

The Possibility of Knowledge

1.1 How-Possible Questions

I'm listening to a baseball game on the radio on the way home from my office. To my surprise, I hear the announcer describing a fielder catching a long fly ball that was just about to hit high up on the fence. I'm puzzled because I happen to know that the fence is twenty feet high. I wonder how, given the height of the fence, it was possible for the fielder to make the catch. I only stop being puzzled when I discover that he used a ladder attached to the scorekeeper's platform to reach the ball.¹ In the meantime, what I want to know is not whether it was possible for him to make a catch twenty feet off the ground (since he actually made it) but how this was possible. This is an example of a how-possible question. To ask a how-possible question is to ask how something which looks impossible given other things that one knows or believes is nevertheless possible.²

How-possible questions matter in philosophy because, as Nozick points out, 'many philosophical problems are ones of understanding how something is or can be possible' (1981: 8). Familiar philosophical how-possible questions include 'how is freedom possible?' and 'how is evil possible?'.³ The first of these questions arises because of the tension between the natural assumption that we are capable of acting freely and the equally natural assumption that all actions are causally determined. But if all actions are causally determined doesn't that mean that freedom of action is impossible? Similarly, if we take it that there is an omniscient, omnipotent and benevolent God doesn't that make it impossible to account for the existence of evil in the world? So if we think that some of our actions are actually free, or that evil actually exists then we ought to be puzzled. How-possible questions are a vivid way of giving expression to this puzzlement. We ask how freedom and evil are possible at the point at which the very existence of freedom and evil begins to strike us as problematic.

On this account, how-possible questions are obstacle-dependent questions. We ask how x is possible when there appears to be an obstacle to the existence of x.⁴ We don't ask how x is possible if there is no perceived obstacle or no inclination to suppose that x is possible. So, for example, we don't ask how baseball is possible or how round squares are possible. Where an obstacle-dependent how-possible question does arise there appear to be two basic strategies for dealing with it. The first is to deny the existence of the obstacle which gave rise to the question. This is an obstacle-dissipating strategy. So, for example, we can try to explain how freedom is possible by denying that all actions are causally determined, or how evil is possible by denying the existence of God.⁵ A different approach would be to argue that freedom is possible even if all actions are causally determined, or that evil is possible even if God exists. These are obstacle-overcoming rather than obstacle-dissipating strategies since they don't straightforwardly deny the existence of the obstacle in either case. What they deny is that the alleged obstacles are insuperable and, in this sense, genuine.⁶

What if it turns out that the obstacle which got the discussion going in the first place can neither be dissipated nor overcome? If we don't want to deny that all actions are causally determined but can't think of a plausible way of reconciling freedom and determinism then it seems that we are going to have to accept that freedom is impossible. This looks like a form of scepticism, so the next question is whether we can live with scepticism. Maybe we can. On the other hand, it might turn out that our intellectual commitment to the possibility of acting freely is stronger than our intellectual commitment to the principle that all actions are causally determined, or to the incompatibility of freedom and determinism. In that case, we should look again at whether the alleged obstacle to freedom can be overcome or dissipated. The worst case scenario in philosophical terms is if scepticism about the possibility of freedom is intolerable but the obstacle continues to strike us as insuperable however hard we try to make it disappear or find a way around or over it.

My concern is with epistemological rather than metaphysical, ethical, or theological how-possible questions. Not surprisingly, discussions of how-possible questions in the theory of knowledge have the same basic structure as discussions of how-possible questions in other branches of philosophy. Epistemological how-possible questions start from some cognitive achievement which they assume to be genuine and ask how that achievement is possible. The most general epistemological how-possible question, which is also often represented as one of the defining questions of epistemology, is:

(HP) How is knowledge possible?

As it stands, this question is entirely general in scope. It asks how it is possible for anyone or anything to know anything about anything. Typically, however, epistemological how-possible questions are more specific. For a start, (HP) is usually read as a question about human knowledge, as asking how it is possible for us and creatures like us to know anything.⁷ But even this looks too general to be manageable. A better approach is to concentrate on specific types of knowledge. For example, we might ask how knowledge of the world around us, as distinct from self-knowledge or knowledge of logic, is possible. Other equally familiar epistemological how-possible questions concern a priori knowledge and knowledge of other minds. In each case, however, the basic idea is the same. We assume, at least to begin with, that actually we have some knowledge of the kind in question, and then try to account for its possibility.

A priori knowledge and knowledge of other minds will be the focus of later chapters.

In this chapter I want to take a closer look at the following question:

(HP_{ew}) How is knowledge of the external world possible?

The 'external world' is the world around us, and to have knowledge of the external world is to have knowledge of the existence and nature of objects, processes, events, or states of affairs which exist independently of human thought or perception.⁸ My aim here is not just to outline

an answer to (HP_{ew}) but also, at the same time, to develop a general approach to answering epistemological how-possible questions which can be used to explain how other kinds of knowledge are possible. The specific approach to epistemological how-possible questions which I will be defending is what I am going to call a multi-levels approach. As we will see, the tackling of obstacles to knowledge, either by dissipating or overcoming them, is one element of this approach but not the only element.

What is the intuitive obstacle to our knowing anything about the world around us which might lead one to ask (HP_{ew})? The fact is that there are many such obstacles, and that there is no hope of our getting to the bottom of every consideration which might conceivably result in someone asking how knowledge of the external world is possible. Obstacles can only be dealt with on a piecemeal basis, as and when they are encountered. If this is right then the important thing is to have a good general sense of how to tackle the most commonly encountered obstacles. For illustrative purposes, therefore, I am going to concentrate on one consideration, or set of considerations, which has led some philosophers to wonder how we can know anything about the world around us. Hopefully, it will turn out that other commonly encountered obstacles can be dealt with in the same way.

The particular route to (HP_{ew}) which I have in mind has as its starting point a ‘how question’ rather than a how-possible question. The idea is this: if someone is said to know that p, where p is a proposition about the external world, we can always ask how he knows it. To ask how he knows it is to ask how he came to know it; it is a question about the source of his knowledge or his route to the knowledge that p.⁹ There need be no implication that the truth of p is something that he couldn’t know. This is what distinguishes a simple how-question, such as ‘how does S know that p?’, from the corresponding how-possible question, ‘how is it possible for S to know that p?’.¹⁰ The how-question might be prompted by nothing more than what Austin calls ‘respectful curiosity’ (1979: 78). How-possible questions, in contrast, are

usually expressions of more than just respectful curiosity. They are not just questions but challenges.

How, then, do we come to know things about the world around us? What is the source of our knowledge of the external world? The obvious answer is that there are lots of different ways of coming to know things about the external world.¹¹ Looking around, talking to people, reading newspapers, doing Google searches, are all ways of acquiring wordly knowledge, so what is the problem? One problem is that not all sources of knowledge are equal. Some are more basic than others, and the most basic general source of all is sense-perception.¹² But why is that a problem? Because, according to one line of thinking, ‘there are certain apparently undeniable facts about sense-perception’ that make it difficult to understand ‘how we could get any knowledge at all of the world around us on the basis of sense-perception’ (Stroud 2000a: 5). If it is difficult to understand how we could get any knowledge of the world around us on the basis of sense-perception then it is difficult to understand how we could get any knowledge of the world around us. And that is why (HP_{ew}) looks like a good question. On this account, the alleged obstacles to the acquisition of perceptual knowledge are among the epistemological obstacles which might reasonably lead us to wonder more generally how any knowledge of the external world is possible.

This way of getting to (HP_{ew}) brings out the limitations of what might otherwise seem the best way of tackling this question. Consider again the relationship between ‘how’ and ‘how-possible’ questions. We can explain how S knows that p by figuring out how he came to know that p, but figuring out how he actually came to know that p is also a way of explaining how it was possible for him to know that p. This suggests that we can explain how knowledge of the external world is possible by identifying ways or means of coming to know something about the world around us.¹³ Let’s call this a Means Response to (HP_{ew}). A means response to a how-possible question regards the identification of one or more of the means by which

something can come about as a means of explaining how it is possible. So, for example, if perceiving is a means of coming to know something about the world around us then it is also a means by which knowledge of the external world is possible.

But this can't be the end of the story if there are apparently undeniable facts about sense-perception which threaten to make the acquisition of perceptual knowledge impossible. Given the obstacles to the acquisition of perceptual knowledge, all that the proposed Means Response to (HP_{ew}) does is to shift the focus of discussion from this question to another how-possible question, namely:

(HP_{pk}) How is perceptual knowledge possible?

What is the relationship between (HP_{pk}) and (HP_{ew})? Suppose that perceptual knowledge is defined as knowledge of the world that has its source in perception or perceptual experience. To define perceptual knowledge in this way is to define it both by reference to its source (perception) and its subject-matter (the external world). In contrast, the knowledge that is at issue in (HP_{ew}) is identified just by its source. This means that (HP_{ew}) and (HP_{pk}) call for somewhat different responses. It is a substantive claim that knowledge of the external world is possible by perceptual means but not that perceptual knowledge is possible by perceptual means. If there is such a thing as perceptual knowledge then, by definition, perception is its source.

It doesn't follow, however, that there is no such thing as a Means Response to (HP_{pk}). Suppose that I know that the cup into which I'm pouring coffee is chipped. How do I know this? By seeing that it is chipped. Seeing that the cup is chipped is a particular way of coming to know that the cup is chipped, and therefore a particular means by which perceptual knowledge is possible. Feeling that the cup is chipped is a different way of coming to know that the cup is chipped and therefore a different means by which perceptual knowledge is possible. These are examples of Means Responses to (HP_{pk}). Such responses are possible

because ‘perception’ is a generic source of knowledge, and perceptual knowledge has different specific sources. To describe a piece of knowledge as ‘perceptual’ is to abstract from differences between different kinds or modes of perception, and that is why there is scope for a Means Response to (HP_{pk}).

Let us now return to the suggestion that it is difficult to understand how perception can provide us with knowledge of the world around us. What is the significance of this suggestion? One might think that its significance for (HP_{ew}) is limited because we could agree that knowledge of the world around us isn’t possible by perceptual means without agreeing that it isn’t possible by other means.¹⁴ On the other hand, it does seem right that perception is, for us, an absolutely fundamental source of knowledge, and that our epistemic situation would be very different from what we normally take it to be if perceptual knowledge isn’t possible. There are many things about the world our knowledge of which is ultimately perceptual, so it is not as if we can easily give up on perceptual knowledge and concentrate on other sources. What this amounts to is the suggestion that if we are seriously interested in answering (HP_{ew}) then we can’t just ignore (HP_{pk}).

This brings us back to the idea that there are obstacles to the acquisition or existence of perceptual knowledge. It remains to be seen what these obstacles are, but the implication of the discussion so far is that explaining how perceptual knowledge is possible has basically got to be a matter of identifying more specific perceptual means by which it is possible and dissipating or overcoming the alleged obstacles. Suppose, for example, that one particular obstacle takes the form of an epistemological requirement on the acquisition of knowledge, and that it is hard to see how perceptual knowledge is possible because it is hard too see how we could possibly meet the suggested requirement. In that case, an obstacle-overcoming response to (HP_{pk}) and, by implication to (HP_{ew}), would be one which shows that we can satisfy the requirement. An obstacle-dissipating response, in contrast, would be one which

makes it plausible that there is no such requirement. Indeed, on a Moorean view, the very fact that a particular epistemological requirement calls the possibility of perceptual knowledge into question is a good reason for rejecting that requirement.

I will come back to this suggestion later in this chapter. First, there are some other matters to discuss. So far, I have suggested that a good answer to an epistemological how-possible question will have two at least dimensions: it will identify specific means by which a particular kind of knowledge is possible and it will remove obstacles to the acquisition of knowledge by the proposed means. Is that enough? According to a position which I'm going to call explanatory minimalism, it is enough: explaining how knowledge of a certain subject-matter is possible is simply a matter of identifying means of knowing about that subject-matter, and showing that there is nothing that stands in the way of our knowing about that subject-matter by the proposed means. Once we done that, we have done everything that can or needs to be done to answer the original how-possible question. In the case of propositions about the external world, therefore, the explanatory minimalist thinks that we have shown how knowledge of such propositions is possible by pointing out that perception is a means of coming to know things about the external world, and showing that there is nothing that stands in the way of our coming to know things about the external world by perceptual means.

Later in this chapter, I will be taking a closer look at explanatory minimalism. I will be defending the suggestion that more can be done to explain how knowledge is possible than it allows. Having identified perception as a means of knowing about the external world, and satisfied ourselves that there is nothing that prevents us from acquiring knowledge of the external world from this source, there is a further question which we might ask. The further question is: what makes it possible for us to acquire knowledge of the things around us by means of perception? This is a question about what might be called the enabling conditions of perceptual knowledge, the conditions under which it is possible for perception to be a source

of knowledge of the things around us. The thought that underpins this question is that there is more to explaining how something is possible than showing that it isn't impossible.¹⁵ What we want is, as it were, a positive explanation of the possibility of perceptual knowledge, that is, an account of what makes perceptual knowledge possible.¹⁶

What we now have, in outline at least, are all the elements of a multi-levels account of the possibility of perceptual knowledge or knowledge of the external world. A multi-levels response to a how-possible question operates on three different levels. Level 1 is the level of means, the level at which means of knowing about a certain subject-matter are identified. In the case of (HP_{ew}) the proposed means might be as generic as 'perception' or 'testimony'. In contrast, (HP_{pk}) calls for the specification of more specific means since perceptual knowledge is defined by its generic source as well as its subject-matter. In neither case, however, can the how-possible question be satisfactorily answered without saying something about means or ways of knowing. Level 2 is the obstacle-removing level, the level at which obstacles to the acquisition of knowledge by the proposed means are overcome or dissipated. Finally, Level 3 is the level at which enabling conditions for knowing by the proposed means are identified. The minimalist wants to stop at Level 2. The contrasting anti-minimalist approach says that we can or should continue to the level of enabling conditions, and that there are worthwhile philosophical questions about perceptual knowledge which are best understood as Level 3 questions.

This is only a sketch of the main elements of a multi-levels response, and it clearly needs a lot of fleshing out. A good way of doing this would be to look at an actual example of a multi-levels response to a how-possible question in the history of philosophy. The particular philosopher who comes to mind in this connection is Kant. He did more than anyone to bring how-possible questions to prominence in philosophy, and his approach to such questions is a multi-levels approach. But his topic, at least in the Introduction to the first Critique, is not the

possibility of perceptual knowledge ,or anything as general as ‘our knowledge of the external world’. Instead, his how-possible questions all boil down to the single question:

(HP_{sap}) How is synthetic a priori knowledge possible?

Nevertheless, if it is true that Kant’s response to (HP_{sap}) is a multi-levels response then it would be worth taking a closer look at what he says. As will shortly become apparent, his response to (HP_{sap}) raises some important general questions about the multi-levels approach, and these questions will need to be addressed before we return to the project of developing a multi-levels response to (HP_{pk}).

1.2 Kant’s Problem

The background to Kant’s interest in (HP_{sap}) is his conviction that pure mathematics and pure physics are bodies of synthetic a priori knowledge. A priori knowledge is knowledge that doesn’t have its source in experience. Synthetic knowledge is non-analytic knowledge, knowledge that doesn’t have its source in ‘the analysis of concepts’ (B23).¹⁷ Our knowledge of the propositions of pure mathematics and pure physics must be a priori, according to Kant, because these propositions are necessarily true, and because experience can’t provide us with knowledge of necessary truths. At the same time, given that the propositions of pure mathematics and pure physics aren’t analytically true, and that conceptual analysis can’t provide us with knowledge of propositions which aren’t analytically true, mathematical and pure scientific knowledge must be synthetic.

Kant remarks that since pure mathematics and pure natural science actually exist, ‘it is quite proper to ask how they are possible; for that they must be possible is proved by the fact that they exist’ (B20). In contrast, metaphysics ought to contain synthetic a priori knowledge, but ‘we cannot assume metaphysics to be an actual science’ (Kant 1977: 275). Thus, as far as metaphysics is concerned, the issue is whether it is possible; in the case of pure mathematics

and pure natural science we only have to think about how they are possible. In the words of the Prolegomena:

We have therefore some, at least uncontested, synthetic a priori knowledge, and need not ask whether it be possible (for it is actual) but how it is possible, in order that we may deduce from the principle that makes the given knowledge possible the possibility of all the rest (1977: 275).

So the idea is that once we understand how uncontested synthetic a priori knowledge is possible, we should be in a better position to establish and explain the existence of more controversial forms of synthetic a priori knowledge, including metaphysical knowledge.

Suppose that Kant is right to assume that pure mathematics and pure physics are bodies of synthetic a priori knowledge. Why, then, is it ‘quite proper to ask how they are possible’? Given my account of how-possible questions, this question only arises if synthetic a priori knowledge is in some way problematic, that is, if there is at least an apparent obstacle to its existence. The apparent obstacle is this: according to Kant, experience and conceptual analysis are two basic sources of human knowledge, but we have already seen that neither experience nor conceptual analysis can yield synthetic a priori knowledge. Experience can’t yield a priori knowledge and conceptual analysis can’t yield synthetic knowledge. That is why it is proper to ask how synthetic a priori knowledge is possible. What gives this question its philosophical bite is the existence of synthetic a priori knowledge that can’t be accounted for by reference to certain presupposed basic sources of knowledge. Let’s call this the problem of sources.

One response to this problem would be to argue that it’s false that neither experience nor conceptual analysis can account for mathematical knowledge. Conceptual analysis can account for it if mathematical truths are analytic rather than synthetic. Alternatively, there is no reason why mathematical knowledge couldn’t come from experience if mathematical

truths aren't necessary, or if it is false that experience can't provide us with knowledge of a necessary truth. Each of these obstacle-dissipating responses to Kant's question amounts to what might be called a presupposed sources solution to the problem of sources. In each case the possibility of mathematical knowledge is accounted for by reference to one of the presupposed sources of knowledge. However, this isn't Kant's own preferred solution. His solution is an additional sources solution since it involves the positing of what he calls 'construction in pure intuition' as an additional source of knowledge by reference to which at least the possibility of geometrical knowledge be accounted for.¹⁸ It is 'additional' not in the sense that mathematicians haven't been constructing in pure intuition all along but in the sense that no account was taken of this source of mathematical knowledge in the discussion leading up to the posing of (HP_{sap}).

Viewed in one way, the additional sources solution looks like an obstacle-overcoming rather than an obstacle-dissipating response to (HP_{sap}) since it doesn't dispute the suggestion that neither experience nor conceptual analysis can account for our mathematical knowledge. It endorses this suggestion but argues that mathematical knowledge is still possible. Viewed in another way, however, Kant's solution to the problem of sources looks more dissipationist. What gives rise to this problem in its sharpest form is the assumption that experience and conceptual analysis are our only sources of knowledge. Kant challenges this assumption. He implies that if we had acknowledged the existence of non-conceptual yet non-empirical sources of knowledge in the first place we wouldn't have been puzzled by the existence of synthetic a priori knowledge. This suggests that there isn't a genuine obstacle to the existence of this kind of knowledge, only a false assumption about the range of epistemic sources that are available to us. Once we reject this assumption the perceived obstacle disappears. So what started out as an obstacle-overcoming exercise has ended up as an exercise in obstacle dissipation.

When Kant asserts that geometrical knowledge is the knowledge gained by reason from the construction of concepts he is reminding his readers that Euclidean geometrical proofs are essentially diagrammatic. To construct a figure in pure intuition is to ‘draw’ it in the imagination, and Kant’s proposal is that the construction of geometrical concepts in pure intuition is a genuinely non-conceptual, non-empirical means of coming to know geometrical truths, and therefore a means of acquiring synthetic a priori knowledge. In my terms, this is a Level 1 response to (HP_{sap}). In keeping with the general idea of a Means Response to a how-possible question, Kant is assuming that the identification of means by which a specific type of knowledge is possible is a means of explaining how knowledge of that type is possible.

Let’s take a closer look at this assumption. The first thing to notice is that the means by which something is possible needn’t be necessary conditions for its possibility. The precise significance of this point will become clearer as we go along. In the meantime, a simple example might help to clarify the distinction between means and necessary conditions: if someone asks how it is possible to get from London to Paris in less than three hours, it would be an acceptable answer to say that one can do this by catching the Eurostar. Catching the Eurostar is a means of getting from London to Paris in less than three hours but obviously not a necessary condition; one can also fly. By the same token, there is a difference between identifying a means achieving of something and showing that it is the only means of achieving it. Catching the Eurostar is a means of getting from London to Paris but not the only means.

The question which this example raises is whether, assuming that one’s sole aim is to answer (HP_{sap}), it matters whether construction in intuition is the only means of acquiring synthetic a priori knowledge. It would seem not. If one lands up asking (HP_{sap}) because one thinks that some of our knowledge is synthetic a priori, and because one can’t figure out how synthetic a priori knowledge is possible, it is enough that intuitive construction is a means of

acquiring synthetic a priori knowledge. If it is the only means then that would obviously be interesting but it is hard to see why a good answer to (HP_{sap}) has got to establish unique means. As far as Kant is concerned, constructing concepts in pure intuition is our only route to geometrical knowledge but not our only route to synthetic a priori knowledge. He thinks that some of our synthetic a priori knowledge is metaphysical rather than mathematical, and that there is no question of construction in pure intuition being the source of our non-mathematical synthetic a priori knowledge. So his basic answer to (HP_{sap}) is that synthetic a priori knowledge is possible by several different means, construction in pure intuition being one of these means.

Even if one sticks to one specific type of synthetic a priori knowledge, say geometrical knowledge, it's not obvious how one would establish unique means. For example, how can we be sure that some form of 'rational insight' couldn't possibly be a source of geometrical knowledge? In reality, however, this question is neither here nor there given that rational insight is not a source of knowledge that is available to us. The lesson is that a good response to a how-possible question needs to identify practical rather than unique means. For the purposes of explaining how geometrical synthetic a priori knowledge is possible, all that matters is that the construction of concepts is a source of this kind of knowledge, and that it is something that we are actually capable of doing.

At this point we run into the following problem: if the recommended answer to 'how is it possible for us to acquire synthetic a priori geometrical knowledge?' is: 'by constructing concepts in pure intuition', one might reasonably think that this only raises a further and no less pressing how-possible question, namely, 'how is it possible for construction in pure intuition to be a source of synthetic a priori knowledge?'. If how-possible questions are obstacle-dependent, then someone who asks this question must have a particular obstacle in mind, and it's not difficult to guess what that obstacle might be. According to Kant, strict

universality is a criterion of a priori knowledge, so a priori geometrical knowledge must be of strictly universal propositions. Yet what we construct in pure intuition are individual figures. So if we want to understand how construction in pure intuition can be a source of a priori geometrical knowledge, we must explain how, as Kant puts it, the construction of a concept by means of a single figure can ‘express universal validity for all possible intuitions which fall under the same concept’ (A713/ B741). Let’s call this the problem of universality for Kant’s theory of geometrical proof.

Kant’s solution to this problem is buried in the depths of the Schematism.¹⁹ Roughly, the idea is that geometrical proofs have universal validity as long as the figures that mathematicians construct are ‘determined by certain universal conditions of construction’ (A714/ B742). These universal conditions of construction are the schemata of geometrical concepts, that is, rules for constructing them. This is now a Level 2 response to (HP_{sap}). Having identified construction in pure intuition as a means of acquiring synthetic a priori knowledge, and acknowledged an apparent conflict between the singularity of constructed figures and the strict universality of a priori geometrical propositions, Kant draws on the Schematism in an attempt to overcome this apparent obstacle to the acquisition of geometrical knowledge by the proposed means. In essence, his proposal is that it is the fact that construction is a rule-governed activity that makes it possible for geometry to discern ‘the universal in the particular’ (A714/ B742).

Since we are presently concerned with the structure of Kant’s account rather than its details, it doesn’t matter very much whether this exercise in obstacle-removal is successful. The important point, as far as Kant is concerned, is that the problem of universality must be soluble because existence of synthetic a priori geometrical knowledge is a given. He is certain that geometry is a body of synthetic a priori knowledge, and that construction in pure intuition is its source. What we have to do, therefore, is to figure out how these things are possible. He

thinks that it's just not an option to end up concluding that we don't have any synthetic a priori knowledge after all, or that we can't get it from construction in pure intuition. In both these respects, Kant's attitude towards (HP_{sap}), at least in the Prolegomena, might be described as Moorean. What this means is Kant regards the claim that synthetic a priori knowledge exists as non-negotiable, somewhat in the way that Moore regards the existence of perceptual knowledge as non-negotiable.²⁰ For Kant, it makes no more sense to question the existence of synthetic a priori knowledge than to question the existence of pure mathematics.

For the moment, this is as much as I want to say about Kant's Level 2 response to (HP_{sap}). Let us now consider whether, if Kant's story so far were convincing, he would need to go any further in relation to (HP_{sap}). The explanatory minimalist thinks not but Kant disagrees. His thought is that even though construction in pure intuition is bona fide source of synthetic a priori geometrical knowledge, there is more that needs to be said. What still needs to be explained is the capacity of construction in pure intuition to provide us with geometrical knowledge. Explaining this isn't just a matter of dealing with actual or apparent obstacles to the acquisition of geometrical knowledge; it's also a matter of explaining what makes it possible for construction in pure intuition to be a source of synthetic a priori geometrical knowledge. What we now have, therefore, is a what-makes-it-possible question rather than a how-possible question. How-possible questions are obstacle-dependent but what-makes-it-possible questions are explanation-seeking.²¹ What they seek is not a way round some specific obstacle but, as it were, a positive explanation of the possibility of acquiring a certain type of knowledge by certain specified means.

Everything now depends on the nature of the required explanation and the reasons for thinking that any such explanation is either necessary or possible. On the first of these issues, there are at least two different things which might be involved in explaining what makes it possible for construction in pure intuition to be a source of geometrical knowledge. On the

one hand, it might involve explaining what makes construction in pure intuition possible. Let's call this a type A explanation. In my terms, a type A explanation is one that seeks to explain the possible occurrence of a certain cognitive activity. On the other hand, one might think that what needs explaining is not just the possibility of constructing figures in pure intuition but also the fact that doing this is a source of a certain kind of knowledge. Let's call this a type B explanation; what makes it a type B explanation is that it seeks to explain the epistemological significance of a certain cognitive activity.

What would a type A explanation of the possibility of constructing figures in pure intuition look like? To explain what makes this possible would be to identify the conditions under which it is possible. The conditions under which it is possible are its enabling conditions. There are many different kinds of type A enabling condition that one might have in mind at this point. For example, suppose that only thinkers with a certain physiology could construct figure in pure intuition. Possession of the appropriate psychology would constitute one kind of enabling condition, a physiological enabling condition. Another line of thinking would be that only thinkers with certain other cognitive capacities, such as the capacity to see or touch, would be able to construct concepts in pure intuition. Such cognitive capacities would constitute cognitive enabling conditions for construction in pure intuition.

What are enabling conditions? In essence, they are a sub-class of necessary conditions. For example, the existence of at least one thinker is a necessary condition for the construction of figures in intuition but it isn't what I am calling an enabling condition. The existence of at least one thinker is necessary for there to be any thinking at all, not just necessary for the kind of thinking that is involved in the construction of geometrical proofs. In contrast, enabling conditions are more specific; enabling conditions for the construction of figures in intuition aren't necessary condition for there to be any cognitive activity whatsoever. Enabling conditions are necessary conditions for achieving something by a particular means. Relatedly,

enabling conditions are background conditions, which may or may not be causal.²² Being an unmarried man is a necessary condition for being a bachelor but being an unmarried man isn't an enabling condition for being a bachelor. Intuitively, the reason is that being an unmarried man isn't a 'background condition' for being a bachelor. Being an unmarried man doesn't just 'enable' one to be a bachelor, it is what being a bachelor consists in.

Burge makes the notion of a 'background condition' a bit more precise in his paper 'Content Preservation'. One of his examples concerns the role of memory in deduction. The important point is that claims about memory or the past needn't figure as premises in one's reasoning even though any reasoning in time must rely on memory. Specifically:

In a deduction, reasoning processes' working properly depends on memory's preserving the results of previous reasoning. But memory's preserving such results does not add to the justificational force of the reasoning. It is rather a background condition for the reasoning's success.... Memory failures that cause demonstrations to fail are failures of background conditions necessary to the proper function of reasoning (1993: 463-4).

What we have here is, in effect, an answer to a what-makes-it possible question. The question is: 'what makes transitions of reason possible?', and the answer is: preservative memory. Something similar can be said about the role of perception in interlocution, since perception is a background condition necessary for the acquisition of beliefs from others:

In interlocution, perception of utterances makes possible the passage of propositional content from one mind to another rather as purely preservative memory makes possible the preservation of propositional content from one time to another. Memory and perception of utterances function similarly, in reasoning and communication respectively. Their correct functioning is necessary for the enterprises they serve (Burge 1993: 481).

But the role of memory and perception in these enterprises is only to ‘preserve and enable’. In my terms, memory is an enabling condition for reasoning and that perception is an enabling condition for interlocution.

With this account of the notion of enabling conditions in place, we can now briefly examine Kant’s type B explanation of what makes it possible for the construction of figures in pure intuition to be a source of geometrical knowledge. The proposed explanation is metaphysical rather than one in terms of cognitive capacities. Geometry is supposed to tell us something about the nature of physical space, so a type B explanation will be an account of what makes it possible for the construction of figures in pure intuition to yield knowledge of the geometry of physical space. According to Kant, what makes this possible is the fact that space itself is an ‘a priori intuition’ that ‘has its seat in the subject only’ (B41). The proposal, in other words, that the transcendental ideality of space is a background enabling condition for the acquisition of geometrical knowledge by means of construction in intuition. That’s why, if we want to explain how geometrical knowledge is possible, we must be transcendental idealists.

How convincing is this argument? We might fail to be convinced by it because we reject its starting point, because we think that Kant was wrong to regard geometry as a body of synthetic a priori knowledge or as providing us with a priori knowledge of physical space. Or we might think that his solution to the problem of universality fails for one reason or another. Finally, we might reject the last step of the argument on the grounds that it fails to establish that it would be impossible to account for synthetic a priori knowledge of physical space outside an idealist framework. We might wonder, for example, why it wouldn’t be possible for us to have innate synthetic a priori knowledge of the geometry of physical space even if space isn’t ideal. If this is a genuine possibility then the argument for idealism doesn’t go through.

We don't need to go into these issues here. Instead, let's focus on the relevance of Kant's position for the dispute between minimalism and anti-minimalism. Suppose, then, that what is at issue is the acquisition of knowledge of kind K by means M. Minimalists shouldn't deny that there might be causal enabling conditions for the acquisition of K by M. What they should say is that the uncovering of such conditions is a matter for empirical science rather than armchair philosophy.²³ Minimalism is therefore the view that distinctively philosophical explanations of the possibility of knowledge can't go beyond Level 2.²⁴ In contrast, moderate anti-minimalism is not sceptical about the existence of enabling conditions for the acquisition of K by means of M which philosophical reflection can bring to light. It agrees, therefore, that philosophical Level 3 explanations are possible. It denies, however, that such explanations are necessary for the purpose of explaining how knowledge of K kind is possible. Finally, extreme anti-minimalism is the view that philosophical Level 3 explanations are both possible and necessary. According to the extreme anti-minimalist, if there are Level 3 conditions for the acquisition of K by M that philosophy, and perhaps only philosophy, can bring to light, then a philosophical explanation of how knowledge of kind K is possible by means M will be incomplete unless it says what these conditions are.

In these terms, Kant's account of what makes it possible for the construction of figures in pure intuition to be a means of acquiring geometrical knowledge looks like an implicit argument for extreme anti-minimalism. The argument goes like this: it is a priori philosophy rather than empirical science that tells us that the ideality of space is what makes it possible for construction in pure intuition to be a source of geometrical knowledge. So it is false that a philosophical account of what makes geometrical knowledge possible can't go beyond Level 2; Kant's is a philosophical account of what makes geometrical knowledge possible that does go beyond Level 2. It is also false that we can adequately explain how geometrical knowledge is possible without saying anything about the metaphysical conditions under which this is

possible. If it really only makes sense to suppose that construction in pure intuition is a source of geometrical knowledge on the assumption space is ideal, how can reference to this fact possibly be an optional extra in a philosophical explanation of the possibility of geometrical knowledge? This, then, is the basis of Kant's extreme anti-minimalism. It is the idea that it is essential for a philosophically satisfying response to (HP_{sap}) to reach all the way down to the level of enabling conditions.

From a multi-levels perspective the problem with this argument is that it calls into question the distinction between Level 2 and Level 3. For surely Kant only thinks that the ideality of space is a necessary condition for the acquisition of geometrical knowledge by means of construction in intuition because he takes it that the mind-independence of space would constitute a kind of obstacle to our acquiring knowledge of its geometrical properties by these means. But this looks like a reason for locating his argument for the ideality of space at Level 2 rather than at Level 3. More generally, if C is an enabling condition for the acquisition of knowledge of kind K by means M then the non-fulfilment of C can't fail to represent an obstacle to the acquisition of K by M. But if the denial of a putative enabling condition is always an obstacle to knowledge doesn't it follow that the suggested distinction between Level 2 and Level 3 explanations is spurious?²⁵ In that case, the only sense in which a philosophically satisfying response to (HP_{sap}) must reach all the way down to the level of enabling conditions is that such an account can't afford to ignore obvious obstacles to the acquisition of synthetic a priori knowledge by means of construction in pure intuition.

What is right about this is that when an enabling condition C for the acquisition of K by M isn't fulfilled, the very fact that it isn't fulfilled becomes an obstacle to the acquisition of K by M. It doesn't follow from this that the point of identifying C must be to defuse some intuitive, pre-existing obstacle to the acquisition of K by M. Suppose, for example, that it is true that in order to be able to construct concepts in intuition one must have the capacity to

see or touch. There is no intuitive obstacle to the acquisition of geometrical knowledge by means of construction in intuition that one would be overcoming or dissipating by identifying such perceptual capacities as background necessary conditions for construction in intuition. If one lacked the relevant perceptual capacities that would make it impossible for one to acquire geometrical knowledge by constructing concepts in intuition but that doesn't make the mere identification of enabling conditions an exercise in obstacle-removal. Identifying what makes it possible for construction in pure intuition to be a source of geometrical knowledge is not primarily a matter of responding to independent reasons for thinking that it couldn't be.

This suggests that there is a distinction between Level 2 and Level 3 explanations of the possibility of knowledge even if, as Kant's argument for the ideality of space illustrates, it is sometimes unclear whether a particular explanation belongs at Level 2 or at Level 3. Once it is agreed, however, that there are some explanations of the possibility of knowledge that clearly belong at Level 3 rather than at Level 2 then moderate anti-minimalism begins to look like a serious option. For example, it is not obvious that someone who thinks that one would not be able to construct concepts in pure intuition if one lacked certain perceptual capacities is committed to thinking that an explanation of the possibility of geometrical knowledge which fails to mention this fact is necessarily inadequate. Given the obstacle-dependence of how-possible questions, an inadequate explanation of how knowledge of geometry is possible is not one that fails to identify enabling conditions. An inadequate explanation is one that fails to come up with viable means of acquiring geometrical knowledge and which therefore fails to solve the problem of sources.

The point of looking at Kant's response to (HP_{sap}) was to flesh out the idea of a multi-levels response to an epistemological how-possible question. Kant's multi-levels response to (HP_{sap}) is illuminating because it shows how obstacle-overcoming can sometimes shade off into obstacle dissipation. It doesn't show that there is no distinction between Level 2 and

Level 3 explanations of the possibility of knowledge but it does suggest that this distinction is not always as sharp as my initial discussion might have suggested. It brings into focus the differences between different kinds of enabling condition, and between different forms of anti-minimalism. It brings out the force of extreme anti-minimalism but fails to establish that, at least in relation to (HP_{sap}), moderate anti-minimalism isn't in the running.

Bearing all of this in mind, let's now go back to (HP_{pk}), and to the project of filling in a multi-levels response to this question. We already have seen how (HP_{ew}), which is where we started, can lead to (HP_{pk}). We have the outlines of a multi-levels response to (HP_{pk}) but very little sense of how to fill in the details. If there are obstacles to the acquisition of knowledge of the external world by perceptual means then we haven't explained how knowledge of the external world is possible unless we have tackled these obstacles. Having done that, we can then consider what makes it possible to acquire perceptual knowledge, and whether we should be minimalists, moderate anti-minimalists or extreme anti-minimalists in relation to (HP_{pk}) and, by implication, in relation to (HP_{ew}).

1.3. Perceptual Knowledge

The first problem we encounter when we try to apply Kant's multiple-levels account of (HP_{sap}) to (HP_{pk}) is this: the possibility of synthetic a priori knowledge needs explaining because neither experience nor conceptual analysis can be its source. In contrast, it is not the case that experience can't be the source of perceptual knowledge, so it seems that (HP_{pk}) is going to have to be motivated by something other than the problem of sources. Indeed, Kant is an example of a philosopher who appears to think that (HP_{pk}) lacks any respectable motivation. In his terms, perceptual knowledge would be synthetic a posteriori knowledge, but 'the possibility of synthetic a posteriori judgements, of those which are gathered from experience.... requires no special explanation; for experience is nothing but a continual joining together (synthesis) of perceptions' (1977: 275). If the possibility of synthetic a

posteriori knowledge requires no special explanation, then (HP_{pk}) simply doesn't arise; there is no obstacle for it to trade on.

The point generalises. Empiricists characterize a priori knowledge 'negatively', as knowledge that doesn't come from experience or that isn't justified by experience. 'Empirical' knowledge is characterized 'positively' as knowledge that comes from experience or that is justified by experience. If what we know about a priori knowledge is that it doesn't come from experience, then an obvious question is:

(HP_{apk}) How is a priori knowledge possible?

Answering this question will be a matter of identifying sources of a priori knowledge, or the means by which it is possible, given that experience isn't its source. What we need, in other words, is at least a Level 1 explanation of the possibility of a priori knowledge. But if empirical knowledge is defined as knowledge that originates in experience, then we already know how it is possible; it's possible by means of experience. The means by which we acquire empirical knowledge is already built into our conception of what makes it empirical knowledge, just as the means by which we acquire perceptual knowledge is built into our conception of what makes it perceptual knowledge. In neither case, therefore, is there any need to look for an explanation of its possibility.

This attempt to bypass (HP_{pk}) fails. For a start, it's easy to imagine a priori knowledge being given the positive and empirical knowledge the negative characterization. A rationalist might stipulate that a priori knowledge is knowledge that is grounded in reason or rational intuition and that empirical knowledge is knowledge that is not so grounded. From this perspective, the pressing question would be:

(HP_{ek}) How is empirical knowledge possible?

This, rather than (HP_{apk}) would now be the pressing question because we would no longer be building the source of empirical knowledge into our conception of what makes it empirical

knowledge. The rationalist would then regard it as a substantive claim that perception is a potential source of empirical knowledge, in the same way that the empiricist regards it as a substantive claim that rational intuition is a potential source of a priori knowledge.

This is just a way of making the obvious point that what strikes one as a worthwhile how-possible question is bound to be influenced by one's background assumptions. Different assumptions throw up different potential obstacles to knowledge and different how-possible questions. Still, it's important not to exaggerate the significance of this consideration. For even if one defines a priori knowledge as knowledge that is grounded in reason, this doesn't mean that (HP_{apk}) can't or doesn't arise. Someone who defines a priori knowledge in this way still needs to explain how reason can be a source of a priori knowledge.²⁶ In explaining how this can be so, one might find oneself running into actual or potential obstacles to the acquisition of a priori knowledge from this source, and this would leave (HP_{apk}) with considerable philosophical bite. By the same token, the fact that empirical or perceptual knowledge is characterized as having its source in experience or perception will still leave (HP_{ek}) and (HP_{pk}) with considerable philosophical bite if there are actual or potential obstacles to the acquisition of knowledge from experience or perception. Either way, defining a type of knowledge by reference to its source needn't prevent one from asking a how-possible question about that type of knowledge; in asking the question, one would be asking how the proposed source can be a source of knowledge of that type.

This attempt to rehabilitate (HP_{pk}) from an empiricist perspective assumes that there are actual or potential obstacles to our coming to know things about the external world by means of perception. What are these obstacles? Many different answers to this question have been proposed but for present purposes I'm going to focus on Stroud's account of one alleged obstacle. Although I will be suggesting that this obstacle isn't genuine, seeing why not will be the first step along the way to the overall approach to (HP_{pk}) that I will be recommending. In

essence, I want to try to make it plausible that there is a viable alternative to explanatory minimalism, and I will do this by developing and defending a moderately anti-minimalist multiple-levels response to Stroud's version of (HP_{pk}).

As we have already seen, Stroud's basic thought is that there are certain apparently undeniable facts about sense-perception that make it difficult to see exactly how sense-perception works to give us knowledge of the world. Specifically, the difficulty is that:

it seems at least possible to perceive what we do without thereby knowing something about the things around us. There have been many versions of that fundamental idea. But whether it is expressed in terms of "ideas" or "experiences" or "sense data" or "appearances" or "takings" or "sensory stimulations", or whatever it might be, the basic idea could be put by saying our knowledge of the world is "underdetermined" by whatever it is that we get through that source of knowledge known as "the senses" or "experience". Given the events or experiences or whatever they might be that serve as the sensory "basis" of our knowledge, it does not follow that something we believe about the world around us is true. The problem is then to explain how we nevertheless know that what we believe about the world is in fact true. Given the apparent "obstacle", how is our knowledge possible? (Stroud 2000a: 5-6)

This passage is in keeping with the idea that (HP_{pk}) is an obstacle-dependent question. Like Kant's worry about synthetic a priori knowledge, Stroud's worry about perceptual knowledge is that it can't easily be accounted for by reference to our presupposed cognitive resources. These resources include perception, but if several different possibilities are compatible with our perceiving what we do then it is hard to see how perception could be a source of knowledge of the things around us. In other words, it is hard to see how perceptual knowledge is possible, and that is why (HP_{pk}) is a genuine question despite the fact that perceptual knowledge is defined as knowledge that we get from perception.

There is, however, one crucial difference between Kant's question about synthetic a priori knowledge and Stroud's question about perceptual knowledge. The difference is that the geometrical version of (HP_{sap}) can be answered by positing construction in pure intuition as an 'additional' source of knowledge, whereas it's obviously not going to be acceptable to respond to (HP_{pk}) by positing additional, non-perceptual sources of perceptual knowledge. It's not as if it would make sense to agree that sense-perception can't give us knowledge of the world and then argue that this doesn't matter because such knowledge is possible by some other means. Trivially, perceptual knowledge is possible only if perception can be a source of knowledge, and it is at least arguable that knowledge of the external world is only possible if perceptual knowledge is possible. So what needs to be made plausible is that perception can provide us with knowledge of the world. This would be a 'presupposed sources' rather than an 'additional sources' response to (HP_{pk}), one which explains how perceptual knowledge is possible by reference to its presupposed, canonical source.

How can perception be a source of knowledge? The obstacle-generating principle that underpins (HP_{pk}) in Stroud's discussion is the principle that:

(U) Our knowledge of the world is underdetermined by what we get through the senses.

If this is the obstacle that stands in the way of the acquisition of perceptual knowledge, then a Level 2 response to (HP_{pk}) can take one of two forms. One possibility would be to accept (U) but argue that it doesn't prevent perception from being a means of coming to know things about the world around us. This would be an obstacle-overcoming response to Stroud's version of (HP_{pk}). The other possibility would be to dispute (U). If this principle is incorrect, then it can't amount to a genuine obstacle to perceptual knowledge. This would be an obstacle-dissipating response to (HP_{pk}). While there might be genuine obstacles to perceptual knowledge, the obstacle-dissipating response to (HP_{pk}) suggests that (U) isn't one of them.

In order to decide between these responses to (HP_{pk}), more needs to be said about (U) and about the sense in which it calls into question the possibility of perceptual knowledge. The point of (U) is to suggest that what we get through the senses is not knowledge of the world but knowledge of something that is epistemically prior to knowledge of the world. Specifically, what we get through the senses is only knowledge of the character of our sensory experiences themselves or knowledge of how things seem to us to be. To say that things of one sort are epistemically prior to things of another sort is to say that ‘things of the first sort are knowable without things of the second sort being known, but not vice versa’ (Stroud 1984: 141). But if knowledge of how things seem to us to be is epistemically prior to knowledge of the world, then the truth of propositions about the external world would need to be inferred from what we get through the senses. Yet inferences from the character of our sensory experiences can’t provide us knowledge of external reality.²⁷ To regard perceptual knowledge as inferential is therefore to call its very possibility into question. So if (U) is a genuine obstacle, then it is one that cannot be overcome; unlike the obstacle to catching a fly ball twenty feet off the ground, there is no ladder that one can use to climb over it.

If the prospects for an obstacle-overcoming response to (HP_{pk}) are really as dim as this argument suggests then an obstacle-dissipating response to (HP_{pk}) must be the way to go. An obstacle-dissipating response to (HP_{pk}) would be one that disputes the existence of the alleged obstacle. Given that ‘the apparent obstacle to our knowledge comes from the doctrine of the epistemic priority of sensory experiences over independently existing objects’ (Stroud 1984: 143), and that (U) is simply an expression of this alleged epistemic priority, the obvious obstacle-dissipating response to (HP_{pk}) would be one that calls this doctrine into question. The claim that needs to be made out, therefore, is that our knowledge of the external world is not underdetermined by what we get through the senses, and that this is so because what we get through the senses is not epistemically prior to knowledge of the world.

To see what such an obstacle-dissipating response to (HP_{pk}) might look like in practice, consider the following example: as I'm pouring myself a cup of coffee I see that the cup is chipped. If I see that the cup is chipped, I know that the cup is chipped. In general, 'S sees that p' entails 'S knows that p'.²⁸ But to know that the cup is chipped is to know something about the external world. So if what I see is that the cup is chipped, then it is not possible for me to perceive what I perceive in this case without thereby knowing something about the external world. Seeing that the cup is chipped is precisely a means of knowing, or of coming to know something about the external world. The knowledge that it yields isn't 'epistemically prior' to the knowledge that the cup is chipped; it is the knowledge that the cup is chipped. It is false, therefore, that our knowledge of the world is always underdetermined by what we get through the senses.²⁹

Seeing that something is the case is a form of what Dretske calls epistemic seeing.³⁰ In epistemic seeing one sees that something is the case, so epistemic seeing is propositional. One can't see that A unless one grasps the proposition A, and one can't grasp this proposition unless one grasps its constituent concepts.³¹ Epistemic seeing is therefore conceptual; I can't see that the cup is chipped unless I have the concepts cup and chipped. Finally, a situation in which one sees that A is 'a type of situation which represents the acquisition of knowledge by visual means' (Dretske 1969: 80). So the sense in which seeing that p is epistemic is that it embodies or involves an epistemic achievement, the acquisition of knowledge. The contrast is with non-epistemic seeing or 'simple' seeing. One can simply see the chipped cup without thereby knowing or believing that it is chipped. In order to see the chipped cup, one doesn't need to have concepts like cup or chipped, and one doesn't need to grasp any proposition.³²

Seeing that the cup is chipped is obviously not the only way of coming to know that it is chipped. One can also come to know that it is chipped by running one's finger along its rim, or by being told that it is chipped. Seeing that A, feeling that A, hearing that A, and so on are

different forms of what might be called epistemic perception, and it seems that we can now explain how perceptual knowledge is possible by drawing on the possibility of epistemic perception. In so far as seeing that A is a way of coming to know that A, perceptual knowledge of the external world is possible by means of epistemic seeing or, more generally, by means of epistemic perception. This is effectively a Level 1 or a Means Response to (HP_{pk}). And if seeing that A is a way of coming to know that A by visual means, then there is nothing that stands in the way of our coming to know things about the external world by means of the senses. This is now a Level 2 response to (HP_{pk}).

In so far as the appeal to epistemic perception is intended as a response to Level 2 (HP_{pk}), is it successful? Does it really dissipate the alleged obstacle to our coming to know things about the external world by means of the senses? On the face of it, there are several reasons why one might fail to be convinced by this attempt at obstacle-dissipation. For example:

(a) One might think that it is just false that ‘S sees that the cup is chipped’ entails ‘S knows that the cup is chipped’. Thus, even if what S sees is that the cup is chipped, what he knows by visual means still falls short of knowledge of the world.

(b) One might deny that it is possible for S to see that the cup is chipped; if S could really see that the cup is chipped, there would be nothing which stands in the way of S’s knowing that the cup is chipped, but this observation leaves (U) and the doctrine of epistemic priority untouched if there is no such thing as epistemic seeing.

(c) One might agree that S can see that the cup is chipped and that ‘S sees that the cup is chipped’ entails ‘S knows that the cup is chipped’, but still deny that this constitutes an adequate answer to (HP_{pk}); the worry here is that we haven’t explained how knowledge of the world is possible just by identifying various different means by which it is possible, such as epistemic perception.

Let's consider these concerns in turn. With regard to (a), it is easy to understand why the entailment from 'S sees that the cup is chipped' to 'S knows that the cup is chipped' might appear problematic. The standard example in this connection is that of a subject who sees that the cup is chipped but who still doesn't know that the cup is chipped because he mistakenly believes that his senses are malfunctioning.³³ If S believes that his senses are malfunctioning then he might very well not believe that the cup is chipped, and if he doesn't believe that the cup is chipped then he doesn't know that it is chipped. But he still sees that the cup is chipped. That is why seeing that the cup is chipped doesn't entail knowing that it is chipped.

Those who think that believing isn't a condition for knowing needn't be worried by such examples because they can describe them as ones in which S sees and thereby knows that the cup is chipped without believing that the cup is chipped.³⁴ But there are also less radical ways of defending the view that seeing that the cup is chipped is a way of knowing that it is chipped. For if S mistakenly believes that his senses are malfunctioning then it is arguable that he fails to see that the cup is chipped. He fails to see that the cup is chipped because he fails to satisfy an intuitive condition on seeing that the cup is chipped: he does not believe the conditions under which he sees the cup are such that the cup would not look the way it looks to him now unless it was chipped.³⁵ So we don't have a case in which S sees but does not know that the cup is chipped. What we have is a case in which S neither sees nor knows that the cup is chipped.

Perhaps, in that case, the problem is, as (b) suggests, that S cannot see that the cup is chipped. One possibility is that there are contingent obstacles which stand in the way of S's seeing any such thing. For example, S might be blind or it might appear to be too dark for S to see anything. If, in such circumstances, S were to assert that the cup is chipped, the question 'how do you know?' would have an obvious point. As Austin points out, 'how' questions are often asked pointedly, the implication being that the person to whom the question is directed

doesn't really know what he claims to know because he is not in a position to know.³⁶ Yet the assumption that S isn't in a position to know might be mistaken. Maybe S isn't blind; perhaps it's just light enough for S to see.

Presumably, those who say that it is not possible for S to see that the cup is chipped do not say this on account of the existence of such contingent obstacles. The obstacles to epistemic seeing which they have in mind are epistemological, and they are allegedly ones which not only prevent S from seeing that the cup is chipped but which also prevent any other person from simply perceiving that something external is the case and thereby knowing that it is the case. For example, it might be claimed that S cannot correctly be said to see that the cup is chipped unless he can eliminate the possibility that he is dreaming or hallucinating, and that these are not possibilities which we are ever in a position to eliminate. On this account, the appeal to the possibility of epistemic perception fails to dissipate the ultimate obstacle to perceptual knowledge since this obstacle is also an obstacle to epistemic perception. Knowing that one is not dreaming at the relevant time is a necessary condition for knowing about the world by means of the senses, and the doctrine of epistemic priority is simply a consequence or reflection of the fact that this is not an epistemological requirement that can be met.³⁷

Faced with this argument, there are basically two ways of pursuing a Level 2 response to (HP_{pk}). One would be to argue that the requirement which calls the possibility of epistemic perception into question is spurious; for S to see that the cup is chipped it must be true that S isn't dreaming or hallucinating, but this doesn't mean that S cannot correctly be said to see that the cup is chipped unless he can eliminate these possibilities. The elimination of these possibilities has a bearing on whether S knows that he sees that the cup is chipped, not on whether he actually sees and thereby knows that the cup is chipped. The other approach would be to accept the epistemological requirement but argue that it can be met because one can know that one is not dreaming. According to McDowell, for example, 'one's knowledge

that one is not dreaming in the relevant sort of situation owes its credentials as knowledge to the fact that one's senses are yielding one knowledge of the environment - something that one does not happen when one is dreaming' (1998a: 238). So if one knows by means of one's senses that the cup is chipped, then one also knows that one is not dreaming.

It might seem that the second of these two approaches fits better with the idea of an obstacle-overcoming response to (HP_{pk}) than with that of an obstacle-dissipating response. Doesn't acceptance of an obstacle-generating epistemological requirement only leave one with the option of trying to overcome it? And doesn't the suggestion that the requirement can be met just amount to the suggestion that the obstacle can in fact be overcome? In fact, this isn't quite right. I have been assuming that obstacle-generating epistemological requirement is just the requirement that:

(D1) In order to know anything about the world by means of the senses one must know that one is not dreaming.

But even if this is a requirement that can be met in the way that McDowell suggests, there is another, more demanding requirement that can't be met in this way. This is the requirement is that:

(D2) In order to know anything about the world by means of the senses one must know that one is not dreaming independently of knowing whatever it is that one takes oneself to know about the world by means of the senses.³⁸

While it might be acceptable to claim conformity to (D1) on the basis that one's senses are in fact giving one knowledge of the environment, it obviously unacceptable to claim conformity to (D2) on this basis. To attempt to rule out the possibility that one is dreaming on this basis would not be to provide independent grounds for thinking that one is not dreaming.

Unlike (D1), therefore (D2) represents an obstacle to the acquisition of knowledge of the world which can't be met or overcome. This strengthens the case for giving up on the idea

of a viable obstacle-overcoming response to (HP_{pk}) . If the obstacle is (D2) rather than (D1), only an obstacle-dissipating response has any chance of succeeding. The position so far is that there is a certain epistemological requirement which can't be met on its own terms and which, if it were genuine, would make the acquisition of knowledge of the world by means of the senses impossible. So what needs to be done to keep alive the possibility of perceptual knowledge is to show that (D2) isn't a genuine requirement. By showing that (D2) doesn't state a necessary condition for the acquisition of knowledge by means of the senses, one would be dissipating rather than overcoming an obstacle to epistemic perception. Having dissipated the alleged obstacle, one could then reasonably continue to insist that seeing that the cup is chipped is a viable means of coming to know that the cup is chipped.

But how does one decide whether a potentially obstacle-generating epistemological requirement such as (D2) is genuine? In deciding whether to accept a particular requirement one factor to take into account is whether it strikes one as intuitively plausible. A related consideration is whether imposition of the requirement would have acceptable consequences. The consequences that are at issue here are epistemological. Epistemological requirements have epistemological consequences, and acceptable epistemological requirements mustn't have unacceptable epistemological consequences. The problematic cases are ones in which a requirement that initially strikes us as plausible turns out to have unacceptable consequences. In such cases, the fact that the requirement has unacceptable consequences needn't deprive it of whatever plausibility it seemed to have before we registered its consequences, so we have either got to live with the consequences or reject a requirement that continues to strike us as plausible.

We now have to explain what it would be for the epistemological consequences of a principle like (D2) to be unacceptable. Here is one proposal: in any serious investigation of the conditions of knowledge, we start off with the idea that there are certain things that we

know, or certain kinds of knowledge that we actually have. We regard some of the knowledge that we take ourselves to have as negotiable and some as more or less non-negotiable. An unacceptable consequence of an epistemological principle would be the undermining of knowledge in the latter category. For example, even if we take ourselves to have some insight into our own motives and desires, it doesn't necessarily count against a theory of mind or knowledge that it implies that many of our beliefs about our deepest motives and desires don't amount to knowledge. Self-knowledge is, to this extent, negotiable. In contrast, there are certain basic forms of perceptual knowledge that are non-negotiable. An example of non-negotiable perceptual knowledge might be my knowledge that the cup into which I am pouring coffee is chipped. The presumption is that I do know this by means of the senses, so it would be an unacceptable consequence of a principle like (D2) if it implies that this is something I couldn't know; the very fact that it carries this implication would be a strike against it.³⁹

Clearly, there is a difference between a piece of knowledge being more or less non-negotiable and its being absolutely non-negotiable. Presumptions to knowledge, even presumptions to basic perceptual knowledge are, in principle, capable of being defeated, but we shouldn't underestimate what it would take to defeat such a presumption. What we would need is an epistemological principle that is at odds with the presumption and that has such overwhelming independent plausibility that even the fact that it threatens to undermine what we previously took to be non-negotiable knowledge doesn't warrant the principles' rejection. This is not how things stand with (D2). Whether or not it is overwhelmingly plausible that in order to know anything about the world by means of the senses one must know that one is not dreaming, it is not overwhelmingly plausible that 'the epistemic status of the thought that one is not dreaming must be established independently of the epistemic status of whatever putative perceptual knowledge of the environment is in question' (McDowell 1998: 225). In

other words, (D2) is negotiable to an extent that the perceptual knowledge whose possibility it calls into question is not, and the fact that the truth of (D2) would make the acquisition of knowledge of the environment by means of the senses impossible is itself a good reason for refusing to endorse this principle. Given that denial of basic perceptual knowledge is an unacceptable consequence of (D2), the right thing to think is that (D2) doesn't state a genuine requirement on perceptual knowledge.

This attempt at obstacle-dissipation might be described as Moorean. In particular, it is reminiscent of what has been described as Moore's 'argument from differential certainty'.⁴⁰ According to this argument, we are much more certain of our basic perceptual knowledge than we are of the premises of any sceptical argument for the impossibility of perceptual knowledge. As Moore recognizes, however, the question which this argument raises is whether the certainty to which it appeals is objective or merely subjective. What it seems to require is the objective certainty or indubitability of one's perceptual knowledge. Yet this is not something that Moore establishes. This shows that the argument from differential certainty has its limitations as an anti-sceptical argument, but there is another way of looking at it. Instead of seeing it as an attempt to refute scepticism, we might simply see it drawing attention to the fact that the epistemological principles we endorse are answerable to our epistemological verdicts in particular cases, and that conflicts between particular verdicts and general principles needn't always be decided in favour of the latter. If we are so strongly committed to a particular verdict as to regard it as non-negotiable, we can assess a general principle by testing its compatibility with that verdict. What is at issue here is commitment rather than certainty, and the question is whether our commitment to thinking that we have some basic perceptual knowledge is, as I have been arguing, stronger than our commitment to (D2).

The sceptic's worry is, of course, that our commitments to particular verdicts might be irrational or unfounded, and that we might therefore be wrong to regard our possession of basic perceptual knowledge as non-negotiable. Since there is no general guarantee that our epistemological commitments are well-founded, it isn't going to be possible to prove to the sceptic's satisfaction that we are actually in possession of any perceptual knowledge. It's fortunate, therefore, that the obstacle-dissipating response to (HP_{pk}) isn't attempting to prove any such thing. The object of the exercise is simply to explain how perceptual knowledge is possible, given that it is possible. By showing that we needn't be committed to obstacle-generating principles such as (D2), we are doing all that needs to be done at Level 2 to secure the possibility of perceptual knowledge. In so far as we have made it plausible that (D2) doesn't state an inescapable requirement on perceptual knowledge, we have left it open that epistemic perception is a possible means of coming to know things about the world around us. We have thereby rebutted the suggestion in (b) that S can't really see that the cup is chipped or know that it is chipped by seeing that it is.

That leaves (c), according to which we still haven't fully explained how perceptual knowledge is possible even after we have both identified epistemic perception as a means of coming to know about the external world and shown that there is nothing that stands in the way of epistemic perception. So what more is required? Just as in the geometrical case the anti-minimalist thinks that what is required is an explanation of what makes it possible for construction in pure intuition to occur and to be a source of synthetic a priori knowledge, so the parallel suggestion in relation to (HP_{pk}) is that the proposed Means Response to this how-possible question needs to be supplemented by answers to two what-makes-it-possible questions. The first is: what makes it possible to perceive that something is the case? The second is: what makes it possible for perceiving that something is the case to be a means of coming to know that it is the case? These are questions about enabling conditions. An answer

to the first question will need to identify type A enabling conditions, that is, the background necessary conditions for the occurrence of epistemic perception. An answer the second question will need to identify type B enabling conditions, that is, the background necessary conditions for epistemic perceiving to be a source of knowledge. The identification of type A and type B enabling conditions is a Level 3 explanation of the possibility of perceptual knowledge, and the point of (c) is to suggest that an adequate answer to (HP_{pk}) must reach all the way to down to this level.

This is effectively an argument for extreme explanatory anti-minimalism in relation to (HP_{pk}), so the issue is whether this form of anti-minimalism is warranted. We need to consider whether Level 3 explanations of the possibility of perceptual knowledge are possible and, if so, whether they are necessary. If they aren't possible then this would count against (c) and in favour of explanatory minimalism. If they are possible but not necessary this would count against (c) and in favour of moderate anti-minimalism. This is the approach for which I want to argue. In opposition to minimalism, this approach allows for the possibility of Level 3 explanations of perceptual knowledge and does not dispute the legitimacy of asking what makes it possible for epistemic perception to occur and to be a source of knowledge. On the other hand, there is no clear sense in which explanations of the possibility of perceptual knowledge that stop at Level 2 are 'incomplete', so there isn't a case for extreme anti-minimalism.

To get a sense of the force of anti-minimalism and of what (c) represents as the limitations of minimalism, it's essential that some of the details of an anti-minimalist account of epistemic perception are filled in. It isn't enough for the anti-minimalist to insist that there are legitimate what-makes-it-possible questions about perceptual knowledge that a philosophically satisfying answer to (HP_{pk}) can or should address. Some indication also needs to be given of what a good answer to these questions might look like in practice, and of how

‘philosophical’ level 3 explanations of the possibility of perceptual knowledge differ from other Level 3 explanations. The next challenge, therefore, is to give a sketch of some of the type A or type B enabling conditions that might figure in a fully-fledged multiple-levels response to (HP_{pk}). Having done this, we will be in a better position to assess the relative merits of moderate and extreme anti-minimalism.

1.4. Anti-Minimalism

What makes it possible to see that the cup is chipped? In the primary sense of seeing that the cup is chipped, this is not something that one could see without seeing the cup. So as long as we stick with the primary use of constructions of the form ‘S sees that b is P’, we now have the proposal that S can’t see that b is P without seeing b. Obviously, not all uses of the ‘sees that’ construction work like this; one doesn’t see that the cup is missing by seeing the cup. This points to the need for a contrast between primary and secondary epistemic seeing, that is, a contrast between ‘the cases where we see that b is P by seeing b itself, and the cases where we see that b is P without seeing b’ (Dretske 1969: 79-80).⁴¹ Primary epistemic seeing is the more fundamental or basic form of epistemic seeing since without it there would be no secondary epistemic seeing either, and no possibility of coming to know that b is P by seeing that b is P. For this reason I’m going to concentrate on the following question: what are the enabling conditions for primary epistemic seeing or for primary epistemic perceiving?

Suppose, then, that I see that the cup is chipped by seeing the cup. If there are enabling conditions for seeing the cup then they are also going to come out as enabling for seeing that the cup is chipped. But cups are objects, and seeing the cup is therefore an example of (visual) object-perception. Indeed, cups are not just objects but specifically material objects, in other words, bounded, three-dimensional space-occupiers. So as long as we are thinking of a case of primary epistemic seeing, the enabling conditions for seeing that the cup is chipped will include any background necessary conditions for seeing material objects. More generally,

given that seeing a material object is not the only way of perceiving one, we now have the proposal that one way, though not the only way, of figuring out what makes primary epistemic perceiving possible is to figure out what makes the perception of material objects possible.

This proposal is along the right lines, but it needs to be qualified in the following respects: to begin with, while it is true that some epistemic seeing involves the seeing of material objects, it is false that in every case in which one sees that b is P by seeing b itself b itself is a material object. At a baseball game one sees that the shadow of a low-flying plane is moving rapidly across the stadium, and one sees that the shadow is moving rapidly across the stadium by seeing the shadow. Yet shadows aren't material objects, even if they are objects in some looser sense of 'object'. So not every case of object-perception is a case in which what is perceived is a material object.

A further complication is that there are plenty of cases in which one sees that b is P by seeing b itself but in which b itself isn't an object at all, not even an object in a loose sense of 'object'. For example, b might be an event; one sees that a game of baseball is in progress by seeing the game but games are temporally extended events rather than objects. What makes it possible to see that b is P by seeing b itself will therefore include what makes it possible to see b itself for many different types of 'b'. This could make things very complicated but it is important at this stage to keep things as simple as possible. That is why it makes sense to concentrate, at least to begin with, on the simplest possible case of primary epistemic seeing, the case in which one sees that b is P by seeing b itself, and in which b itself is a material object. This isn't just a case of primary epistemic seeing but a case of what I'm going to call basic primary epistemic seeing. Seeing that the cup is chipped is a case of this kind. Later, I will consider the consequence of lifting the restriction to material objects as well as the consequences of lifting the restriction to visual perception.

We are now looking for an account of what makes basic epistemic seeing possible, and the present suggestion is that the enabling conditions for seeing that b is P will include the enabling conditions for seeing b itself. The enabling conditions for seeing b itself will be type A enabling conditions for seeing that b is P, that is, conditions which must be fulfilled for primary epistemic seeing to occur. What are the enabling conditions for seeing b itself in the case in which b itself is a material object? Some of these conditions have to do with the physical environment. For example, it must be light enough for b to be seen and there mustn't be anything blocking the perceiver's view of b. Others have to do with the workings of the perceiver's cognitive apparatus. Unless one's eyes and brain are functioning properly, one wouldn't be able to see b or, for that matter, anything else. So the enabling conditions for the perception of objects by sight include physiological conditions as well as environmental conditions.

Physiological and environmental enabling conditions are causally necessary conditions. From a naturalistic perspective, the enabling conditions for object-perception are always causally necessary conditions, and it is reasonable to suppose that there are many such conditions. It also seems a reasonable assumption that causally necessary conditions can only be discovered empirically. The implication is that the project of explaining what makes it possible to see a material object, and therefore the project of uncovering enabling conditions of basic primary epistemic seeing, can't be completed by armchair philosophy. On the contrary, it now appears that it is the business of empirical science to reveal the background necessary conditions for seeing that b is P by seeing b itself.

On this account, there is nothing wrong with the project of explaining what makes primary epistemic seeing possible as long as this project is conceived of naturalistically. But once it is conceived of in this way, it's no longer obvious that philosophy has much to contribute to it. In a way, therefore, this can be seen as a vindication of a kind of explanatory

minimalism; the thought is that philosophical explanation comes to an end at Level 2, and that any remaining questions about the means by which it is possible for us to come to know things about the external world are scientific questions. Philosophy can get one as far as the idea that a Level 3 explanation of what makes it possible to see that b is P will need to incorporate an explanation of what makes it possible to see b but it can't explain what makes it possible to see b.

This alliance between naturalism and minimalism shows that a defensible anti-minimalism will need to do more than insist on the necessity or possibility of explaining what makes perceptual knowledge possible. I'm taking it that an anti-minimalist is someone who thinks that distinctively philosophical explanations of the possibility of perceptual knowledge are necessary or at least possible. As far as basic epistemic seeing is concerned, therefore, the issue isn't whether there are questions about its causal enabling conditions that can only be answered empirically. That is not in dispute. The issue is whether, as the anti-minimalist insists, there are other questions about what makes this kind of seeing possible that can, or can only, be answered by means of a priori philosophical reflection. Minimalism is only in trouble if there are such questions.

Suppose, then, that we say that enabling conditions that can only be established by some form of a priori reflection are strongly a priori conditions. Weakly a priori conditions are ones that can be established without empirical investigation. In other words, it isn't written into the very idea of a weakly a priori enabling condition that such conditions can't also be established by empirical means. In these terms, minimalism can be understood as denying that there are any strongly or weakly a priori enabling conditions for this kind of knowledge. That is why, on the assumption that what is distinctive of philosophical explanation is that it is non-empirical, minimalism thinks that all such explanation comes to an end at Level 2. If this is right then there is no need for anti-minimalists to demonstrate that

there are strongly a priori enabling conditions for perceptual knowledge. In order to undermine minimalism they only need to make it plausible that there are enabling conditions that can be established without any empirical investigation.

The particular version of moderate anti-minimalism that I want to explore is Kantian in inspiration. To get a flavour of it, let's go back to the example of seeing that a particular cup is chipped by seeing the cup itself. Usually when one sees a cup one doesn't just see the cup. Typically, the cup is one among a range of things that one also sees, and seeing the cup involves being able to differentiate or distinguish it from these other things. If there are other cups in one's field of vision, one must be able to differentiate one cup from another. If one is holding the cup, one must be able to differentiate it from one's hand, and so on. The required differentiation is visual, and anything that is a background necessary condition for visual differentiation is also going to come out as background necessary condition for the perception of objects by sight. So if there are enabling conditions for the visual differentiation of objects that can be established by armchair reflection, without any empirical investigation, then this would count against what I have been calling explanatory minimalism.

In arguing in this way, one is not committed to thinking that it isn't possible to see an object without differentiating it from its immediate surroundings. Imagine seeing an object b which is in contact with another similar object c. When one sees b one fails to see it as distinct from c; in this sense one fails to differentiate b from its surroundings but this leaves open the possibility that what one sees in this case is b. At the very least, further work needs to be done to make it plausible that a failure visually to differentiate b from c amounts to a failure to see b.⁴² It's clear, however, that paradigmatic cases of object-perception are ones in which what one perceives is differentiated from its surroundings. For example, when one is pouring coffee into a cup it is perceptually manifest to one where the cup ends and the rest of the world begins; that is why the coffee ends up in the cup. So does a priori philosophy have

anything useful to say about the background conditions under which this kind of differentiating object-perception is possible?

Kant's proposal is that the perception of space is a background necessary condition for visual object-perception. The 'perception of space' is a cognitive capacity, the capacity to perceive spatial properties such as shape and location. Kant's idea is that possession of this cognitive capacity is an enabling condition the perception of objects by sight because it is an enabling condition for visual differentiation. If this claim is correct, and if seeing that b is P involves visually differentiating b from its surroundings, we can conclude that the perception of space is a type A cognitive enabling condition for seeing that b is P. By making it possible for one to see b, the perception of space makes it possible for one to see that b is P. Kant also thinks that the link between the perception of space and visual differentiation can be established non-empirically. That is why, on his view, the perception of space is not just an enabling condition but an a priori enabling condition for differentiating visual perception. The identification of this condition therefore serves as an example of what philosophy, as distinct from empirical science, can achieve in this area.

Why should one think that the perception of space is a background necessary condition for differentiating visual perception? One thought is that in order to see an object b as differentiated from another object c one must see b and c as being in different places.⁴³ The perception of place, and therefore of space, serves as the means by which b is differentiated from other things in its environment. But one couldn't see where b is unless one has the capacity to perceive spatially. So possession of this capacity comes out as a background necessary condition for differentiating perception as long as it is also true that the perception of place is not just a means but the only means of differentiating b from its environment. If it differentiating perception needn't be a form of spatial perception, then one would need to find

other reasons for thinking that the perception of space is a cognitive enabling condition for seeing that b is P.

Is Kant's conception of the role of spatial perception in primary epistemic seeing defensible? I will have much more to say about this in chapter 3. For example, one might wonder whether the perception of place can really be as important for visual differentiation as the thesis suggests. When I pour myself a cup of coffee, I certainly see the cup as distinct from my hand, yet it is implausible that I see them as distinct things by seeing them in different places. I see them as distinct because they look distinct, and looking distinct in this sense seems to have little to do with the perception of a difference in location. When senses other than sight are considered, the problem is even more acute. One can hear two people arguing in the next room as distinct from each other just on the basis of their voices. The perception of a difference in location has little to do with it, since I might fail to hear them as being in different places in the next room.

These observations suggest that Kant's thesis faces some formidable challenges. It is worth emphasizing, however, that it could be true that one can't see an object without perceiving any of its spatial properties even if it is false that the perception of space is the key to visual differentiation. Where the object is a material object, one might think that one couldn't see it at all without seeing it as shaped, located or extended in space. So the perception of space could still be a background enabling condition for the perception of b itself regardless of whether it is the key to differentiating b from other objects. If this is right, there would again be a question about the role of space in the perception of objects by non-visual modes of perception, and about the extent to which it is an empirical question whether one can see a material object without seeing any of its spatial properties. For the moment, we can just bracket these questions. Right now all I am trying to do is to give an illustration of what might count as a type A enabling condition for primary epistemic seeing, on the

assumption that primary epistemic seeing is a means of coming to know things about the world around us.

Are there any other enabling conditions for primary epistemic seeing which have not yet been mentioned? When I introduced the notion of epistemic seeing I said that this form of seeing is conceptual; for example, in order to see that the cup is chipped one must have the concepts cup and chipped. These are examples of empirical concepts, that is, concepts that have their source in experience or that can be derived from experience. Hence, if there are background necessary conditions for the possession or acquisition of empirical concepts, one would be entitled to regard them as enabling conditions for seeing that the cup is chipped. More generally, given that concepts are necessary for epistemic seeing, one would expect a Level 3 explanation of what makes it possible for one to see that b is P to incorporate an account of what makes it possible for one to have concepts like b and P.

What are the background conditions under which the possession or acquisition of empirical concepts is possible? Again, the naturalist or minimalist claims that this is best understood as an empirical question about causal enabling conditions. Such conditions might be physiological or biological in nature, and it is not for armchair philosophy to tell us what they are. In contrast, the anti-minimalist denies that Level 3 questions about the concepts that figure in epistemic seeing are necessarily scientific or empirical, even though some of them undoubtedly are. As far as the anti-minimalist is concerned, there are enabling conditions for the possession and acquisition of empirical concepts that can be discovered by armchair philosophical reflection and that are therefore at least weakly if not also strongly a priori. Again, the intended upshot is that philosophical explanation needn't, and perhaps shouldn't, come to an end at Level 2.

One distinctively philosophical Level 3 proposal is that in order to have concepts like cup and chipped one must have lots of other concepts. Possession of a network of interrelated

concepts is a background necessary condition for the possession individual concepts like cup and chipped. While it's not uncontroversial whether this holistic constraint on concept-possession is correct, the anti-minimalist's idea is that its correctness or otherwise can't be settled empirically; only a philosophical theory of concepts can do that. And the same goes for another putative enabling condition for concept-possession. On a linguistic conception of concepts, only a creature with a language can have concepts. Concepts must be exercisable in judgements, and if the capacity to judge depends on language then so do the concepts that figure in one's judgements. If this is right, mastery of a language is required in order to see that the cup is chipped.

I'm not going to have very much to say about these alleged holistic and linguistic constraints on concept-possession. Instead, I'm going to concentrate on Kant's conception of what makes it possible for one to acquire and possess concepts like cup and chipped. It's worth noticing that the holistic and linguistic constraints aren't just constraints on empirical concepts; for example, if it is in the nature of concepts to presuppose mastery of a language, then this is presumably going to be true of all concepts, empirical or otherwise. The Kantian proposal which I will be discussing makes more of the distinction between empirical and non-empirical concepts. In so far as empirical concepts are ones that have their source in experience or that can be derived from experience, non-empirical or a priori concepts are ones that don't have their source in experience or that can't be derived from experience. Kant's inventory of a priori concepts includes the so-called 'categories' or 'pure concepts of understanding'. His thesis is that the categorial thinking, thinking by means of concepts like substance, unity, plurality and causality, is an a priori enabling condition for the possession and acquisition of empirical concepts like cup and chipped.⁴⁴ This kind of thinking makes it possible for one to have and acquire such concepts, and therefore also makes epistemic seeing

possible. Without the categories in the background, one couldn't have or acquire concepts like cup and chipped, and without these concepts one couldn't see that the cup is chipped.

As we will see in a chapter 4, this Kantian explanation of the possibility of empirical concepts and, by extension, of the possibility of epistemic seeing faces a range of challenges and objections that are no less formidable than those facing Kant's Level 3 explanation of the possibility of object-perception. The main worry isn't that Kant is attempting to establish by means of a priori reflection claims that ought to be established empirically. The worry is rather that it's far from obvious that Kant's claims about the link between empirical concepts and the categorial thinking are actually correct. In the end, I will argue that only very watered down versions of these claims have any chance of being defensible. They are nevertheless worth discussing both because of their historical interest and because they provide an excellent illustration of what happens when anti-minimalism is taken too far. Kant's thinking about these matters is underpinned by an extreme form of anti-minimalism, and reflecting on the failings of this approach will make the virtues of the moderate anti-minimalism that I want to defend very much clearer than they would otherwise be.

So much for type A enabling conditions of epistemic seeing. What is there to say about its type B enabling conditions? An account of these conditions would need to explain what makes it possible for seeing that something is the case to be a way of coming to know that it is the case. The minimalist says that all that can be done in this connection is to point out that 'S sees that b is P' entails 'S knows that b is P'. There is nothing that 'makes it possible' for 'S sees that b is P' to entail 'S knows that b is P', so it doesn't make sense to look for type B enabling conditions for epistemic seeing. If explaining how epistemic seeing can be a source of knowledge is a matter of explaining how 'S sees that b is P' can entail 'S knows that b is P', then minimalism denies that any such an explanation is either necessary or possible.

This form of minimalism is related to a more general form of minimalism in epistemology. In Knowledge and its Limits, for example, Williamson argues that the project of trying to fix non-circular necessary and sufficient conditions for propositional knowledge is doomed, and that the concept knows cannot be analysed into more basic concepts.⁴⁵ Instead he gives a modest positive account of this concept according to which ‘if one knows that A, then there is a specific way in which one knows; one can see or remember.... that A’ (2000: 34). ‘Sees that...’, ‘remembers that...’ and ‘knows that...’ are all examples of factive mental state operators (FMSOs). If ϕ is an FMSO, ‘S ϕ s that A’ entails ‘S grasps the proposition that A’, and the inference from ‘S ϕ s that A’ to ‘A’ is deductively valid. In addition, ϕ is semantically unanalysable. In these terms, ‘the proposal is that if ϕ is any FMSO, then ‘S ϕ s that A’ entails ‘S knows that A. If you see that it is raining, then you know that it is raining. If you remember that it was raining, then you know that it was raining’ (2000: 37). The implication of this account of the sense in which ‘seeing that A is a way of knowing that A’ (2000: 38) is that the concepts knows can effectively be characterized as the determinable of which such specific ways of knowing are the determinations.

Part of what makes this a form of epistemological minimalism is that it offers no analysis or explanation of the link between seeing that A and knowing that A. In contrast, anti-minimalism does seek to explain this link. It points out, for example, that seeing that A isn’t a way of knowing that A in the sense in which remembering that A is a way of knowing that A. Seeing that A is a way of coming to know that A, a way of acquiring this knowledge, whereas remembering that A is a way of retaining the knowledge that A. To explain the link between seeing that A and knowing that A would therefore be to explain how seeing that A can be a way of coming to know that A. This is where type B enabling conditions come into the picture. The anti-minimalist’s proposal is that we understand how epistemic or ‘factive’ seeing can be a source of knowledge by identifying the background conditions that must be

met for seeing that something is the case to be a way of coming to know that it is the case. Once we have identified these conditions, we might also find ourselves being able to say something illuminating about why 'S sees that A' entails 'S knows that A'.

What are the background necessary conditions for seeing that A to be a way of coming to know that A? Are these conditions any different from the enabling conditions for seeing that A? We can start to make some progress with these questions by drawing on Dretske's account of epistemic seeing in Seeing and Knowing. Dretske describes himself as aiming to provide 'an analytic description of those states of affairs which are described by statements of the form "S sees that b is P" in so far as they tell us how S knows that b is P' (1969: 81). The account proceeds by specifying necessary and sufficient conditions for epistemic seeing that are adequate to the epistemic implications of this mode of perception. Thus, for S to see that b is P it is necessary and sufficient that (i) b is P, (ii) S sees b, (iii) the conditions under which S sees b are such that b would not look the way it looks now to S unless it was P, and (iv) S, believing the conditions are as described in (iii), takes b to be P. When conditions (i) through (iv) are fulfilled, S has in the way that b looks to him a conclusive reason for believing that b is P, and 'it is the conclusiveness of this reason which supports the entailment between "S sees that b is P" and "S knows that b is P"' (1969: 124).

This is in a way an anti-minimalist account of epistemic seeing but work needs to be done to transform it into a full-blown anti-minimalist account. It is anti-minimalist to the extent that it tries to analyse the claim that S sees that b is P and thereby to explain the entailment between 'S sees that b is P' and 'S knows that b is P'. It also employs the notion of a background enabling condition since these are the conditions that are referred in conditions (iii) and (iv). On the other hand, Dretske makes it clear that his enabling conditions are causal enabling conditions.⁴⁶ To get to the idea of a priori enabling conditions for the acquisition of knowledge by visual means, we don't need to take a stand on the issue of whether Dretske's

conditions are genuinely sufficient. For as long as it's plausible that S must see b in order to know by visual means that b is P, we can proceed to identify a priori enabling conditions for epistemic seeing on this basis.

Suppose that b itself is an object. In that case, coming to know that b is P by seeing that b is P requires a capacity for object-perception. So if there are a priori enabling conditions of object-perception they will also be a priori enabling conditions for knowing that b is P by seeing that b is P. We have already briefly examined the proposal that a capacity for spatial perception is an a priori enabling condition for the perception of objects, and therefore a type A enabling condition for epistemic seeing. But anything that is an enabling condition for the perception of objects is also going to be a type B enabling condition for epistemic seeing, given that object-perception is involved in the acquisition of knowledge by primary epistemic seeing. What we now have, therefore, is the possibility that the perception of space is both a type A and a type B enabling condition for epistemic seeing.

This overlap between type A and type B enabling conditions should come as no surprise. Since it is in the nature of epistemic seeing that it has epistemic implications, one would expect the background conditions under which it is possible for epistemic seeing to occur to be closely related to, if not identical with, the conditions under it is possible for epistemic seeing to be a source of knowledge. In other words, what makes it possible for there to be such a thing as epistemic seeing can't be sharply distinguished from what makes it possible for there to be such a thing as knowing that something is the case by seeing that it is the case. There remains a notional difference between type A and type B enabling conditions, but a notional difference, a distinction at the level of sense, needn't amount to a real difference, a distinction at the level of reference.

We are now in a position to consider the extent to which Kant's claims about the enabling conditions of perceptual knowledge vindicate anti-minimalism. I have said that an

anti-minimalist is someone who thinks that distinctively philosophical Level 3 explanations of the possibility of perceptual knowledge are necessary or at least possible. One way of showing that such explanations are possible would be to produce one. This is what Kant does or at least purports to do. If the perception of space and the categories are enabling conditions for epistemic perception that can be established non-empirically, then minimalism is wrong to claim that distinctively philosophical Level 3 explanations of what makes perceptual knowledge possible can't be given. For minimalism to be in the running, it would have to be the case that Kant fails to identify Level 3 conditions for perceptual knowledge. This could either be because spatial perception and the categories aren't background necessary conditions for the acquisition of perceptual knowledge or because they aren't background necessary conditions for the acquisition of perceptual knowledge that can be established non-empirically. I will consider these possibilities in detail in later chapters, though it's worth pointing out that the failure of Kant's Level 3 explanations wouldn't necessarily mean that such explanations can't be given.

Suppose, then, that we are persuaded on the basis of Kant's discussion that in principle Level 3 explanations of the possibility of perceptual knowledge can be given. This would be a problem for minimalism, but where does it leave the debate between moderate and extreme anti-minimalism? Extreme anti-minimalists think that in the absence of a Level 3 explanation we can't reasonably claim to have explained how perceptual knowledge is possible, and that this is the sense in which Level 2 explanations aren't good enough. In contrast, moderate anti-minimalists insist that Level 3 explanations are possible while denying they are necessary. Their point is that once we have reached Level 2 and identified epistemic seeing as a means of coming to know things about the world around us we have already done everything that needs to be done to explain how perceptual knowledge is possible; we could go further but we don't need to.

How is this dispute to be resolved? Consider this analogy: I ask how it is possible to get from London to Paris in less than three hours and the answer I get is that it's possible to do this by catching the Eurostar. Should I be satisfied by this answer? An extreme anti-minimalist in this context is someone who thinks that I shouldn't be satisfied and that more can and needs to be done to answer the how-possible question. Catching the Eurostar is a means of getting from London to Paris, but what makes it possible to get from London to Paris by train? This is the further what-makes-it-possible question to which extreme anti-minimalism demands an answer. The answer, or at least an answer, to this question is that the existence of the Channel Tunnel is what makes it possible to get from London to Paris by train. In the absence of the Channel Tunnel, or some such link between England and France, going by train would not be a means of reaching Paris from London, so the existence of such a link is an enabling condition for getting from London to Paris by train. According to the extreme anti-minimalist, I haven't fully understood how it is possible to reach Paris from London in less than three hours unless I recognize this enabling condition, just as I haven't fully understood how it is possible to arrive at synthetic a priori geometrical knowledge by constructing figures in pure intuition unless I recognize the ideality of space.

This is what moderate anti-minimalists find implausible. Their thought is that there is no obvious sense in which a failure to say anything about the background necessary conditions for crossing the English Channel by train constitutes a failure to give a 'complete' answer to the how-possible question. There are lots of factors that make it possible to reach Paris from London by train – the existence of the Tunnel, the existence of trains, and so on – and some of these factors can be established without empirical investigation while others can only be empirically. Presumably I can know a priori that the existence of trains is necessary for getting anywhere by train but I can't know a priori that the existence of a cross-channel tunnel is an enabling for getting from London to Paris by train. Yet there is no need to go into

any of this if all one wants to know is how it is possible to get from London to Paris in less than three hours. As far as answering the how-possible question is concerned the what-makes-it-possible question is an optional extra even though it's one to which answers can be given if someone insists on asking it. If I know that it's possible to get from London to Paris in less than three hours by catching the Eurostar then I know how it's possible to get from London to Paris in less than three hours.

Moderate anti-minimalism's take on (HP_{pk}) is similar to its take on the Eurostar. Again the idea is that once I understand that it's possible to know that the cup is chipped by seeing that it is chipped I understand how this particular piece of knowledge is possible. Just as nothing needs to be said about the existence of the Channel Tunnel in order to explain how it's possible to get from London in Paris in less than three hours, so nothing needs to be said about the enabling conditions of epistemic perception in order to explain how perceptual knowledge is possible. The questions that get addressed at Level 3 of the multiple levels response are perfectly legitimate and interesting but it would be wrong to claim that Level 2 explanations are, in any straightforward sense, incomplete.

None of this amounts to a knockdown argument against extreme anti-minimalism but it does put considerable pressure on this defenders of this position to explain why we should be any less satisfied with a Level 2 response to (HP_{pk}) than with a Level 2 response to many other how-possible questions, including the one about getting to Paris from London. While it's possible to get oneself into a frame of mind in which questions about enabling conditions can seem pressing it's also quite easy to see why someone might think that this question has already been answered at Level 2. If, as I have claimed, (HP_{pk}) is an obstacle-dependent question the key to answering it is to remove the various obstacles that have been alleged to stand in the way of the acquisition of perceptual knowledge. Since this is what happens at Level 2 there is no obvious need to go any further. From this standpoint, the moderate anti-

minimalist's insistence that (HP_{pk}) can but needn't evolve into a what-makes-it-possible question appears to be entirely justified. The onus is on extreme anti-minimalism to make it plausible that we haven't explained how perceptual knowledge is possible unless we have explained what makes it possible, and we have so far failed to find any decisive arguments in favour of this approach.

These considerations also have a bearing on the worry that Level 3 explanations are fundamentally no different from Level 2 explanations. The worry was that the denial of a putative enabling condition is always an obstacle to knowledge, and that what goes on at Level 3 of a multi-levels response to a how-possible question is therefore just as much an exercise in obstacle-removal as what goes on at Level 2. My initial response to this worry was to argue that the point of identifying enabling conditions C for the acquisition of knowledge of kind K by means M needn't be to address any intuitive obstacle to the acquisition of K by M. If C is not fulfilled then that becomes an obstacle to the acquisition of K by means of M but it doesn't follow that Level 3 explanations are Level 2 explanations by another name. The Kantian account role of spatial perception and categorial thinking helps to make this point. There is no intuitive obstacle to epistemic seeing, to knowing that the cup is chipped by seeing that it is chipped, that is overcome or dissipated by the observation that the perception of space and categorial thinking are background necessary conditions for epistemic seeing. Without these cognitive capacities one wouldn't be able to see that the cup is chipped but what Kant is trying to do by talking about what makes perceptual knowledge possible is not to show that perceptual knowledge is not impossible; what he is after is a better understanding of the cognitive foundations of this kind of knowledge.

The lesson, once again, is that the only version of anti-minimalism to which we should be willing to commit ourselves in relation to (HP_{pk}) is moderate anti-minimalism. So the issue is not whether it's possible to answer (HP_{pk}) without identifying a priori enabling conditions

for perceptual knowledge but whether perceptual knowledge has any enabling conditions that can be established non-empirically. What we need to consider, therefore, is whether spatial perception and categorial thinking are enabling conditions for the acquisition of perceptual knowledge and, if so, whether they are a priori enabling conditions. As we will see in chapters 3 and 4, the answer to both of these questions is ‘yes’. It isn’t true, therefore, that armchair philosophy can’t tell us anything about what makes perceptual knowledge possible.

But before looking in more detail at the role of space and categorial thinking in the acquisition of perceptual knowledge there is another methodological question issue that needs to be addressed. I have written at length about how-possible questions in epistemology but have so far said nothing about so-called ‘transcendental arguments’, that is, arguments which specify necessary conditions of the possibility of thought, experience or knowledge. Yet it has often been suggested transcendental arguments are the best way of answering epistemological how-possible questions. So the next chapter is about transcendental arguments. We need to consider whether such arguments are any good, whether they have any bearing on questions like (HP_{pk}) and, if so, how they relate to the multi-levels approach to this question that I have been recommending.

¹ This is William Dray's example. See Dray 1957: 158. Dray got it from a magazine report on a baseball game in Canada. My account of how-possible questions owes a lot to Dray's seminal discussion, but he focuses on the philosophy of history rather than epistemology.

² As Dray puts it, 'explanation is called for because what happened seemed impossible in the circumstances' (1957: 160). Dray's conception of how-possible questions has also been taken up by Nozick and Stroud. See Nozick 1981: 8-11, and Stroud 1984: 144.

³ These are Nozick's examples. See Nozick 1981: 8-9.

⁴ This is what Nozick is getting at in the following passage: 'the form of [how-possible] questions is: how is one thing possible, given (or supposing) certain other things? Some statements r_1, \dots, r_n are assumed or accepted or taken for granted, and there is a tension between these statements and another statement p ; they appear to exclude p 's holding true. Let us term the r_1 apparent excluders (of p). Since the statement p also is accepted, we face the question of how p is possible, given its apparent excluders' (1981: 9). Nozick's 'excluders' are my 'obstacles'. Although, as Nozick points out, logical incompatibility is the strongest mode of exclusion, it isn't the kind of incompatibility that is at issue in all how-possible questions. Dray's example clearly illustrates this point. In general, I take obstacle-dependence to a pragmatic rather than a semantic phenomenon.

⁵ Denying the existence of a putative obstacle doesn't always do the trick. Sometimes, it just creates another obstacle. For example, suppose we deny that all actions are causally determined. Does that mean that actions occur randomly? If so, that is just as problematic for attributions of moral responsibility as the truth of determinism.

⁶ If denying that an obstacle is genuine is the same as denying that it exists then the distinction between the two obstacle-removing strategies will have to be explained slightly differently. In practice, however, the distinction between dissipating and overcoming an obstacle isn't difficult to understand. It corresponds to Nozick's distinction between two ways of dealing

with an apparent incompatibility between a statement p and its apparent excluders. First, ‘one of the apparent excluders can be denied, or there can be a denial of their conjunction all together’ (1981: 10). This amounts to what I am calling ‘obstacle-dissipation’. Second, ‘each of the apparent excluders can continue to be maintained, while their apparent incompatibility with p is removed’ (ibid). To do this is to overcome a supposed obstacle.

⁷ The restriction to human knowledge should come as no surprise if we understand epistemology as ‘the philosophical study of human knowledge’ and therefore as seeking to understand ‘what human knowledge is and how it comes to be’ (Stroud 2000c: 99). The focus on human knowledge also plays an important part in some accounts of scepticism. In his ‘Defence of Skepticism’, for example, Unger argues specifically for the thesis that ‘every human being knows, at best, hardly anything to be so’ (1971: 198).

⁸ Michael Williams objects that knowledge of the external world isn’t a ‘kind of knowledge, which we might assess or explain as a whole’, and that in this sense there is ‘no such thing as knowledge of the external world’ (1996: xii). I don’t find Williams’ arguments for this view convincing but won’t discuss them here. If Williams is right, however, then (HP_{ew}) is not a good question.

⁹ As Austin remarks, the question ‘how do you know?’ is one to which (some) claims to knowledge are ‘directly exposed’ (1979: 77). See Hampshire 1979 for more on this theme, and for a defence of the idea that there are some knowledge claims in relation to which the how-question would be ‘at least absurd, and perhaps unintelligible’ (1979: 282).

¹⁰ Dray has some useful things to say about the differences between simple how-questions and how-possible questions. See Dray 1957: 166.

¹¹ Cf. Stroud 2000a: 3. The importance of recognizing that there are different ‘pathways’ to knowledge has also been emphasized by Alvin Goldman. See the preface to Goldman 2002.

¹² This is something that Stroud frequently emphasizes, as in this characteristic passage: ‘speaking “anthropologically”, we can say that human beings gain knowledge of the world through sense-perception. And when we look more closely into how sense-perception works, and what exactly we perceive, it can become difficult to see how perceptual knowledge of the world is possible’ (2000d: 123).

¹³ Stroud mentions ‘ways of coming to know something’ (2000a: 3) in one of his accounts of knowledge and scepticism. Peacocke also employs the notion of a ‘way of coming to know that something is the case’ in his account of a priori knowledge. See Peacocke 2000. Ways of coming to know are what Goldman calls ‘pathways to knowledge’. See n. 11 above.

¹⁴ The underlying question is this: if knowledge of the external world is attainable by ‘a wide variety of sometimes independent and sometimes interconnected pathways’ (Goldman 2002: xi), and if perception is just one of these pathways, why should we think that knowledge of the external world wouldn’t be possible just because perceptual knowledge isn’t possible? Why couldn’t we rely on other sources to provide us with knowledge of the external world? The answer to this question depends, in part, on why perceptual knowledge is thought to be impossible. Suppose that it is thought to be impossible because we are incapable of perceiving external objects. But if we can’t perceive external objects then, on some views of thought, we can’t think about them either, and if we can’t think about them then we can’t have knowledge of them. In this way, an alleged obstacle to perceptual knowledge transforms itself into an obstacle to our having any knowledge of the external world, even if perception isn’t the only pathway to this kind of knowledge. This would be one way of justifying Stroud’s insistence on the centrality of perceptual knowledge.

¹⁵ Nozick appears to have something similar in mind when he remarks that ‘the task of explaining how p is possible is not exhausted by the rearguard action of meeting arguments from its apparent excluders’ (1981: 11).

¹⁶ What-makes-possible explanations are also the focus of Peacocke's moderate rationalist account of a priori knowledge. There is a priori knowledge only if there are a priori ways of coming to know a proposition, and rationalism is the view that a priori ways of coming to know are grounded in the understanding. This leaves the rationalist with an explanatory task: 'the task is to say what it is about understanding that makes a priori knowledge possible' (Peacocke 2000: 257). There is more on all of this below, in chapter 6.

¹⁷ All references in this form are to Kant 1932.

¹⁸ Kant defends the view that 'mathematical knowledge is the knowledge gained by reason from the construction of concepts' (A713/ B741) in the section of the Critique of Pure Reason called 'The Discipline of Pure Reason in its Dogmatic Employment'.

¹⁹ See, especially, A141/ B180.

²⁰ On the non-negotiability of perceptual knowledge see Moore 1922.

²¹ See McDowell 1998: 444 for the notion of an explanation-seeking question.

²² Enabling conditions figure in Dretske's account of 'epistemic seeing' (seeing that b is P) and in Searle's account of intentionality. Both writers stress that enabling conditions, as they conceive of them, are background conditions. See Dretske 1969: 82-3 and Searle 1983: 141-59. But both Dretske and Searle take it that enabling conditions are causally necessary conditions. See, for example, Searle 1983: 157-8. This isn't built into my account of enabling conditions, which is not to deny that some enabling conditions are causal.

²³ This would presumably be Quine's view. See Quine 1969. What is effectively Quine's approach to the study of enabling conditions is helpfully summarized in the following terms by Stroud: 'It is not that we have any difficulty understanding the general idea of a study of the conditions of human knowledge – an investigation of those characteristics of human organisms that make it possible for them to come to know things about what is going on around them. But the best way to carry out such a study would seem to be by observing

human beings and trying to understand how they work. It would be an empirical investigation – which is not to say that it would be easy to carry out’ (1984: 160).

²⁴ This assumes, of course, that distinctively philosophical explanations are armchair explanations, and that philosophy is therefore discontinuous with empirical science.

²⁵ Thanks to Bill Brewer and Tim Williamson for raising this question.

²⁶ Peacocke 2000 is helpful on this point.

²⁷ This claim is controversial but let’s not worry about that here.

²⁸ This point has been emphasized by, among others, Dretske, Williamson and Stroud. See Dretske 1969: 124, Williamson 2000: 37, and Stroud 2004: 167.

²⁹ This is the gist of what Stroud calls the ‘most straightforward answer’ to (HP_{pk}) and, by implication, to (U). This answer says that ‘one sees that there is a table in the room, not that one infers that there is a table in the room from something else. And to see that p is to know that p. Whoever sees that p thereby knows that p. Whoever sees that there is a table in the room knows that there are external things’ (2004: 167).

³⁰ Dretske gives a detailed account of epistemic seeing in his 1969: 78-139.

³¹ Williamson also makes this point. To quote his example, ‘a normal observer in normal conditions who has no concept of chess can see a situation in which Olga is playing chess, by looking in the right direction, but cannot see that Olga is playing chess, because he does not know what he sees to be a situation in which Olga is playing chess’ (2000: 38).

³² See Dretske 2000a for an account of simple seeing.

³³ This is McDowell’s example. See his 1998b: 390 n. 37. He implies that in such cases the subject is in a position to know that the cup is chipped without actually knowing that the cup is chipped. But if one can see that the cup is chipped without actually knowing that it is chipped then it is hard to see how seeing that p can be a way of knowing that p.

³⁴ This is Williamson's view. He claims that cases such as the one I have described 'put more pressure on the link between knowing and believing or having justification than they do on the link between perceiving or remembering and knowing' (2000: 38).

³⁵ This is one of Dretske's conditions on epistemic seeing. A subject S sees that b is P in a primary epistemic way only if 'the conditions under which S sees_n b are such that b would not look, L, the way it looks now to S unless it was P (1969: 82). The subscript makes it clear that non-epistemic or simple seeing is a component of epistemic seeing. See below, 1.4, for more on the distinction between primary and secondary epistemic seeing.

³⁶ See Austin 1979: 78 for an account of the pointed use of 'how' questions.

³⁷ This line of argument is set out in Stroud 1984 and 2000d. Take a case in which any one of us would think that there is a fire in the fireplace right before us, and that we know that it is there because we see that it is there. The introduction of 'alternative, uneliminated possibilities' leads us to conclude that 'whatever we see to be so in that case, we do not simply see that there is a fire there' (Stroud 2000d: 131).

³⁸ Cf. McDowell 1998a: 238-9.

³⁹ Moore argues in this way in some of his writings on scepticism. See, for example, 1953: 121-2.

⁴⁰ This is what Baldwin calls Moore's argument. See Baldwin 1990: 269-74 for a critical discussion of this argument. Also relevant is Pryor 2000.

⁴¹ As Dretske points out, this account of the distinction between primary and secondary epistemic seeing isn't quite accurate because there are cases in which one sees that b is P 'in virtue of the way other objects look or behave when b is P' (1969: 153). These are cases of secondary epistemic seeing even if they are ones in which one sees b itself. The reason is that seeing b itself is incidental. For example, when I insert a toothpick into the middle of a cake I see the cake but it is the way the toothpick looks that enables me to see that the cake is done.

⁴² See Dretske 1969: 18-32 and chapter 3 below for further discussion of this issue.

⁴³ Allison defends something along these lines in his 1983: 83. There is much more on this issue in Warren 1998 and chapter 3 below.

⁴⁴ Longuenesse attributes something along these lines to Kant. See Longuenesse 1998.

⁴⁵ See Williamson 2000: 27-33.

⁴⁶ See Dretske 1969: 82-4.