

Still, the source of philosophical intuitions is one thing, their justification another. Even if philosophical intuitions do not derive from experience, it may still be that they can only be justified a posteriori. This is my view. Note that in general the epistemological status of encapsulated assumptions is not high. The function of cognitive mechanisms that embody encapsulated assumptions is to deliver judgements about particular cases quickly and efficiently. Because of this, the relevant assumptions are standardly rules of thumb that work well enough in most cases but are not strictly accurate, in the way illustrated by

the familiar perceptual examples. If the cognitive mechanisms behind philosophical intuitions are at all similar, we should expect encapsulated philosophical assumptions to have a similar status. They may work well enough for practical purposes, but they may not be strictly accurate and may lead us astray in certain cases. If we are to be confident about these assumptions, we will need to make them explicit and subject them to proper a posteriori evaluation.

I have already argued that it is at least epistemologically possible that there may be inaccuracies in the assumption that knowledge must be non-accidental. Nor is it hard to think of real cases of mistaken assumptions which were once encapsulated and therefore seemed immune to imaginable counter-examples. Descartes thought it unimaginable that a purely mechanical being could reason. Kant thought it unimaginable that parallel lines could meet. Many people still find it unimaginable that temporal succession could be relative, or that time could have a beginning.

The recent findings of 'experimental philosophy' are relevant here. They indicate that many central philosophical intuitions, including those invoked by Gettier and Kripke, are by no means universal, but rather peculiar to certain cultures and social classes. (Knobe and Nichols eds 2008.) At one level, it is not always clear what to make of these findings. Presented as a challenge to 'conceptual analysis', they invite the response that the variability of intuitions only establishes the philosophically insignificant point that different groups of people express different concepts by words like 'knowledge' and 'name'. However, the variability of intuitions is clearly more significant if philosophical intuitions are substantial claims whose truth is not analytically guaranteed. In that case, the variability of the intuitions is in tension with their reliability. If different people have opposed philosophical intuitions, then it cannot be that intuitions of this kind are always true. This reinforces the point that an a priori provenance for philosophical assumptions does not amount to an a priori justification. As before, the justification of such assumptions requires that we subject them to proper a posteriori examination.

It may seem as if I am here backtracking on my earlier enthusiasm for armchair philosophy. Among philosophers who agree with me that philosophical intuitions are synthetic, we can distinguish two broad positions. There are those who think that philosophical intuitions are little more than manifestations of naïve folklore, and should therefore carry little weight in serious philosophical discussion. According to this point of view, philosophers should turn away from intuitions and instead engage with serious empirical theories. (Cf. Kornblith 2002, Knobe and Nichols eds 2008.) On the other side are philosophers like Timothy Williamson, and perhaps Alvin Goldman, who think that philosophical intuitions are by and large reliable, and that the findings of the experimental philosophers are not as worrying as they appear. (For example, Williamson suggests that the special training of philosophers may make them sensitive to niceties that escape the untrained—2007, 191.) My line of argument in this section so far may seem to place me on the former side and thus against armchair investigation.

However, I think this conclusion is based on a false dichotomy. Just because I am doubtful about the authority of philosophical intuitions, it doesn't mean that I have to reject the method of reasoning about merely possible cases. Armchair thinking can be useful, even if the intuitions involved are unreliable. Go back to the idea, briefly aired earlier, that philosophy is characteristically concerned with theoretical tangles. We find our thinking pulled in opposing directions and cannot see how to resolve the tension. Often part of our predicament is that we don't know what assumptions are directing our thinking. We end up with conflicting judgements, but are unclear about what led us there. In such cases thought experiments can bring the implicit principles behind our conflicting judgements to the surface. They make it clear what intuitive general assumptions are governing our thinking and so allow us to subject these assumptions to explicit examination. Nothing in this requires that thought-experimental thinking is generally reliable. When some explicit prior theory conflicts with an intuitive judgement elicited by a thought experiment, this needn't always result in the rejection of the theory. We can also end up rejecting the implicit assumptions behind the thought-experimental intuition.

Just this pattern is displayed by some of the most famous and important thought experiments in science. Consider the 'tower argument' against Copernicanism, which appeals to the intuition that an object dropped from a moving source will be 'left behind' as it falls. Or take the Einsteinian argument against the completeness of the Copenhagen interpretation of quantum mechanics, which appeals to the intuition that spacelike separated events cannot be co-ordinated without a common cause. In cases like these, the assumptions generating the thought experiments eventually came to be recognized as mistaken. But this certainly did not mean that the thought experiments were worthless. Both the tower and Einstein arguments were hugely important in the history of science. But showing us which of our implicit assumptions conflicted with new theoretical ideas, they led to crucial new advances. Galileo responded to the tower argument with his innovatory formulation of a principle of inertia, and J.S. Bell to the Einstein argument with his derivation of the eponymous inequality whose experimental confirmation ruled out local hidden variable theories.

It is not hard to think of similar philosophical cases. The worth of philosophical thought experiments does not always require that the intuitions they elicit are sound. In some cases, of course, the intuitions will be correct. I don't think we should really harbour any serious doubts about the Gettier and Kripke intuitions. But in other cases thought experiments can clarify the issues even if the accompanying intuitions point us in the wrong direction.

Consider the classic Lockean set-up where someone's memories are transferred to a new body. We all have an intuition that the person goes with the memories, not the old body, as evidenced by our reactions to the many fictions which trade on just this kind of scenario. But few philosophers of personal identity would nowadays hold that this intuition is decisive in favour of Lockeanism. We need to follow through the implications of the Lockean views and assess the overall resulting theory against its competitors, and in this context the initial intuition is indecisive. But, for all that, it would be hard to deny that Locke's thought experiment has led to advances in our understanding of personal identity. Again, consider the intuition that conscious

properties are ontologically distinct from physical ones, as displayed in our immediate reaction to zombie scenarios. Here too few would suppose that these intuitions are decisive in refuting physicalism. But at the same time even physicalists will allow that reflection on zombie cases has helped to clarify what is at issue in the mind-brain debate. (I shall return to this particular example in the next section.)

So my view is that philosophical intuitions do no qualify as knowledge until they have been subject to serious a posteriori assessment. Philosophers need to articulate their intuitions in order to understand the source of their theoretical difficulties. But since these intuitions are standardly nothing more the encapsulated rules of thumb we happen to have grown up with, we should not place any great epistemological weight on them until they have been properly evaluated against experience. In saying this, I do not mean to imply that all philosophical claims need to be assessed directly against specific empirical findings from empirical disciplines. A synthetic theory can be vindicated a posteriori even though it has no specific empirical evidence to call its own, on the grounds that it provides a more coherent and natural overall account than the alternatives.

As it happens, I do think that specific empirical findings bear directly on a surprisingly wide range of philosophical issues. These include not just topics from philosophy of science, such as the logic of natural selection or the interpretation of quantum mechanics, but also such central and traditional topics as the nature of causation and the relation between mind and brain. Still, I am happy to allow that there are other central philosophical issues, such as the nature of persisting objects or realism about properties, where the philosophical claims float free of any specific matters investigated by the empirical sciences.[1] In such cases, we will then have no alternative but to evaluate alternative philosophical positions by comparing their overall coherence and naturalness. Still, this too is an a posteriori procedure, akin to the method by which we compare alternative scientific theories that are underdetermined by the evidence. When we prefer Copernicus to Ptolemy, or special relativity with the Lorentz-Fitzgerald reworking of classical mechanics, it is not because of any specific empirical findings, but because they

[1] But see Maudlin, 2007, who brings scientific considerations to bear even on these two topics.

are more in accord with our general a posteriori understanding of the way the world works. I see no reason to doubt that the most abstract philosophical issues are to be decided in the same way.[2]

Philosophy and Necessity

Can the account of philosophy offered so far can accommodate the modal dimension of philosophical knowledge? It is sometimes said that the difference between philosophy and science is that philosophy seeks necessary truths where science trades in contingencies. (Thus Russell: '[a philosophical proposition] must not deal specially with things on the surface of the earth, or with the solar system, or with any other portion of space and time. . . . A philosophical proposition must be applicable to everything that exists or may exist' 1914, p 110.) This modal view of philosophy might seem to be in tension with my account of philosophy as synthetic and a posteriori. Don't we need a priori analysis to uncover necessary truths?

But of course this line of thought is far too quick. There no reason why necessities should not be synthetic and a posteriori. Empirical science provides plenty of familiar examples. Water is H_20. Heat is molecular motion. Stars are made of hot gas. Halley's comet is made of rock and ice. All these claims are necessary, but clearly they are not knowable a priori on some analytic basis.

These claims are necessary because they use rigid terminology to report on facts of identity or constitution. All claims of these kinds are necessary, notwithstanding any synthetic a posteriori status they may have. It is a nice question, worthy of further discussion, why facts like these should count as necessary, while truths about spatiotemporal location, say, do not. But this is not the place to pursue this issue. For

[2] Some readers might wish to query whether choices between underdetermined scientific theories should count as a posteriori. I think that they should (1993, Ch 5), but perhaps I can let the point pass here, and settle instead for the observation that empirically underdetermined philosophical theory-choices are made on the same grounds as scientific ones.

present purposes the important point is simply that the necessity of claims of these kinds is perfectly consistent with their synthetic a posteriori status.

The central questions of philosophy are almost entirely concerned with issues of identity and constitution. When we ask about knowledge, names, persons, persisting objects, free will, causation, and so on, we are seeking to understand the nature of these categories. We want to know whether knowledge is the same as true justified belief, whether naming involves descriptive content, whether persisting objects are composed of temporal parts, and so on. Any truths we might establish about such matters will inevitably be necessary rather than contingent, even if they are also a posteriori and synthetic.

The answers to the central questions of philosophy may be necessary, but that is no reason to suppose that philosophy is here concerned with necessity per se rather than actuality. Consider empirical science once more. As I have just observed, many of the claims established by science are necessary. But it would be odd to infer from this that empirical science is aiming to explore some wider modal realm rather than simply to understand the actual world. When science investigates the chemical make-up of water, or the composition of the stars, it is primarily concerned with how things are in this world ('with things on the surface of the earth, or with the solar system, or with any other portion of space and time. . .') That these discoveries have implications about the contents of other possible worlds, so to speak, is an inevitable side-effect of the content of these claims, but not something that we need regard science as actively seeking.

I say the same about the central areas of philosophy. Our primary philosophical concern is to find out about things in this world. We want to know about such actual categories as knowledge, free will, persons, and so on—kinds that exist and make a difference in this world. Of course, given that answers to our questions will normally take the form of claims about identity and constitution, philosophical knowledge will also place constraints on what is necessary and possible. But there is no reason to regard such modal corollaries as our main aim. We are first seeking to understand this world, and are only derivatively concerned with modal matters. We want to know whether p, not

whether necessarily p. That the former implies the latter does not make the latter our focus of interest, any more that my interest in whether you are 47 years old makes me interested in whether your age is a prime number.

Of course, some philosophers are specifically interested in modal questions as such. They are interested in whether necessary truths are necessarily necessary, or in whether modal claims commit us to an ontology of possible worlds, or in the connection between metaphysical and conceptual necessity, or indeed in why facts of identity and constitution but not spatio-temporal location should count as necessary, and so forth. There are certainly substantial philosophical issues worthy of serious discussion.[3] But most central philosophical questions are not of this form. The study of modality is a specialist subject within philosophy, engendered by specific theoretical interests. There is no reason to suppose that an interest in modality infects all of philosophy, even if all philosophical claims have modal implications.

Having said this, it is worth recognizing that it is often heuristically useful to focus on modal implications, even in cases where our real interest is in non-modal matters. Given the immediate modal upshot of claims of identity and constitution, it is sometimes easier to articulate our thinking by starting with the modal consequences rather than their this-worldly counterparts. Take the relation between individual objects and their property instantiations. In the actual world there is a one-to-one correspondence between objects and sets of property instantiations. But is this a matter of identity, as in the 'bundle theory' of objects, or mere association? A good way to clarify our thinking on this issue is to consider the modal question of whether there could be a world in which this blue

[3] This branch of philosophy obviously demands a qualification to the third of my initial theses—philosophers of modality are certainly concerned to understand modality per se, even if other philosophers are not. But it may still satisfy my other two theses by being synthetic and a posteriori. As before, however, we cannot expect decide these issues in the absence of an agreed view about the nature of modal claims.

cup, say, acquired all the properties of that red one, and vice versa. To the extent this strikes us as possible, then we are thinking of objects as distinct from their property instantiations; while if it seems that this is not a real possibility, then we are identifying objects with their property instantiations.

I am not of course here suggesting that such modal intuitions are somehow a privileged route to the truth. Whether we are right to think of objects as bundles of properties, say, would remain a substantial further issue, even after modal reflection has made it clear that this is our intuitive view. The role of the modal reflection is merely to clarify the content of our intuitive commitments in cases where thinking about actuality alone leaves them unclear, not to decide the substantial issues. From this perspective, modal thinking is a special case of the kind of thought-experimental reflection described in previous sections. It is a useful way of identifying the implicit assumptions that drive our reasoning. Once these assumptions have been identified, we are then is a position to subject them to serious a posteriori evaluation.

Let me conclude with one further example. Consider the relation between conscious mental properties and brain properties. Let us agree that pairs of these properties go hand-in-hand in the actual world. Still, is this association due to the identity of the relevant properties, or merely to a correlation between distinct properties? Well, ask yourself whether there could possibly be a being with all your brain properties but who lacks your conscious properties. If you think that such zombies are possible, then you must be of the view that conscious properties are distinct from brain properties in this world. Conversely, if you think that conscious properties are in actuality one and the same as brain properties, then you won't think that zombies are so much as possible.

Many recent writers look at this thought experiment differently. They think we can start with our concepts of conscious and brain states, proceed to the point that zombies are conceivable, somehow move from this to their possibility, and thence end up with the conclusion that conscious and brain properties are distinct in the actual world. (Chalmers 1996, Bealer 2002.) I don't think that this works at all. (Papineau 2007.) The

interesting thing about zombies isn't that we can conceive them—after all, we can conceive lots of things that aren't possible—but that they strike us as possible. This shows us something rather surprising, namely that at an intuitive level we are all dualists about the mind-brain relation.

Of course, it is one thing to identify this intuition, and another to justify it. As I have argued throughout, philosophical intuitions need a posteriori backing before we can place confidence in them. In this case it seems clear that the a posteriori evidence counts against the intuition. (Papineau 2002, Appendix.) Still, this is not the place to pursue this issue, which is in any case independent of my present point—which is that in most familiar cases the purpose of modal reflection is not to find out about other possible worlds per se, but simply to clarify our pre-theoretical assumptions about the actual world.[4]

References

Bealer, G. 2002 'Modal Epistemology and the Rationalist Renaissance' in Gendler, T. and Hawthorne, J. Conceivability and Possibility Oxford: Oxford University Press

Brandom, R. 1994. Making it Explicit Cambridge, Mass: Harvard University Press

Brandom, R. 2001 'Reason, Expression, and the Philosophical Enterprise' in Ragland, C. and Heidt, S. (eds.) What Is Philosophy? New Haven, Conn: Yale University Press. 74-95

Chalmers, D. 1996 The Conscious Mind Oxford: Oxford University Press

[4] I have presented versions of this material in many places, including a seminar on the philosophy of philosophy in London in 2008, and I would like to thank all those who responded on these occasions. I can particularly remember helpful comments from George Bealer, David Chalmers, Keith Hossack, Fraser McBride, Tom Pink, Andrea Sangiovanni, Gabriel Segal, Jonathan Shaffer, Barry Smith, Stephen Stich, Scott Sturgeon, Celia Teixeira, Mark Textor, Lee Walters, Tim Williamson and Crispin Wright.

Galileo, G. 1638 Discourses Concerning Two New Sciences translated by Drake, S. 1974 Madison: University of Wisconsin Press

Goldman, A. 2007 'Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status' Grazer Philosophische Studien 74 1–26

Jackson, F. 1998 From Metaphysics to Ethics: A Defence of Conceptual Analysis Oxford: Oxford University Press

Joyce, R. 2006 'Metaethics and the Empirical Sciences' Philosophical Explorations 9 133-148

Knobe, J. and Nichols, S. eds 2008 Experimental Philosophy New York: Oxford University Press

Kornblith, H. 2002 Knowledge and its Place in Nature Oxford: Oxford University Press

Maudlin, T. 2007 The Metaphysics within Physics New York: Oxford University Press

Papineau, D. 1993 Philosophical Naturalism Oxford: Blackwell

Papineau, D. 2002 Thinking about Consciousness Oxford: Oxford University Press

Papineau, D. 2007 'Kripke's Argument is Ad Hominem Not Two-Dimensional' in Philosophical Perspectives 21 475-494

Russell, B. 1914 'On Scientific Method in Philosophy' reprinted in his Mysticism and Logic 1917 London: Longmans

Williamson, T. 2007 The Philosophy of Philosophy Oxford: Oxford University Press