

Philosophy and the Understanding of Ignorance

Bernard Williams

Our subject is what we do not know; and this is a meeting under the auspices of philosophy—a meeting, indeed, designed to advance UNESCO's programme in philosophy. The conjunction of this subject and these auspices already confronts us with certain questions. How is philosophy related to ignorance?

One question I shall not pursue is whether philosophy itself just is a kind of ignorance, and whether there is such a thing as philosophical knowledge at all. It has often been said, particularly by positivists, that philosophy is virtually by definition a home of ignorance, that it consists of questions which we do not know how to answer by established forms of enquiry. On this account, questions that have previously been part of philosophy may mature into questions for the sciences or for other disciplines capable of accumulating knowledge. Thus some questions that have belonged to philosophy have moved into physics, others into linguistics, others into psychology. At the present time, questions about mind and body, perpetual concerns of philosophy, are (some would say) moving into the realm of cognitive science.

The processes presented in this picture do have to be distinguished from something else, namely the fact that, at least in the English language, the word "philosophy" has a more restricted reference than it used to have. Newton's great book was a contribution to a subject then called "natural philosophy," but those who practised the subject under that name were capable of drawing distinctions between that subject and the kind of inquiry that we would now call "philosophy," and drawing it on lines broadly familiar today. It is not this verbal point that I have in mind, but rather processes by which questions develop from a status in which we do not know how to answer them, and they belong to

what is indisputably philosophy, to a status in which they have become a proper subject for systematic, and perhaps scientific, enquiry. The picture that emphasizes these processes might, then, suggest that not only does philosophy contain no knowledge, but it is virtually defined as not doing so. As I said, this is not a question that I want to take much further, since in general I believe that philosophy is not at its most interesting when it is talking about itself. But it perhaps can at least be said in leaving this topic that there is a certain paradox in thinking that what these processes demonstrate is that philosophy offers no knowledge of any kind. There certainly are developments in which—to put it very roughly—questions move out of philosophy into other fields of enquiry. But these developments do not leave philosophy entirely passive in relation to them; it is not merely the wine cellar in which a question matures until it reaches the state in which it can be put on the market of science. Nor is it a place in which questions wait until other sciences have developed, by their own processes, to a point at which they are capable of picking them up. Philosophy itself contributes to these processes, and indeed it can contribute so much that it is very artificial to say that precisely the same question has graduated from philosophy to science. Questions are reformulated and redefined under such processes, and those refinements and redefinitions, whatever precisely their status, are certainly a product of philosophy.

What we should consider for rather longer is the subject, not of philosophy as ignorance, but of what, rather, philosophy might conceivably tell us about ignorance more generally. This is a subject on which philosophy has over the centuries shown a remarkably high degree of ambition: from the ancient world through early modern philosophy to contemporary studies, philosophers have been keen to suggest that we know little or nothing. I do not want to deny the philosophical importance of these skeptical arguments, but their importance seems to me to lie much more in what they may be able to tell us about the nature or the basis of knowledge, than in any actual determination of what it is that we do not know. The reason for this lies in their extreme generality: typically, they try to show us that we know nothing, or know only the contents of immediate experience, or know only some simple necessary truths. Such conclusions do not make the boundary

between what is known and what is not known at all interesting, in particular because they do not represent it as a boundary that could in principle move. So while the possibility of skepticism remains a challenge that a theory of knowledge has to deal with in one way or another—if only, in the manner of some contemporary theories, by turning their back on it—it does not, by its very nature, shed any very interesting light on the question of what we know as opposed to what we do not know.

This is perhaps only an example of a more general point, that skepticism is more interesting and more disquieting if it is based on more particular kinds of consideration. What is often called skepticism “about other minds” does, once again, represent an important area of philosophical reflection, but when it is conducted at a level of very high generality, it does not offer us any very anxious ground of concern. Someone who is genuinely worried whether he or she knows that another person is in pain, even if that person is writhing on the floor with a knife in his leg, is someone who should be referred for clinical treatment. This, like serious skepticism about the external world, is a kind of skepticism which, as Descartes himself insisted, should be reserved for a reflective philosophical exercise. But this is not the only type of skepticism about other minds, and to insist upon skepticism at such a high level of generality actually serves to disguise the disquieting force of more particular manifestations of skepticism about other minds. While it is absurd, at a practical rather than a purely theoretical level, to wonder whether other people have any emotions, feelings or sensations at all, it is not at all absurd to entertain real doubt about what the character of their feelings is, or to raise the question of how much we know or ever could know about someone else’s inner life once one gets beyond the familiar features of it that are, as a matter of mutual human understanding, genuinely given to us.

The same point applies to skepticism about our knowledge of the past, and also skepticism about the physical world. To raise the question of whether we know anything about the past *at all* is indeed a pathological state if it is considered as more than a device for investigating such knowledge. But, once again, to insist on simply this kind of skepticism, and hence, in reaction to that,

on a straightforward rejection of it, is to obscure the extent and depth of skepticism—a skepticism that may be entirely justified—that one may experience with regard to historical narratives. There is indeed a deep and structural problem about the credentials of such narratives, and one fails to grasp the force of that problem if one raises only the absolutely general problem of skepticism about the past; or if, as is often the case, one extends the skepticism appropriate to historical narrative even to the least ambitiously reported elements of a chronicle. It is certainly true that an earthquake struck Los Angeles in January 1994, and that Julius Caesar was killed in Rome in March 44 BC, and it is only if that level of affirmation about the past is acceptable that one can get on to raise the really interesting questions of how much we do not know, as opposed to what we do know, about the past.

Our concerns about what we do not know can get a real and compelling grip on us only if there are some things that we do know. This is why traditional philosophical skepticism, suggesting as it does that we may know nothing at all (or nothing about other minds, or nothing about the past) is not compelling, in the sense, at least, of attacking our assurance about those things. Still more, we should not be compelled—we cannot be compelled—by very general reflections directed not simply against knowledge but against truth. Recent forms of skepticism, drawing in many cases on a very partial reading of Nietzsche, have tried to discredit the notion of truth altogether. In doing this, they typically take on a tone of mild heroism about their project of uncovering our illusions (as they are inevitably, but on their own account misguidedly, tempted to call them).

What is disquieting about such positions is not so much their self-refutation, as their false promise of discomfort. What casts suspicion on everything casts suspicion on nothing: even the common or garden paranoiac needs his exercise book of carefully researched facts. Our suspect assurances will be undermined, as they are in Nietzsche's own practice, only by an interpretative attention which is selectively directed, and which accepts the materials that are needed if its direction of attention is to be intelligible. In the case of history, we can have doubts about our understanding of the past only if we have a past, and we have a past (as Wittgenstein empha-

sized) only if there are some things that certainly happened in it. It is only because we can accept large numbers of facts about the past, many of them in themselves very boring, that we can confront the genuinely disturbing suggestion that historical understanding requires narrative, and narrative demands closure, and closure in history is always a fiction and often a lie.

In the case of the physical world as such, the boundary between compelling versions of skepticism and noncompelling versions of it takes a different shape. A blank skepticism about the external world does not constitute a compelling skepticism, in the sense that it gives one any reason to worry whether everything we think we know about the external world might not be false. Equally, there can be no reason to suppose that we do not know a very large number of truths based on observation, whether the observation involves scientific apparatus, or is unassisted by apparatus. The interesting question arises with regard to our theoretical statements about nature, and the theoretical entities that are typically introduced by those theories. It is still very much a live issue in the philosophy of science, to what extent our theoretical understandings of nature constitute knowledge, but there are two different levels at which such concerns arise. On the one hand, there is the kind of question often expressed in terms of the choice between instrumentalism and realism, and this asks us whether *any* theory could constitute knowledge, or whether our real access to knowledge about nature is confined simply to the level of what is observable (where what is "observable" is itself an issue that involves some extremely pressing difficulties). On the other hand, we may ask whether some such theories *rather than others* constitute knowledge. At this level, it is no longer a question of whether there is something inherent in such theories that prevents them (as opposed to observations) from constituting knowledge; it is a question, rather, of whether some scientific theories make a better claim to that status than others.

It may be thought that this latter kind of question is not itself philosophical, but rather is precisely the kind of question that forms the substance of scientific practice, which is concerned with advancing and preferring some theories to others. However, I think it is a mistake to associate the term "philosophical" always

with the most general kinds of question that can arise with regard to theories or hypotheses. To put it another way, there are areas of science in which the boundaries between the philosophy of science and science itself are pretty arbitrarily drawn. A notable example of this is quantum mechanics, in which a major concern precisely is the conceptual constraints that may be appropriate to the descriptions and explanations that it offers. While no-one can discuss these matters without being well informed about the contemporary currency of such physical theory, it becomes beyond that point fairly arbitrary in many cases to determine whether somebody is discussing the philosophical theory of quantum mechanics, or is doing quantum mechanics at a very highly theoretical level.

Some of the philosophical issues that I have mentioned themselves constitute questions to which, as it seems to me, we do not know the answers, but of which we might have hopes, if not of acquiring answers, of at least advancing our understanding in the coming years. This seems to me true with regard to the issues that I have just mentioned of the relation between observation and theory in the philosophy of science, and also with regard to the structure of a historical narrative and its relations to what I have called the materials of chronicle. We may reasonably be said not at present to know how best to discuss these subjects, and it does not seem to me foolish to suppose that we might come to understand these things better. These subjects can philosophical subjects, so if it is correct that we can hope to make progress with them, these will be examples of questions *within* philosophy to which we might come to have better answers than we have at present. For that very reason, we should not, as I have already suggested, pay too much attention to arguments that are designed to show that we can never come to know anything in philosophy. The phrase *philosophia perennis*, one might say, expresses pessimism. Nor is the relevance of philosophy to our topic to be found principally in its old invitation to suppose that we do not know anything at all.

There is, however, a quite different kind of philosophical argument that bears upon our theme, and which I think is more interestingly relevant to it. This kind of argument suggests that, while we can no doubt state some things that we do not know, we cannot in general state with confidence what it is that we do not know.

There are limits in principle to the extent that we can know what we do not know. The point can be demonstrated, in a rather restricted way, by the following argument.¹ Suppose that I am looking at a room with a large number of people in it. I do not know how many people there are in the room and I am not in a position to count them. But I do know (for instance) that there are fewer than a thousand people, and more generally, for various numbers n , I know that there are fewer than n people. When n gets rather smaller (closer to what is in fact the actual number of people), it ceases to be the case that I know that there are less than n people present. However, it is obvious that there is no particular number n such that I know that there are less than n people present but don't know that there are less than $(n-1)$ people present: so, for some numbers, it must be the case that I don't know whether I don't know that there are less than that number of persons present.

Now this is of course an extremely artificial and regimented example, but it illustrates a much more general point about the knowledge of ignorance. In general there are important limits to the knowledge of ignorance, and this point is implicit in the notion of a margin of error, a notion which is itself involved in the concept of knowledge. If I do not know in every case what I do not know, it is also true that, even if I do know something, I do not necessarily know that I know it: the principle so beloved of certain philosophers, that if I know, I must know that I know (the so-called "KK principle"), is certainly false. This has important consequences for any enquiry of the kind that we are addressing. We may indeed be able to mention some things that we do not know; I have already mentioned one or two such things, and I shall go on to mention one or two others. In some cases, it is possible to specify what it is that one does not know in such a way that one can know that one does not know it: so, in the previous example, I do know that I do not know *the exact number of people in the room*. Similarly, I, personally, express knowledge if I say that I do not know the name of the present Archbishop of Milan, or of Leonardo da Vinci's father. But we do not always know that we do not know certain things, and, as has already been argued, it can be shown that it is impossible that in every case of our not knowing something, we know that that is the situation.

There are other and less formal kinds of argument that lead to the same conclusion. Among the things that we do not know are things that people used to know, and also things that people might in the future come to know. There are some important asymmetries between these two. One asymmetry is that, with regard to things that people used to know but which we do not know, we suppose in general that we possess the terms in which such knowledge might be expressed. We know, or know of, the terms in which past people would have expressed their knowledge, terms which we might at some level be able to understand. This is because our picture of knowledge that has got lost is a picture of knowledge expressed (very broadly speaking) in terms which themselves have not got lost.

With regard to knowledge that people may acquire in the future, however, we do not necessarily have any such conception. It was a positivist error, to which no-one now is attached, to suppose that the fundamental vocabulary or conceptual resources of science are fixed, and that what will be discovered in the future can only be new facts or theories expressible in that same vocabulary. On the contrary, we believe that theoretical advances typically consist of introducing new concepts, and that those concepts may not be strictly commensurable with concepts that we presently have. I do not think that this need lead to a radical relativism; but it does mean that future science may contain theoretical innovations which, as things are, we could not understand at all. It is a disputed issue whether there might be such innovations which we could not in principle come to understand. Perhaps, as some have suggested, our idea of a possible language must be the idea of a language which we could in principle come to understand. But we need not engage with that issue; our present question merely concerns the possibility that future discoveries may be expressible only in a language which we now as a matter of fact cannot understand, and that certainly could not be ruled out by any argument of principle. Such future discoveries, we are assuming, would be *discoveries*, which is to say that they could constitute knowledge. But we cannot know what that knowledge would be, for the radical reason that we have no ways of expressing it; consequently we cannot know what it is exactly that, in lacking that

knowledge, we do not know. We can in various degrees locate it, and we must locate it if we are to identify it as a possible scientific discovery—we must locate it in a space of problems, for instance. But this is a long way from our knowing in any exact terms what it is that we do not know.

This consideration, that one cannot foresee the terms in which future discoveries might be expressed, applies particularly to certain areas of scientific enquiry which presently invite the thought that there is not only a great deal that we do not know—that is obviously true of every area—but that there might be a certain insecurity to the knowledge that we hope we may already possess. I cannot pretend to any expertise at all in theoretical physics, and what I have to say about such subjects is of the most unprofessional kind. But the expert opinion of others who are better placed to understand the current situation can lead one to think that with regard both to scientific cosmology and to particle physics (fields closely allied to each other) the structures of theory that we have are in more than one way rather perilously related to what we definitely know. In the case of cosmology, the current conclusions or speculations may be rather extensively cantilevered out from observations and interpretations of those observations, to such an extent that the alteration of an assumption near the start of the argument might lead to radical readjustments of the theoretical picture as a whole. Among these conclusions, perhaps, are the orders of magnitude of time and distance that are invoked in the theory, and it may be reasonable to think that in these and in other respects some radical revision of current theory may await us.

The situation is in this respect different from that in other areas of science, which are themselves just as theoretical. (This illustrates the point I mentioned earlier, that philosophical distinctions between theory and observation do not coincide with the questions that interest us most in this discussion.) I think it is generally agreed that there is no question that our understanding of the structure of DNA and the mechanics of the transmission of genetic information could turn out to be other than fundamentally correct. Our understanding of this seems now to have attained the point at which much of it constitutes part of the data of any future theory rather than part of the content that could be replaced by future theory.

The case of particle physics, to the (again, very limited) extent that I understand the situation, is in some ways similar to that of cosmology, but here even those who are very confident that present theory is basically sound are conscious of a definite limitation on the increase of knowledge, which lies in financial and practical limits to controlled experimentation. I suppose that it is possible that the European super-collider, if indeed it is finally built, may be the last such machine ever to be built on earth, and we know already that not every question that can be raised within the present structure of particle physics, even supposing that structure to continue without radical revision, can be answered by that machine. This is an area in which theory has far outrun any foreseeable experiment. If this is so, then our knowledge of our ignorance will have a special structure; it will reflect the fact that, even if our questions are well posed, the amount of energy (and hence the resources) that would be needed to conduct experiments that might answer them simply outruns anything that we could bring to bear on the question. Such high energies are of course at work in the universe, and the ultimate laboratories of particle physics are in the stars; but we cannot conduct controlled experiments with the stars, and the gap between what we can sensibly construct for the purposes of experiment, and the energies required by experiments that could answer our questions, may be forever unbridgeable.

These speculations about the sciences have, as I have said, very little authority. However, they may at least illustrate a general point which is perhaps likely to be overlooked by philosophy, that it is possible to gain from inside a science itself some idea of the questions that it might or might not be able to answer, even though we grant the point that the terms in which an answer might be given may be to varying degrees not known to us. This apparent paradox can be resolved because to some extent we can form an idea of the experimental situation in which an answer would have to be pursued, if any answer were to be forthcoming, and we may be able to see that we cannot get to that situation, or that it is unlikely that we shall be able to do so. All of this must, inevitably, be speculative because, in not foreseeing future theory, we equally cannot foresee future experiment; perhaps less expensive ways of investigating nature on the micro-scale may come to be available,

which at present we cannot conceive. Certainly in these areas, our thoughts about what we do not know must be structured both by considerations of what experimental routes are or might become available to us, and also, as notably in the case of cosmology, by the question of the degree to which the most ambitious parts of the theory are, as I have put it, "cantilevered out" from the undisputed observational material; to such an extent that there is a possibility that under revision of some element in the structure, the theory might dramatically implode, and alter and change into something very different. This once again illustrates the point that it is a mistake for philosophers simply to discuss the status of theoretical science *as such*. Particle physics, cosmology, and indeed molecular biology are all theoretical sciences, but their relations to future experiment and refutation, and correspondingly their relations to our present ignorance, are very different from one another.

There is a very dramatic example of current ignorance the relations of which both to philosophy and to scientific enquiry are very different from what has so far been discussed. The main difference is that in this case it is a philosophical question how far our ignorance has anything to do with scientific enquiry at all. This is the problem of consciousness, the question of how we may explain, or even adequately describe, the difference between creatures *for whom things exist or happen in certain ways*, and creatures who lack any such experience. A great deal of work has been done in this area in recent years, and is now being done, and it can at least be said that we do possess more material, above all at the neurophysiological level, than we did before, which might prove relevant to a solution of this problem if a solution of this problem is possible. My extremely cautious formulation of this fact is dictated by the remarkably wide range of opinions that it is still possible to hold on this topic. At one extreme there are people, such as Daniel Dennett, who would claim at least in outline actually to have solved the problem. Then there are those, perhaps the majority, who think the problem has not been solved, but that some combination of physiological research, work in cognitive science, and philosophical clarification should or may eventually bring us the solution. Again, there are those, notably Thomas Nagel, who believe that at the present time we not only have no way of relat-

ing the facts of consciousness, the first personal-experiences of a conscious agent, to changes in the brain and the nervous system, but that we lack even any idea of what a theory that related those two things in a perspicuous and explanatory way might look like. Some who think this nevertheless suppose that this may be a situation such as that which has preceded other scientific advances, in which a step which has seemed incomprehensible to us has in fact been taken by means of a new conceptual invention.

At the furthest extreme of pessimism on the subject of consciousness is Colin McGinn, who suspects that the problem of giving a coherent and explanatory account of human consciousness is insoluble for a quite special reason, namely that the structure of the brain is such that it cannot possibly grasp this aspect of its own operations. This conclusion naturally, and perhaps healthily, attracts skepticism, because it sounds a great deal too much like previous claims about what might prove scientifically unintelligible: in particular the so-called problem of the nature of life, which gave rise to the pessimistic and obscurantist position of vitalism. However, this analogy does not necessarily lead one to discount McGinn's pessimism (which is in any case better argued than any vitalist position was): it also leads us to a further thought about what might happen to this problem. In a certain sense, the problem of the nature of life—the problem which vitalism declared insoluble—was never exactly solved: rather, we have learned so much about the operations of living things that the problem in that form has gone away. As things stand, it seems to us as though the problem of the nature of consciousness could not be like that, since consciousness seems so present and manifest a phenomenon we cannot understand how the question of its nature could, by an enormous elaboration of physiological and psychological understanding, seem ultimately to have evaporated. But perhaps that impression itself is a function of the present state of our understanding or rather of our lack of it.

Certainly the problem of consciousness is one that combines in the highest degree the various kinds of doubt that can constitute our admissions of our own ignorance. We are not agreed that there is a problem; or, if there is, whether it has been solved; or, if it has not, whether it is soluble; nor, if it is soluble, whether the

present obstacles to our solving it are technical, theoretical, or conceptual. In addition, we are not agreed whether the problem is of a kind which, even if it cannot directly be solved, might eventually turn out to have gone away. It is hard to think of any other problem-area in which so many impressions of the nature of our ignorance can coexist.

There is one last set of problems which, particularly perhaps as we discuss these matters under the auspices of UNESCO, should be mentioned as peculiarly defying our understanding and revealing our ignorance: the problems, that is to say, of how to live together. No-one is going to deny, presumably, that there is such a problem, in the sense that various groups constantly, repeatedly, all over the world, find it remarkably difficult to live with one another. In this sense, certainly, there is something we do not know—how to live together, except under a variety of fairly favorable circumstances. However, it is a different claim, and it might be thought on reflection rather an optimistic claim, that this represents an intellectual problem: the problem, as we might put it, that we do not know why we do not know how to live with each other. This suggests that there is something to be found out about the causes of conflict, something which is presently hidden from us and which, when it is found out, may enable us to negotiate and progressively eliminate those conflicts. Perhaps there is some such thing which is presently hidden from us, and certainly we should not relax our efforts in asking what it might be, seeking the help of psychology, anthropology, history and perhaps biology in so doing. However, we cannot be sure that there is an intellectual problem which takes the form of finding some central explanation which is hidden. Perhaps, rather, we already know most of what is to be known at a general level about the causes of human conflict, and there is nothing very deep or extensive, which we do not already recognize, to be learned about it. What we need to do is rather to organize the resources which, in general terms, we already know to be necessary to deal with such conflicts, in so far as they can be dealt with, understanding each in terms of its own circumstances. If we cannot mobilize the resources, or it is not the sort of conflict that will respond to any resources that we might mobilize, we shall not suppose that there is some further, poten-

tially revealing thing we do not know. We shall have to reconcile ourselves to a perfectly obvious thing, that we do not know how to deal with the conflict.

This second, and bleaker, account, we do not necessarily have to accept. Perhaps the happier idea, that there is still some important thing to be learned about why human beings are so disposed to hate and kill one another, has some promise. If it has, we would certainly like to know that it has. If the bleaker story is true, however, perhaps we do not want to know that it is. With this, the most pressing of all our questions, the position is as it often is with matters that come close to our interests: we cannot know whether we really want to overcome our ignorance until we have done so.

Notes

1. I owe this argument to T. Williamson, who develops it in a more precise and richer form in his book *Vagueness*, London, 1995.