

Knowing by Imagining

Timothy Williamson

1. Imagining is often contrasted with knowing. When you know nothing about something, you have to imagine it instead. Knowledge deals in facts, imagination in fictions. This paper sketches a way of thinking about the imagination on which that stereotypical contrast is utterly misleading. Far from being the opposite of knowing, imagining has the basic function of providing a means to knowledge — and not primarily to knowledge of the deep, elusive sort that we may hope to gain from great works of fiction, but knowledge of far more mundane, widespread matters of immediate practical relevance.

The assimilation of imagining and knowing is a card that can be played for either realist or idealist ends. The idealist wants to shift our conception of knowing towards our conception of imagining: somehow even knowing deals in fictions. By contrast, the realist wants to shift our conception of imagining towards our conception of knowing: somehow even imagining deals in facts. The spirit of this paper is firmly realist. It aims more to rethink imagination than to rethink knowledge. Consequently, it is *not*

appropriate to start by defining ‘imagination’. Any such definition would be premature in advance of inquiry. Without one, we can still make progress with that inquiry on the basis of our capacity to recognize clear cases of imagination. What unifies them (if anything), and therefore which less clear cases should be grouped with them, are to be identified in the course of inquiry, not prejudged from the beginning.

A more useful starting-point is to ask oneself why the elaborate capacity for imagining that normal humans possess should have arisen in our evolutionary history. Although this paper will not attempt any detailed evolutionary considerations, one feature of the view to be proposed is that it makes the evolutionary advantage of having a good imagination obvious. That contributes to the view’s explanatory power, and so to its abductive confirmation.

The reference to humans’ evolutionary history does not imply that the imagination is an exclusively human capacity. It is quite plausible that some non-human animals engage in at least primitive imaginative exercises, and thereby gain a similar advantage. In this paper, however, the examples will largely concern the human imagination, since we know so much more about it.

Of course, that we possess an imagination does not prove that it confers or once conferred some evolutionary advantage on either the individual or the species. In principle, the imagination might have arisen as an accidental by-product of something else that did confer an evolutionary advantage at one or other level. Sometimes, for structural reasons, the easiest way for evolution to develop the capacity to do something useful involves developing the capacity to do something useless too. This paper does not take for granted that the imagination confers an evolutionary advantage on either the

individual or the species. Rather, it proposes a speculative view of the imagination on which it *does* confer an evolutionary advantage on both the individual and the species, and that aspect of the view's explanatory power contributes to its confirmation.

There are many kinds of evolutionary advantage. Presumably, peacocks have unwieldily magnificent tails because peahens prefer peacocks so endowed. One could develop a similar hypothesis about the imagination. Other things being equal, the more imaginative you are, the better your seduction technique: if you dance, talk, and do other things more imaginatively, you are more exciting to be with. In *One Thousand and One Arabian Nights*, Scheherazade's skill at telling stories keeps her alive and eventually enables her to pass on her genes. On the view to be proposed here, however, a good imagination also confers much more straightforward and direct advantages which do not depend on whether other members of the species are attracted by works of your imagination.

2. Consider a small group of our distant human ancestors, travelling across a difficult landscape previously unknown to them. How might they find an imagination useful?

One obvious answer is that an imagination will alert them to various potential dangers and opportunities. They are about to enter a forest. They imagine wolves in the forest; warned of the danger, they keep a sharper look-out for signs of wolves. They imagine edible berries in the forest; alerted to the potential opportunity, they look about for bushes of the right kind. In both cases, their imagination enables them to prepare for practically relevant possibilities, helping them avoid dangers and take advantage of opportunities.

To serve that purpose well, the imagination must be both selective and reality-oriented. They could imagine the wolves bringing them food to eat, but doing so would be a waste of time, and a distraction from more practically relevant possibilities. An imagination that clutters up the mind with a bewildering plethora of wildly unlikely scenarios is almost as bad as no imagination at all. It is better to have an imagination that concentrates on fewer and more likely scenarios. One's imagination should not be completely independent of one's knowledge of what the world is like.

Such examples already suggest a distinction between two modes in which the imagination can operate: voluntary and involuntary. When you encounter a problem to which there is no obvious solution, you can turn your imagination to thinking up ways of solving it and to thinking through their consequences. Our little group may do that when they come to a deep river that they must cross. Such uses of the imagination are often voluntary. But sometimes we need our imaginations to work involuntarily too. If the group is absorbed in joking together as it enters the forest, it may be saved by an imagination that breaks into the stream of consciousness with reminders of dangers and opportunities.

Imagination resembles *attention* in having both voluntary and involuntary modes of operation, and for similar reasons. We need to be able to focus our attention voluntarily on something: for example, to set ourselves to watch a hole in case an animal comes out. But we also need our attention to be capable of involuntary switches: for example, to be drawn by a slight movement at the periphery of vision, perhaps a predator or prey, despite our previous intention of watching something else. At least metaphorically, one might regard imagination as a form of attention to possibilities.

Since members of the group have imaginations, they can be expected to use them, voluntarily or even involuntarily, in ways that serve no immediate practical purpose too, for instance by inventing and telling stories. In the long run, of course, such uses of the imagination may help the group bond. Moreover, even the most playful uses of the imagination may also help to exercise the capacity and keep it in good trim, ready for more practical applications, just as a cat practises its hunting skills by playing with a mouse once caught, letting it go and recapturing it.

3. So far, the imagination has been described primarily as *raising possibilities*, rather than assessing the truth-values of propositions. One might therefore be tempted to suppose that the proper role of the imagination belongs, in Hans Reichenbach's terminology, to the *context of discovery*, rather than the *context of justification*. On a simple-minded version of the distinction, one dreams up scientific theories in the context of discovery, but assesses them as true or false, or as probable or improbable on the evidence, in the context of justification. Rationality is essential in the context of justification, but not in the context of discovery. If you came up with the theory in the first place under the influence of drink or drugs, never mind, but you must sober up or come down from your high before you reach a conclusion as to its status on the evidence. Similarly, someone might think, the imagination has done its work once it has delivered enough proposition to consciousness, to be entertained there. On this view, the imagination plays no further role in the assessment of those propositions as true or false, or as probable or improbable on the evidence.

The considerations of the previous section already include signs that such a view is inadequate. The need for the imagination to be selective and reality-oriented in bringing material to consciousness suggests that it must be capable of some sort of rational responsiveness to evidence. But in many examples we can see the imagination playing a far more direct role in the assessment of truth-value.

Think of a hunter who finds his way obstructed by a mountain stream rushing between the rocks. He reaches the only place in the vicinity where jumping the stream might be feasible. The best scenario for him is to jump and succeed in getting across the stream. Then he can continue on his way with little loss of time or energy. The worst scenario for him is to jump and fail to get across the stream, for then he will probably be drowned or smashed on the rocks. If he does not jump, and goes another way instead, he suffers a great loss of time and energy, but does not incur imminent death or injury; that is the intermediate scenario. Thus it is vitally important for the hunter to know whether he can jump the stream, whether he would succeed if he did try to jump it, before he decides whether to attempt the jump. Since the method of trial and error is too risky a way of finding out whether he can jump the stream, he needs a way of finding out whether he can do it in advance of trying. He can remember some of his past jumps, but he cannot remember failing with a jump that was clearly easier than this one, or succeeding with a jump that was clearly harder. He has to consider not only the width of the stream, but also the awkwardness of the place from which he would have to launch himself, the slipperiness of the rocks on which he would have to land, how tired he is, and so on. How should he try to determine whether he would succeed?

There is a natural method human of gauging one's capacities in such situations. One *imagines* oneself trying. If one then imagines oneself succeeding, one judges that if one tried, one would succeed. If instead one imagines oneself failing, one judges that if one tried, one would fail. If one is still uncertain, one repeats the thought experiment, perhaps many times. If our hunter cannot resolve the uncertainty, he will presumably take the long way round — unless he is being chased by a sabre-toothed tiger, in which case he may jump anyway.

A traditional stereotype of the imagination can make reliance on such an imaginative exercise look like a mad way of making up one's mind. For however difficult the jump, one can *imagine* succeeding with it, and however easy the jump, one can *imagine* failing with it. How can one learn anything relevant from what one chooses to imagine? Such incomprehension indicates neglect of the distinction in section 2 between voluntary and involuntary exercises of the imagination. When the hunter makes himself imagine trying to jump the stream, his imagination operates in voluntary mode. But he neither makes himself imagine succeeding nor makes himself imagine failing. Rather, having forced the initial conditions, he lets the rest of the imaginative exercise unfold without further interference. For that remainder, his imagination operates in involuntary mode. He imagines the antecedent of the conditional voluntarily, the consequent involuntarily. Left to itself, the imagination develops the scenario in a reality-oriented way, by default.

Obviously, nothing has been said to guarantee that the imagination will reach a true answer. In some cases, it reaches a false one. The imagination, like perception, memory, and every other generic human cognitive capacity, is fallible. In other cases,

when too much uncertainty is registered, the imagination reaches no answer at all. Nevertheless, under suitable conditions, the method constitutes a reliable way of forming a true belief as to what would happen in hypothetical circumstances. I have used it myself in situations like that described, and lived to tell the tale, as have other people I know. Indeed, there is no good reason to deny that, under suitable conditions, the method enables one to *know* what would happen in the hypothetical circumstances, because the conditional connection is *safe* from failure. Of course, hardly anything has been said so far to *explain* the method's reliability, except for the hint in section 1 of evolutionary pressures. Later sections will suggest some links between the method and other normal human cognitive processes. First, however, it will be useful to have some more examples, to illustrate the range of cases over which the method is applicable.

The method can be applied with complex iterations. Suppose that now the hunter's way to the next valley is blocked by broken cliffs. Can he climb through or up them? The price of trying and failing is again high: perhaps becoming exhausted, perhaps getting stuck halfway up. The hunter stares at the rock-face, trying to trace a route all the way to the top in his imagination, testing each step for feasibility by imagining what it would involve. No more rational method of answering the question is available to him.

Even in the modern world, important decision-making often relies on knowledge or beliefs acquired through the imagination. For example, you are looking round a house, wondering whether to buy it. You want to know whether, if you lived in it, you would like doing so. You voluntarily imagine the antecedent of the conditional; your final decision may depend on what consequent you involuntarily imagine.

Not all the knowledge gained from such imaginative exercises concerns the capacities and dispositions of agents. For example, if you look at a piece of furniture and then at a doorway, you can sometimes come to know whether the former would go through the latter, without measuring either, by imagining trying to get the former through the latter. You gain knowledge of spatial relations between the doorway and the furniture.

The examples so far may appear to involve an essential role for mental imagery, in some sense. But even if that appearance is veridical, we should not over-generalize to the conclusion that all imagining involves imagery. For example, you are very busy, and wonder whether to postpone a lunch appointment with a friend. You want to know whether, if you did postpone the appointment, she would be upset. You voluntarily imagine the antecedent of the conditional. You *might* then involuntarily form a visual image of your friend with a composed or disappointed face. But no such imagery is *necessary* for imagining her reacting with composure or disappointment. Again, suppose that a politician is trying to work out what his core supporters would do at the next election if he voted for gun control. He imagines their reactions, but doing so need not involve mental imagery.

Given this heterogeneity of knowledge-yielding imaginative exercises, we need a more abstract characterization of what is going on.

4. As Frank Ramsey observed, how we evaluate conditionals is closely tied to how we update our beliefs on new information. For example, suppose that someone tells a shepherd ‘The sheep have broken out of the pen and disappeared’, and the shepherd

concludes ‘The sheep have gone down to the river’. Then presumably, even if the shepherd had not been given the testimony, he could still have reached the indicative conditional conclusion ‘If the sheep have broken out of the pen and disappeared, they have gone down to the river’ (or the subjunctive conditional conclusion ‘If the sheep had broken out of the pen and disappeared, they would have gone down to the river’, depending on subtle differences in the cognitive process). The testimony in the first case is the same as the antecedent of the conditional in the second case, and the unconditional conclusion in the first case is the same as the consequent of the conditional in the second case. If we regard the shepherd’s updating of his beliefs in the first case as an online process, then we can regard his evaluation of the conditional in the second case as the corresponding offline process. If he accepted the conditional on the basis of an imaginative exercise similar to those in section 3, as we may assume, then that imaginative exercise is the offline analogue of online updating. Very roughly, the online and offline processes take the same input — ‘The sheep have broken out of the pen and disappeared’ — and deliver the same output — ‘The sheep have gone down to the river’ — by the same means. One process is online and the other offline in virtue of the different sources of the input. If we apply the term ‘imagine’ on the basis of the source of the input, we shall classify only the offline process as an imaginative exercise. If we apply the term ‘imagine’ on the basis of the processing between input and output, we shall classify the online process as an imaginative exercise too. It is more important to see the underlying cognitive similarity than to decide exactly how to use the word ‘imagine’.

In many of the examples in section 3, the cognitive process took a mix of online and offline input. When the hunter imagines himself trying to jump the stream, he also

has to look carefully at its banks in front of him, to tailor his imaginative exercise as exactly as he can to their actual contours. Voluntarily in imagination, he somehow adds his jumping to the perceived scene. But that element of offline input is enough to make it a clear case of imagination.

Typically, there are further differences between the online and offline processes. When the new information to be updated on derives from sensory perception, we are hard put to articulate it verbally in its full specificity, to be the antecedent of a conditional. Moreover, when we learn something by perception or testimony, we usually learn other things too. When you learn by sight that the sheep have gone, you usually also learn that *you can see that* the sheep have gone. When you learn by testimony that the sheep have gone, you usually also learn that *you have been told that* the sheep have gone. But often we do not build that extra information into the antecedent of the corresponding conditional. For conditionals such as ‘If the sheep have gone, I can see that the sheep have gone’ and ‘If the sheep have gone, I have been told that the sheep have gone’ are far from trivial. Nevertheless, these differences of detail do not undermine the strong cognitive analogy between the online and offline processes.

One consequence of the analogy is that any scepticism about the offline processes is liable to generalize to the online processes too. For example, someone who doubts that in suitable conditions the shepherd’s imaginative exercise enables him to know that if the sheep have broken out of the pen and disappeared, they have gone down to the river should also doubt that in suitable conditions, on being told that the sheep have broken out of the pen and disappeared, the shepherd can know that they have gone down to the river. That looks like the thin end of the wedge for a far more general scepticism.

Much of our knowledge of the future depends on cognitive processes relevantly similar to imaginative exercises. For consider some cognitive process by which we form expectations about the future based on our knowledge of the present. Although we know far less about the future than we do about the recent past, we are not entirely ignorant of the future. The process may be more or less hard-wired into our brains, such as one for forming expectations about where a moving object in our visual field will be in a fraction of a second's time. Alternatively, the process may involve complex conscious reasoning, such as one for forming expectations about the political situation a year from now. Our imagination enables us to apply such a cognitive process offline, to imagined input about a time t , to reach a conditional conclusion of the form 'If X were to obtain at t , then Y would obtain at $t + 1$ '. Similarly, we can reach a conditional conclusion of the form 'If Y were to obtain at $t + 1$, then Z would obtain at $t + 2$ '. Under conditions that make it legitimate to assume transitivity for the conditionals at issue, we can derive the further conclusion 'If X were to obtain at t , then Z would obtain at $t + 2$ ', and so on. If these offline imaginative exercises are unreliable, the likeliest explanation is that the corresponding online processes for forming expectations about the future are also unreliable.

An analogous point applies to large tracts of our knowledge of the past. Consider cases where we use inference to the best explanation to solve a crime or to interpret an archaeological site. For such abductive reasoning, we need auxiliary conditional premises of the form 'If the explanans had obtained in the past, then the explanandum would obtain now'. We may have to obtain those conditionals in the same way as before, by an

imaginative exercise. Unless one has general sceptical inclinations, it is unwise to deny that, in suitable conditions, imaginative exercises are a source of knowledge.

Could someone argue that what have here been called 'imaginative exercises' are really just inductive inferences? Most of them depend somehow on past experience, and go beyond it non-deductively. If that suffices for a cognitive process to be an inductive inference, then they are inductive inferences. But they do not depend on the subject's *remembering* the relevant past experiences. What matters is whether they have made the subject skilful enough in performing imaginative exercise itself. It is irrelevant to the process whether the subject can assemble the particular premises of the supposed inductive inference. Nor is it remotely clear in the given cases how to fill in 'F' and 'G' in the conclusion of the supposed inductive inference, 'All Fs are Gs' (or 'Most Fs are Gs', for that matter). Thus it is also irrelevant to the process whether the subject can formulate the general conclusion of the supposed inductive inference. The imaginative exercises are inductive inferences only in a sense so loose as to be entirely unhelpful.

5. We can test the cognitive view of the imagination by confronting a genuine structural difference between imagination and perception, one sometimes used to contrast perception as reality-directed with imagination as supposedly not so. The structured difference is this. If you perceive an F, then there is an F that you perceive. Macbeth literally saw a dagger only if there was a dagger that he saw; otherwise he only seemed to himself to be seeing a dagger. By contrast, you may literally *imagine* an F even if there is no F that you imagine. I am imagining a golden mountain in Austria. Since there is no golden mountain at all in Austria, *a fortiori* there is no golden mountain in Austria that I

am imagining. What should we make of this difference? I will argue that it is predicted by the cognitive account of the imagination sketched above, and so confirms rather than disconfirms that account.

First, consider the involuntary imaginative exercise of developing an initial hypothetical supposition, in effect answering a ‘What if?’ question. On the cognitive account, various offline cognitive procedures add further conclusions to a pool that starts with the initial supposition. The process can be iterated indefinitely. Although most of the procedures are non-deductive, reasoning by deductive logic is an especially good procedure from a cognitive point of view, since it guarantees truth-preservation. The deductive aspect of the whole process will look something like the method of tableaux in first-order logic, by which the consequences of the initial premises are teased out — although in our imaginative exercises the aim is not usually to reach a contradiction. For each logical constant, there is a rule for extending a branch of a tableau which contains a formula whose main connective is that constant. For present purposes, the rule that matters is the one for the quantifier \exists (‘something’). It says that if a branch contains a formula of the form $\exists x \Phi(x)$, then one may add to it the formula $\Phi(a)$, where a is an individual constant that has not previously appeared on that branch. One can then apply the rule for the main connective in $\Phi(a)$ (if it has one), and so on. Informally, we can regard a as an ‘arbitrary name’ for a satisfier of the formula $\Phi(x)$ as a value of the variable x — if there are several satisfiers, it does not matter which one. The point of the ban on previous occurrences of a in the branch is to avoid prejudging anything else about the satisfier. If an interpretation I verifies every formula on the branch before $\Phi(a)$, then it verifies $\exists x \Phi(x)$, so some interpretation differing from I at most in making an object o

the value of x verifies $\Phi(x)$, so modifying I to make o the denotation of a yields an interpretation that verifies every formula on the branch up to and including $\Phi(a)$. The formula $\Phi(a)$ has the form of a formula expressing a singular proposition about a particular object, even if the individual constant a does not really denote anything: after all, $\exists x \Phi(x)$ may be false on its intended interpretation.

Now read $\Phi(x)$ as ' x is a golden mountain in Austria'. Thus $\exists x \Phi(x)$ expresses the proposition that there is a golden mountain in Austria. If an imaginative exercise starts from that false hypothesis, the cognitive account of the imagination (defeasibly) predicts that developing it properly will involve applying some informal analogue of the tableau rule for \exists , and thereby adding some informal analogue of $\Phi(a)$ to the development. That is, one will do something formally similar to thinking of some particular object that *it* is a golden mountain in Austria, even though really there is no particular object of which one is thinking that *it* is a golden mountain in Austria. But that is a good description of a case in which one imagines a golden mountain in Austria, even though there is no golden mountain in Austria that one imagines. Of course, the description does not entail that one forms a mental image of a golden mountain in Austria, but we saw in section 3 that forming mental imagery is not necessary for imagining. Although one may imagine a golden mountain in Austria by forming a mental image of one, the informal analogue of $\Phi(a)$ may be a mental image. Thus the very feature that was supposed to differentiate imagining from reality-directed attitudes is predicted by the account of imagining as a reality-directed attitude.

In more general terms, the felt tendency of the imagination to be specific, to fill in details, is partly explained by the hypothesis that it uses something like the tableau rule

for the quantifier \exists in developing an initial supposition. The explanation can be strengthened with the additional hypothesis that the imagination also uses something like the tableau rule for disjunction, \vee , in developing the initial supposition. The rule says that if a branch contains a formula of the form $\Phi \vee \Psi$, then one may divide it into two sub-branches, one containing all formulas already on the branch and Φ , the other containing all formulas already on the branch and Ψ . Informally, the idea is to explore the two (not mutually exclusive) ways in which the disjunction may hold separately. This corresponds to the dissatisfaction one tends to feel about, for instance, just imagining that the keys are in either the kitchen or the bedroom while neither imagining that they are in the kitchen nor imagining that they are in the bedroom.

It is almost trivial that the imagination uses something like the other standard tableau rules in developing an initial supposition. For example, the rule for conjunction, \wedge , says that if a branch contains a formula of the form $\Phi \wedge \Psi$, then one may add to it both the formula Φ and the formula Ψ . Similarly, it would be hard to imagine that Mary is tall and thin without also imagining that she is tall and imagining that she is thin. On the cognitive view of the imagination, one would expect it to use something like the tableau rules in developing an initial supposition. The imagination conforms with that prediction.

6. The cognitive view of the imagination does *not* predict that it will be cognitively reliable only for tasks just like those it evolved to serve. Its tendency to use something like rules of deductive logic is an example to the contrary, since they are quite generally truth-preserving. The simplest forms of reasoning to implement that are truth-preserving for practical matters are truth-preserving for all matters.

Once we appreciate the role of the imagination as a standard means for evaluating conditionals and modal claims, we should be much less inclined to regard the use of thought experiments in philosophy (or natural science) as constituting any highly distinctive method. Still less should we be tempted to characterize such a method in terms of cooked-up categories such as ‘philosophical intuition’, which serve mainly to obscure the similarities between thought experiments and more routine exercises of the imagination in virtually every branch of human inquiry. We simply reserve the term ‘thought experiment’ for the more elaborate and eye-catching members of the kind.

One might suppose that, as science progresses, the role of the imagination will increasingly be confined to the context of discovery, and that in the context of justification it will gradually be replaced by more rigorous methods. But there is evidence to the contrary. For rigorous science relies on mathematics, and so indirectly on the axioms or first principles of mathematics. But when one examines the justifications mathematicians give of their first principles, such as axioms of set theory, one finds unashamed appeals to the imagination. Things are complicated, because the justification is abductive: it involves the derivability of standard ‘working’ mathematics from the candidate first principles — and, one hopes, the non-derivability of contradictions. But it also involves the way in which the candidate first principles fit together as parts of an intrinsically simple, elegant, and unified picture of mathematical reality. Those theoretical virtues are needed to trump rival gruesomely gerrymandered pictures, for example, one on which the hierarchy of sets is truncated at an arbitrarily chosen large cardinal, in the usual manner characteristic of abductive theory comparison. Moreover, since by Gödel’s second incompleteness theorem the axioms can be proved jointly

consistent only on the basis of even stronger axioms that are inconsistent if the former are, much of mathematicians' confidence in the consistency of their axioms comes from the way in which they form such a coherent picture. Of course, they also have inductive evidence of their consistency, because no one has ever yet derived a contradiction from them, but mathematicians put comparatively little weight on such inductive evidence. They have much less confidence in the consistency of alternative axiom sets from which no one has ever yet derived a contradiction, but which do not seem to form a coherent picture. Without an imagination, one would be in no position to judge whether some candidate first principles form a coherent picture of mathematical reality. For the foreseeable future at the very least, imagination will play a crucial role in the context of scientific justification, not only in the context of scientific discovery.

Whatever the function or evolutionary origin of our capacities, we are not forbidden to use them for other ends, including playful ones. Imagination has enabled humans to produce works of art, great and awful. Indeed, the combination of voluntary hypothetical suppositions and involuntary developments of those suppositions is reminiscent of the way some novelists describe novel-writing. An imagination is cognitively powerful only if it is capable of producing and developing fictions. But if we try to understand the imagination while taking for granted that fiction is its central or typical business, we go as badly wrong as we would if we tried to understand arms and legs while taking for granted that dancing is their central or typical business.

Acknowledgements

This paper is based on my talk to the 2012 Claremont conference on Knowledge through Imagination. Other versions of the material were presented at a workshop at the Inter-University Centre in Dubrovnik and at the following universities: Belgrade, Bergen, Birmingham, Boğaziçi (Istanbul), Cambridge, Edinburgh, Gothenburg, Liverpool, London (King's College), Manchester, Oxford, Warwick (apologies to any I have forgotten). I am grateful to the audiences at all these events for stimulating discussion. Despite all that help, the piece remains speculative and programmatic, so I have not dignified it with footnotes. It develops ideas in my books *The Philosophy of Philosophy* (Oxford: Blackwell, 2007) and, to a much lesser extent, *Knowledge and its Limits* (Oxford: Oxford University Press, 2000).