
The future of philosophy

John R. Searle

Department of Philosophy, University of California—Berkeley, Berkeley, CA 94720, USA

There is no sharp dividing line between science and philosophy, but philosophical problems tend to have three special features. First, they tend to concern large frameworks rather than specific questions within the framework. Second, they are questions for which there is no generally accepted method of solution. And third they tend to involve conceptual issues. For these reasons a philosophical problem such as the nature of life can become a scientific problem if it is put into a shape where it admits of scientific resolution. Philosophy in the 20th century was characterized by a concern with logic and language, which is markedly different from the concerns of earlier centuries of philosophy. However, it shared with the European philosophical tradition since the 17th century an excessive concern with issues in the theory of knowledge and with scepticism. As the century ends, we can see that scepticism no longer occupies centre stage, and this enables us to have a more constructive approach to philosophical problems than was possible for earlier generations. This situation is somewhat analogous to the shift from the sceptical concerns of Socrates and Plato to the constructive philosophical enterprise of Aristotle. With that in mind, we can discuss the prospects for the following six philosophical areas: (i) the traditional mind–body problem; (ii) the philosophy of mind and cognitive science; (iii) the philosophy of language; (iv) the philosophy of society; (v) ethics and practical reason; (vi) the philosophy of science.

The general theme of these investigations, I believe, is that the appraisal of the true significance of issues in the philosophy of knowledge enables us to have a more constructive account of various other philosophical problems than has typically been possible for the past three centuries.

Keywords: philosophy; science; mind–body problem; cognitive science; epistemology; ethics

1. PHILOSOPHY AND SCIENCE

Because this article is intended for a predominantly scientific audience, I will begin by explaining some of the similarities and differences between science and philosophy. There is no sharp dividing line between the two. Both, in principle, are universal in subject matter and both aim at the truth. However, though there is no sharp dividing line, there are important differences in method, style and presuppositions. Philosophical problems tend to have three related features that scientific problems do not have. First, philosophy is in large part concerned with questions that we have not yet found a satisfactory and systematic way to answer. Second, philosophical questions tend to be what I will call ‘framework’ questions; that is, they tend to deal with large frameworks of phenomena, rather than with specific individual questions. And third, philosophical questions are typically about conceptual issues; they are often questions about our concepts and the relationship between our concepts and the world they represent.

These differences will become clearer if we consider actual examples. The question ‘What is the cause of cancer?’ is a scientific and not a philosophical question. The question ‘What is the nature of causation?’ is a philosophical and not a scientific question. Similarly the question ‘How many neurotransmitters are there?’ is a scientific and not a philosophical question; but the question ‘What is the relationship between mind and body?’ is still, in large part, a philosophical question. In each case

the philosophical questions cannot be settled by the simple application of either experimental or mathematical methods, they are about large frameworks and they involve conceptual issues. Sometimes major scientific advances are contributions to both science and philosophy because they involve changes in frameworks and revision of concepts. Einstein’s relativity theory is an obvious 20th-century example.

Because philosophy deals with framework questions and with questions that we do not know how to answer systematically, it tends to stand in a peculiar relationship to the natural sciences. As soon as we can revise and formulate a philosophical question to the point that we can find a systematic way to answer it, it ceases to be philosophical and becomes scientific. Something very much like this happened to the problem of life. It was once considered a philosophical problem how ‘inert’ matter could become ‘alive’. As we came to understand the molecular biological mechanisms of life, this ceased to be a philosophical question and became a matter of established scientific fact. It is hard for us today to recover the intensity with which this issue was once debated. The point is not so much that the mechanists won and the vitalists lost, but that we came to have a much richer concept of the biological mechanisms of life and heredity. I hope a similar thing will happen to the problem of consciousness and its relation to brain processes. As I write this it is still regarded by many as a philosophical question, but I believe with recent progress in neurobiology and with a philosophical critique of the

traditional categories of the mental and the physical, we are getting closer to being able to find a systematic scientific way to answer this question. In which case it will, like the problem of life, cease to be 'philosophical' and will become 'scientific'. These features of philosophical questions, that they tend to be framework questions and tend not to lend themselves to systematic empirical research, explains why science is always 'right' and philosophy is always 'wrong'. As soon as we find a systematic way to answer a question, and get an answer that all competent investigators in the field can agree is the correct answer, we stop calling it 'philosophical' and start calling it 'scientific'. These differences do not have the result that in philosophy anything goes, that one can say anything and make any speculation that one likes. On the contrary, precisely because we lack established empirical or mathematical methods for investigating philosophical problems, we have to be all the more rigorous and precise in our philosophical analyses.

It might seem, from what I have said, that eventually philosophy will cease to exist as a discipline as we find a systematic scientific way to answer all philosophical questions. This has been the dream of philosophers, I believe, since the time of the ancient Greeks, but in fact we have not had much success in getting rid of philosophy by solving all philosophical problems. A generation ago it was widely believed that we had at last discovered, through the efforts of Wittgenstein, Austin and other 'linguistic philosophers', systematic methods for solving philosophical questions, and it seemed to some philosophers that we might be able to solve all the questions within a few lifetimes. Austin, for example, believed that there were about a thousand philosophical questions left, and with systematic research, we should be able to solve all of them. I do not think anyone believes that today. Only a small number of the philosophical problems left us by the preceding centuries, going back to the Greek philosophers, have been amenable to scientific, mathematical and linguistic solutions. The question as to the nature of life, I believe, has been finally resolved and is no longer a philosophical question. I hope something like this will happen to the so-called mind-body problem in the 21st century. However, a very large number of other questions left us by the ancient Greeks, such as 'What is the nature of justice?', 'What is a good society?', 'What is the proper aim and goal of human life?', 'What is the nature of language and meaning?', 'What is the nature of truth?', are still very much with us as philosophical questions. I would estimate that about 90 per cent of the philosophical problems left us by the Greeks are still with us, and that we have not yet found a scientific, linguistic or mathematical way to answer them. Furthermore, new philosophical problems are constantly being thrown up and whole new areas of philosophy invented. The Greeks could not possibly have had the sort of philosophical problems we have had in getting a correct philosophical interpretation of the results of quantum mechanics, Gödel's theorem or the set theoretical paradoxes. Nor did they have such subjects as the philosophy of language or the philosophy of mind as we think of them. It seems that even at the end of the 21st century we shall still have a very large number of philosophical problems left.

2. 20TH CENTURY PHILOSOPHY: LOGIC AND LANGUAGE

So far I have been discussing certain general features of philosophy as an area of investigation and how it differs from the sciences. Before I try to make some projections about the future of philosophy, I need to say at least a few words about some special features of 20th-century philosophy. If we are going to examine the possibilities of philosophy in the 21st century, we need to know something about the jumping-off place from which we will be entering that new century. Philosophy in the 20th century has many special features, but the single most decisive difference between 20th-century philosophy and earlier epochs, is the central role of logic and language in both the methods and the subject matter. This new era actually began in 1879, when a little-known professor of mathematics in Jena named Gottlob Frege revolutionized the subject of logic and effectively invented the philosophy of language. From the time of Aristotle until 1879, the Aristotelian theory of the syllogism dominated logic to the extent that the theory of the syllogism was taken to be more or less coextensive with the discipline of logic. Even a philosopher as great as Kant could say that by the 18th century, logic as a subject was essentially complete. With the completion of the theory of the syllogism, there was nothing further to be done in logic. However, Frege revolutionized logic by inventing what came to be known as the predicate calculus, or quantificational logic, the logic of the quantifier expressions 'There is some x such that' and 'For all x , x is such that'. Frege's logic is vastly more powerful than traditional Aristotelian logic, and it is now so much a part of contemporary life that we are almost unaware of its special revolutionary features. For example, the use of the predicate calculus in computer science is now simply taken for granted, and it is hard to imagine how you would do modern computational theory without quantificational logic and set theory.

In the course of developing logic, Frege also, more or less inadvertently, invented the subject of the philosophy of language. Previous philosophers, again beginning with the Greeks, had been interested in language, but the general attitude was that we could take language for granted and get on with the more interesting philosophical issues. The picture of language and meaning that pervades, e.g. the British empiricists, Locke, Berkeley and Hume, is that words get their meanings by standing for ideas in the mind, and that the ideas in the mind stand for objects in the world by way of resemblance. For example, the word 'chair' stands for a mental picture that I have of chairs, and the mental picture stands for real chairs in the world, by way of resemblance. The mental picture looks like real chairs. Frege (as well as, much later, Wittgenstein) argued that this whole approach is bankrupt and worked out a much richer, though still inadequate, philosophy of language.

No one paid much attention to Frege except for some European mathematicians, and a young English philosopher named Bertrand Russell. The distinctive 20th-century style of doing philosophy started with Russell's famous article, published in *Mind* in 1905, called 'On denoting', which applied Frege's methods to the special problems of analysing sentences in ordinary language.

Frege himself hated ordinary language. He thought it was incoherent and self-contradictory, and that we would be much better off with a logically perfect language of the sort that he had invented. Russell was no fan of ordinary language, but he thought that the ambiguities and vagaries of ordinary language could be cleaned up by analysing ordinary-language sentences in the predicate calculus. The point for the present discussion is this. Philosophy in the 20th century had three new features as a result of the revolution brought about by Frege, Russell and Russell's student Wittgenstein, as well as Russell's colleague G. E. Moore. First, Frege's logic gave us a much more powerful tool for analysing logical relations, and for the discussion of philosophical problems in general, than previous generations of philosophers had. Second, the philosophical analysis of language itself became a central—indeed some would say *the* central—problem in philosophy. What exactly is the relationship between language and reality? How is it that words can stand for things in the real world? What exactly is the nature of truth and reference? Third, language was not only a subject of philosophical investigation, but the analysis of language was taken to be an essential tool in investigating other areas of philosophy. For this reason we need to distinguish between the philosophy of language and linguistic philosophy. The philosophy of language deals with certain general features of language, such as truth and meaning, whereas linguistic philosophy uses the methods of linguistic analysis to try to solve traditional problems. So, for example, a problem I mentioned earlier concerning the nature of causation was treated by 20th-century philosophers as a matter of analysing the use of the concept of causation both in the sciences and in ordinary life. What exactly is meant by saying that A causes B? Can we get an analysis of the causal relationship in terms of more fundamental features? This was felt by many philosophers of the 20th century to be not so much a revolutionary change in philosophy, but rather a matter of making patterns of analysis that had already been present in philosophy clearer and more precise. Thus Hume tried to analyse the notion of causation by examining the ideas about causation that he had in his mind. The 20th-century philosopher also proceeds by analysis, but instead of analysing the ideas of causation in his mind, he analyses the language we use in stating causal facts about the world.

I would not wish to give the impression that philosophy has been or is now a unified subject. There are many different schools, methods and approaches in philosophy, and the one that I have described is usually called 'analytic philosophy'. It is not the only way of doing philosophy, but in Britain, the USA and other English-speaking countries, there is no question that it has become the dominant approach to philosophy, and it is the prevailing approach in literally all of our major universities. There are other approaches, such as existentialism and phenomenology. Indeed, phenomenology and its successors can properly be described as the more influential method of doing philosophy in certain European countries, especially France. This is not the place to try to explain the differences between so called 'Continental' philosophy and analytic philosophy, but one crucial difference for the purpose of this article is that analytic

philosophers tend to be very much concerned with science and to see philosophy as aiming for exactly the same sort of objective truth that one gets in the sciences. In my experience, Continental philosophers—with some notable exceptions—tend to see philosophy as less like the sciences and more like a branch of literature, or at least closely allied to the study of literature and literary theory.

One further feature of 20th century philosophy needs to be mentioned. I have said that philosophers in the 20th century showed a special obsession with language, but the study of language as a discipline was itself revolutionized by the work of Noam Chomsky and others, beginning in the late 1950s. The primary thrust of Chomsky's work was and still is in the syntax of natural languages. What exactly are the rules by which humans construct sentences in the various natural languages? And what are the rules that are common to all natural languages, the rules of 'universal grammar'? But philosophers were always more interested in semantics and in pragmatics, than they were in syntax. Semantics, on a standard definition, deals with the truth conditions of sentences: under what conditions is a sentence true or false? And pragmatics deals with the use of sentences in actual human situations, the use of sentences to give orders, make statements, give promises, etc. It seemed to a number of philosophers of language, myself included, that we should attempt to achieve a unification of Chomsky's syntax with the results of the researches that were going on in semantics and pragmatics. I believe that this effort has proven to be a failure. Though Chomsky did indeed revolutionize the subject of linguistics, it is not at all clear, at the end of the century, what the solid results of this revolution are. As far as I can tell there is not a single rule of syntax that all, or even most, competent linguists are prepared to agree is a rule.

In the middle years of the century, in the decades after the Second World War, optimism about using logic and language as the primary tools of philosophy ran at its highest. And indeed it seems to me that a great deal of progress was in fact achieved in those decades. Much of the optimism and self-confidence of the period derived from the belief in two linguistic distinctions. These are (i) the distinction between analytical and synthetic propositions, and (ii) the distinction between descriptive and evaluative utterances. If you accept these two distinctions in their pure forms—and many philosophers did—they seem to define the nature of philosophy and to determine its specific research programme. The first distinction between analytic and synthetic statements is between, on the one hand, those statements that are true or false by definition, such as statements in logic and mathematics and such commonsense tautologies as 'All bachelors are unmarried'; and, on the other hand, statements that are true or false as a matter of fact in the world, such as the statements of the natural sciences and such statements about contingent facts in the world as, for example, 'Most bachelors drink beer'. The second distinction, between descriptive and evaluative statements, is between those statements that describe states of affairs in the world and thus can be literally true or false, and those that serve to express our feelings, attitudes and evaluations, and, thus, according to the theory, cannot be literally true or false. An example of a descriptive statement would be 'The

incidents of crimes of violence have decreased in the last decade'; an example of an evaluative statement would be 'It is wrong to commit a crime'. The descriptive class includes both analytic and synthetic statements. According to those who accepted this theory, statements in the sciences and mathematics are descriptive because they describe matters of objective fact; whereas statements in ethics and aesthetics are evaluative because they are used to express feelings and attitudes, and to guide behaviour rather than to state facts. For those who accepted these distinctions, and they were the mainstream views in the middle decades of the century, the distinctions defined the nature of philosophy. Philosophers aimed at the truth and so were not in the business of making evaluations or value judgements of any kind. Telling people how to live is not the job of the professional philosopher. But the truths of philosophy are not contingent synthetic truths of the sort one finds in the natural sciences either. They are necessary analytic truths about concepts. The philosopher's task, like that of the logician and mathematician, is to state necessary analytic truths. His truths are conceptual, explicating puzzling philosophical concepts such as causation, knowledge, justice or truth itself.

Philosophy so construed is defined as conceptual analysis, and much of the optimism of the middle decades of the century derived from the conviction that philosophy now had a well-defined research project and well-defined methods for achieving results.

Confidence in these distinctions is now seriously weakened. Language does not seem so neat or simple that we can divide utterances into these simple categories of analytic and synthetic, descriptive and evaluative. In part because of a loss of confidence in the adequacy of these distinctions, the general optimism that we might solve all, or even most, philosophical problems using the methods of conceptual analysis has now abated. The upshot is that philosophy is less self-confident than it was in the 1950s and 1960s, but at the same time, it is much more interesting. All sorts of questions that were not regarded as really possible philosophical questions in the heyday of language analysis have now become possible, and I will say something more about these shortly. But at the same time there is less confidence about the possibility of getting definitive solutions to traditional philosophical problems using the methods of linguistic analysis.

There is another important development of 20th-century philosophy that I am less confident about, but which in the end may be its most important result. For the three centuries after Descartes, from the middle 17th to the late 20th century, the single greatest preoccupation of philosophers was with the problems of knowledge and scepticism. Descartes made epistemology—the theory of knowledge—central to philosophy. For Descartes the primary question was what sort of solid foundational grounding can we give to our claims to knowledge, in the sciences, in common sense, in religion, in mathematics, etc. Subsequent great philosophers, such as Locke, Berkeley, Hume, Leibniz, Spinoza and Kant, felt that Descartes' attempt to answer scepticism was inadequate, but Descartes' problem remained uppermost in their philosophical work. Locke, for example, took the main question of philosophy to be 'What is the nature and

extent of human knowledge?' Hume ended up with a much more radical scepticism than Descartes ever envisaged, but he felt that we could live with this scepticism by adopting a completely naturalistic attitude toward ourselves and the world. We just have to accept the limitations of our knowledge, recognize that we do not really know very much and go on as if we did know a great deal, even though we can offer no justification for the assumptions we make about the world. Kant read Hume and felt that it awakened him from his 'dogmatic slumber'. He made a heroic effort to overcome Hume's scepticism, but it too, I think, was a failure. In the 20th century, as I have said, the primary interest of philosophers was in language and meaning, and not with knowledge and its justification. In short, Descartes' question was 'How do you know?' and later Russell and Moore turned that into the question 'What do you mean?' Nonetheless, philosophers such as Russell, Moore, Wittgenstein and Austin devoted a great deal of intellectual effort to trying to overcome scepticism using linguistic methods. Though the primary thrust of their analyses was on language and meaning, much of the point of the analysis of language and meaning was to explain and justify the notions of truth, evidence and knowledge. I believe, and I sincerely hope, that this whole epoch has finally come to an end. Of course in philosophy nothing is ever finished once and for all, but my interpretation of the present intellectual scene, and my hope for the next century, is that we may simply relinquish our obsession with scepticism and get on with the more constructive aspects of philosophy.

The obsession with epistemology, and its endemic obsession with overcoming scepticism, led to a second feature of philosophy in the three centuries after Descartes. For many philosophers real progress required logical reduction. To understand a phenomenon we had to reduce it to simpler phenomena in the sense that we had to show how statements about the puzzling phenomenon could be logically derived from statements about epistemically simpler and more primitive phenomena. Thus many empirically minded philosophers thought that the only way to understand human mental states was to reduce them to behaviour (behaviourism). Analogously many philosophers thought that in order to understand empirical reality we had to reduce it to sensory experiences (phenomenalism). A natural consequence of the obsession with epistemology was to see the solution to the sceptical problem in reductionism. So there were, in my view, twin errors that pervaded philosophy and which I hope we have now overcome. These are scepticism and an inappropriate extension of reductionism.

I cannot overestimate the extent to which the epistemic bias has infected the practice of philosophy for nearly 400 years. Even in subjects that would appear to have only a fairly remote connection with epistemology, the epistemic question became central to the entire subject. This was nowhere more obvious than in the case of ethics and political philosophy. You might think that the question 'How do we know?' would not figure large in these disciplines, but in fact the central question of ethics in this epoch has been 'How can we have objective knowledge in ethics? How can we get the kind of epistemic certainty in our ethical judgements that we strive for in our scientific

judgements?' It did not even seem possible to our philosophical parents and grandparents that there could be a more fundamental question in ethics than this one. Indeed, for those who accepted the distinction between descriptive and evaluative, the result of philosophical analysis of ethical discourse was sceptical. According to this view it is impossible to have objective knowledge in ethics because ethical statements cannot be objectively true or false. A similar epistemic bias affected political philosophy. The question again was 'How can we be certain, how can we have epistemic objectivity about our political judgements and our claims of political obligation?' Just as ethics was afflicted with a form of scepticism, so political philosophy fell into the doldrums because of the same sort of scepticism. Political philosophy was revolutionized and revitalized by the publication of John Rawls's *A theory of justice* (1972), about which I will say more later.

Nowhere was the epistemic bias more blatant than in the philosophy of language. Frege did not have primarily epistemic worries about meaning, but his followers in the 20th century turned questions about meanings into questions about knowledge of meanings. This was, in my view, a disastrous error, but it is an error that continues to this day. There is an entire movement in the philosophy of language that thinks the central question is: What sort of evidence does a hearer have when he attributes meaning to a speaker of a language? What sort of evidence do I have that when you utter the word 'rabbit' you mean what I mean by 'rabbit', for example? And the answer to this question is, again in my view mistakenly, taken to be not merely an epistemic point about how we decide questions of meaning, but the key to understanding the very nature of meaning. Meaning is analysed completely into the sorts of evidence that hearers can have about what the speaker means. Many influential philosophers have thought that the epistemic question already gave us an answer to the ontological question, that the facts about meaning were entirely constituted by the evidence we could have about meaning. I believe this view is as mistaken in the philosophy of language as it is in the sciences and philosophy generally. It is as if knowledge in physics were supposed to be knowledge entirely of experiments and meter readings, since we use experiments and meter readings to test our knowledge of the physical universe. Analogously, it is equally a mistake to suppose that facts about meaning are facts about circumstances in which people utter expressions, since we use circumstances in which they utter expressions as evidence to make judgements about what they mean. I believe this epistemic bias is nothing less than the philosophical error of our epoch, and I will have more to say about it in the next section.

I have a specific intellectual objective in making the proposal that we should abandon scepticism and reductionism. I believe we cannot get a satisfactory constructive analysis of language, mind, society, rationality, political justice, etc., until we abandon our obsession with the idea that the presupposition of all investigation is first to provide a justification for the very possibility of knowledge, and that real advances in philosophical knowledge in general require the reduction of higher-level phenomena to more epistemically fundamental

phenomena. The way to deal with scepticism is not to try to refute it on its own terms, but to overcome it in such a way that we can go on to deal with the problems at hand. As I said earlier, I am not certain that this is where we are, but it certainly is where I am in my own intellectual development. On my interpretation of the contemporary philosophical scene, scepticism has finally ceased to be a primary concern of philosophers, and reductionism has in general failed. The situation we are in is somewhat analogous to the situation of the Greeks at the time of the transition from Socrates and Plato to Aristotle. Socrates and Plato took scepticism very seriously and struggled with piecemeal issues. Aristotle did not regard the sceptical paradoxes as a serious threat to his overall enterprise of attempting to do systematic, constructive, theoretical philosophy. I think we now have the tools to move into a 21st-century version of an Aristotelian phase. Wittgenstein, one of the most important philosophers of the 20th century thought that general theories in philosophy were impossible. Paradoxically, by helping to clear the field of sceptical worries Wittgenstein did as much as anybody to make general philosophical theories possible.

3. SIX PROBLEM AREAS

Because of the nature of the subject, I do not believe it is possible to project a future course of philosophy with anything like the confidence that one can project the future course of the sciences—though, of course, that is not at all an easy thing to do in itself. What I will, therefore, do here is take about a half a dozen areas of philosophical investigation which are very much alive at the present moment, and discuss their present status and future prospects. In some cases I feel confident enough to make some guesses about what I think will happen, in others I can only make critical remarks and expressions of hope for future research in the coming decades.

(a) *The traditional mind-body problem*

I begin with the traditional mind-body problem, because I believe it is the contemporary philosophical problem most amenable to scientific solution: What exactly are the relations between consciousness and the brain? It seems to me the neurosciences have now progressed to the point that we can address this as a straight neurobiological problem, and indeed several neurobiologists are doing precisely that. In its simplest form, the question is how exactly do neurobiological processes in the brain cause conscious states and processes, and how exactly are those conscious states and processes realized in the brain?

So stated, this looks like an empirical scientific problem. It looks similar to such problems as 'How exactly do biochemical processes at the level of cells cause cancer?' and 'How exactly does the genetic structure of a zygote produce the phenotypical traits of a mature organism?'

However, there are a number of purely philosophical obstacles to getting a satisfactory neurobiological solution to the problem of consciousness, and I have to devote some space at least to trying to remove some of the worst of these obstacles.

The single most important obstacle to getting a solution to the traditional mind-brain problem is the

persistence of a set of traditional but obsolete categories of mind and body, matter and spirit, mental and physical. As long as we continue to talk and think as if the mental and the physical were separate metaphysical realms, the relation of the brain to consciousness will forever seem mysterious, and we will not have a satisfactory explanation of the relation of neuron firings to consciousness. The first step on the road to philosophical and scientific progress in these areas is to forget about the tradition of Cartesian dualism and just remind ourselves that mental phenomena are ordinary biological phenomena in the same sense as photosynthesis or digestion. We must stop worrying about how the brain *could* cause consciousness and begin with the plain fact that it *does*. The notions of both mental and physical as they are traditionally defined need to be abandoned, as we reconcile ourselves to the fact that we live in one world, and all the features of the world, from quarks and electrons to nation-states and balance of payments problems are, in their different ways, part of that one world. I find it truly amazing that the obsolete categories of mind and matter continue to impede progress. Many scientists feel that they can only investigate the 'physical' realm and are reluctant to face consciousness on its own terms because it seems not to be physical but to be 'mental', and several prominent philosophers think it is impossible for us to understand the relations of mind to brain. Just as Einstein made a conceptual change to break the distinction between space and time, so we need a similar conceptual change to break the bifurcation of mental and physical.

Related to the difficulty brought about by accepting the traditional categories is a straight logical fallacy that I need to expose. Consciousness is, by definition, subjective, in the sense that for a conscious state to exist it has to be experienced by some conscious subject. Consciousness in this sense has a first-person ontology in that it only exists from the point of view of a human or animal subject, an 'I', who has the conscious experience. Science is not used to dealing with phenomena that have a first-person ontology. By tradition, science deals with phenomena that are 'objective', and avoids anything that is 'subjective'. Indeed, many philosophers and scientists feel that because science is, by definition, objective, there can be no such thing as a science of consciousness, because consciousness is subjective. This whole argument rests on a massive confusion, which is one of the most persistent confusions in our intellectual civilization. There are two quite distinct senses of the distinction between objective and subjective. In one sense, which I will call the epistemological sense, there is a distinction between objective knowledge and subjective matters of opinion. If I say, for example, 'Rembrandt was born in 1606', that statement is epistemically objective in the sense that it can be established as true or false independently of the attitudes, feelings, opinions or prejudices of the agents investigating the question. If I say 'Rembrandt was a better painter than Rubens', that claim is not a matter of objective knowledge, but is a matter of subjective opinion. But in addition to the distinction between epistemically objective and subjective claims, there is a distinction between entities in the world that have an objective existence, such as mountains and molecules, and entities that have a subjective existence, such as pains and tickles. I call this distinction

in modes of existence, the ontological sense of the objective–subjective distinction.

Science is indeed epistemically objective in the sense that scientists attempt to establish truths that can be verified independently of the attitudes and prejudices of the scientists. But epistemic objectivity of method does not preclude ontological subjectivity of subject matter. Thus there is no objection in principle to having an epistemically objective science of an ontologically subjective domain, such as human consciousness.

Another difficulty encountered by a science of subjectivity is the difficulty in verifying claims about human and animal consciousness. In the case of humans, unless we perform experiments on ourselves individually, our only conclusive evidence for the presence and nature of consciousness is what the subject says and does, and subjects are notoriously unreliable. In the case of animals, we are in an even worse situation, because we have to rely on just the animal's behaviour in response to stimuli. We cannot get any statements from the animal about its conscious states.

I think this is a real difficulty, but I would point out that it is no more an obstacle in principle than the difficulties encountered in other forms of scientific investigation where we have to rely on indirect means of verifying our claims. We have no way of observing black holes, and indeed, strictly speaking, we have no way of directly observing atomic and subatomic particles. Nonetheless, we have quite well-established scientific accounts of these domains, and the difficulties in verifying hypotheses in these areas should give us a model for verifying hypotheses in the area of the study of human and animal subjectivity. The 'privacy' of human and animal consciousness does not make a science of consciousness impossible. As far as 'methodology' is concerned, in real sciences methodological questions always have the same answer. To find out how the world works, you have to use any weapon you can lay your hands on, and stick with any weapon that seems to work.

Assuming, then, that we are not worried about the problem of objectivity and subjectivity, and that we are prepared to seek indirect methods of verification of hypotheses concerning consciousness, how should we proceed? Most scientific research today into the problem of consciousness seems to me to be based on a mistake. The scientists in question characteristically adopt what I will call the building-block theory of consciousness, and they conduct their investigation accordingly. On the building-block theory, we should think of our conscious field as made up of various building blocks, such as visual experience, auditory experience, tactile experience, the stream of thought, etc. The task of a scientific theory of consciousness would be to find the neurobiological correlate of consciousness (nowadays called the NCC) and, on the building-block theory, if we could find the NCC for even one building block, such as the NCC for colour vision, that would in all likelihood give us a clue to the building blocks for the other sensory modalities and for the stream of thought. This research programme may turn out to be right in the end. Nonetheless, it seems to me doubtful as a way to proceed in the present situation, for the following reason. I said above that the essence of consciousness was subjectivity. There is a certain

subjective, qualitative feel to every conscious state. One aspect of this subjectivity, and it is a necessary aspect, is that conscious states always come to us in a unified form. We do not perceive just the colour or the shape, or the sound, of an object, we perceive all of these simultaneously in a unified, conscious experience. The subjectivity of consciousness implies unity. They are not two separate features, but two aspects of the same feature.

Now, that being the case, it seems to me the NCC we are looking for is not the NCC for the various building blocks of colour, taste, sound, etc., but rather what I will call the basal, or background, conscious field, which is the presupposition of having any conscious experience in the first place. We should think of my present conscious field not as made up of various building blocks, but rather as a unified field, which is modified in specific ways by the various stimuli that I and other human beings receive. Because we have pretty good evidence from lesion studies that consciousness is not distributed over the entire brain, and because we also have good evidence that consciousness exists in both hemispheres, I think what we should look for now is the kind of neurobiological processes that will produce a unified field of consciousness. These, as far as I can tell, are likely to be for the most part in the thalamocortical system. My hypothesis, then, is that looking for the NCCs of building blocks is barking up the wrong tree, and that we should instead look for the correlate of the unified field of consciousness in massive synchronized patterns of neuron firing.

(b) *The philosophy of mind and cognitive science*

The mind-body problem is one part of a much broader set of issues, known collectively as the philosophy of mind. This includes not only the traditional mind-body problem, but the whole conglomeration of problems dealing with the nature of mind and consciousness, of perception and intentionality of intentional action and thought. A very curious thing has happened in the past two or three decades—the philosophy of mind has moved to the centre of philosophy. Several other important branches of philosophy, such as epistemology, metaphysics, the philosophy of action and even the philosophy of language, are now treated as dependent on, and in some cases even as branches of, the philosophy of mind. Whereas 50 years ago the philosophy of language was considered ‘first philosophy’, now it is the philosophy of mind. There are a number of reasons for this change, but two stand out.

First, it has become more and more obvious to a lot of philosophers that our understanding of the issues in a lot of subjects—the nature of meaning, rationality and language in general—presupposes an understanding of the most fundamental mental processes. For example, the way language represents reality is dependent on the more biologically fundamental ways in which the mind represents reality and, indeed, linguistic representation is a vastly more powerful extension of the more basic mental representations such as beliefs, desires and intentions.

Second, the rise of the new discipline of cognitive science has opened to philosophy whole areas of research into human cognition in all its forms. Cognitive science was invented by an interdisciplinary group, consisting of

philosophers who objected to the persistence of behaviourism in psychology, together with like-minded cognitive psychologists, linguists, anthropologists and computer scientists. I believe the most active and fruitful general area of research today in philosophy is in the general cognitive science domain.

The basic subject matter of cognitive science is intentionality in all of its forms. ‘Intentionality’ is a technical term used by philosophers to refer to all of those mental phenomena that refer to, or are about, objects and states of affairs in the world. ‘Intentionality’ so defined has no special connection with intending in the ordinary sense in which I intend to go to the movies tonight. Intending is just one kind of intentionality among others. Intentionality so defined includes at least beliefs, desires, memories, perceptions, intentions (in the ordinary sense), intentional actions and emotions.

Paradoxically, cognitive science was founded on a mistake. There is nothing necessarily fatal about founding an academic subject on a mistake, indeed many disciplines were founded on mistakes. Chemistry, for example, was founded on alchemy. However, a persistent adherence to the mistake is at best inefficient and an obstacle to progress. In the case of cognitive science the mistake was to suppose that the brain is a digital computer and the mind is a computer program.

There are a number of ways to demonstrate that this is a mistake, but the simplest is to point out that the implemented computer program is defined entirely in terms of symbolic or syntactical processes, independent of the physics of the hardware. Minds, on the other hand, contain more than symbolic or syntactical components, they contain actual mental states with semantic content in the form of thoughts, feelings, etc., and these are caused by quite specific neurobiological processes in the brain. The mind could not consist in a program because the syntactical operations of the program are not by themselves sufficient to guarantee the semantic contents of actual mental processes. I demonstrated this years ago with the so-called Chinese Room Argument.

A debate continues about this and other versions of the computational theory of the mind. Some people think that the introduction of computers that use parallel distributed processing (PDP), sometimes also called ‘connectionism’, would answer the objections I just stated. But I do not see how the introduction of the connectionist arguments makes any difference. The problem is that any computation that can be carried out on a connectionist program can also be carried out on a traditional Von Neumann system. We know from mathematical results that any function that is computable at all is computable on a universal Turing machine. In that sense no new computational capacity is added by the connectionist architecture, though the connectionist systems can be made to work faster, because they have several different computational processes acting in parallel and interacting with each other. Because the computational powers of the connectionist system are no greater than the traditional Von Neumann system, if we claim superiority for the connectionist system, there must be some other feature of the system that is being appealed to. But the only other feature of the connectionist system would have to be in the hardware implementation, which operates in parallel

rather than in series. But if we claim that the connectionist architecture rather than connectionist computations are responsible for mental processes, we are no longer advancing the computational theory of the mind, but are engaging in neurobiological speculation. With this hypothesis we have abandoned the computational theory of the mind in favour of speculative neurobiology.

What is actually happening in cognitive science is a paradigm shift away from the computational model of the mind and towards a much more neurobiologically based conception of the mind. For reasons that should be clear by now, I welcome this development. As we come to understand more about the operations of the brain it seems to me that we will succeed in gradually replacing computational cognitive science with cognitive neuroscience. Indeed I believe this transformation is already taking place.

Advances in cognitive neuroscience are likely to create more philosophical problems than they solve. For example, to what extent will an increased understanding of brain operations force us to make conceptual revisions in our commonsense vocabulary for describing mental processes as they occur in thought and action? In the simplest and easiest cases we can simply assimilate the cognitive neuroscience discoveries to our existing conceptual apparatus. Thus, we do not make a major shift in our concept of memory when we introduce the sorts of distinctions that neurobiological investigation has made apparent to us. We now in popular speech distinguish between short-term and long-term memory, and no doubt as our investigation proceeds, we will have further distinctions. Perhaps the concept of iconic memory is already passing into the general speech of educated people. But in some cases it seems we are forced to make conceptual revisions. I have thought for a long time that the commonsense conception of memory as a storehouse of previous experience and knowledge is both psychologically and biologically inadequate. My impression is that contemporary research bears me out on this. We have to have a conception of memory as a creative process rather than simply a retrieval process. Some philosophers think even more radical revisions than this will be forced upon us by the neurobiological discoveries of the future.

Another set of philosophical problems arises when we begin to examine the relationships between the developmental evidence regarding mental phenomena and the mental phenomena as they occur in mature adults. Very young children apparently have a different conception of the relation of belief to truth from that which adults have. How seriously should we take these differences? Do we need to enrich our theory of intentionality by incorporating the developmental data? We do not yet know the answer to any of these questions, and my point in raising them here is to call attention to the fact that once we have removed the philosophical error of supposing that the brain is a digital computer, and once we have a more mature and sophisticated cognitive neuroscience, we still have to deal with a number of philosophical questions.

(c) *The philosophy of language*

I said that the philosophy of language was the centre of philosophy for most of the 20th century. Indeed, as I

remarked, during the first three-quarters of the 20th century, the philosophy of language was taken to be 'first philosophy'. But as the century comes to an end that is changing. Less is happening in the philosophy of language now than in the philosophy of mind, for example, and I believe that some of the currently most influential research programmes have reached a kind of dead end. Why? There are many reasons of which I will mention only three.

First, the more successful branches of the philosophy of language are now passing into the science of linguistics. The sort of research that I and others did 30 years ago on the theory of speech acts and on the use of language is now becoming a part of linguistics called 'pragmatics', which has its own corner in linguistics, with its own journals, annual meetings, etc. In short, this part of the philosophy of language is gradually being kicked out of philosophy, upstairs into the social sciences. I welcome this development, and I believe that it is an example of the sort of phenomenon that I described in the early part of this article, where I explained that as areas of investigation arrive at established methodologies for their research, they tend to be thought of as more scientific and less philosophical.

Second, one of the main research programmes in the philosophy of language suffers from the epistemic obsession that I have been castigating. A commitment to a certain form of empiricism, and in some cases even behaviourism, led some prominent philosophers to try to give an analysis of meaning according to which the hearer is engaged in an epistemic task of trying to figure out what the speaker means either by looking at his behaviour in response to stimulus or by looking at the conditions under which he would hold a sentence to be true. The idea is that if we could describe how the hearer solves the epistemic problem we will thereby analyse meaning.

This work, I believe, is going nowhere, because its obsession with how we know what a speaker means obscures the distinction between *how* the hearer knows what the speaker means and *what* it is that the hearer knows. I think that epistemology plays the same role in the philosophy of language as it does, for example, in geology. The geologist is interested in such things as tectonic plates, sedimentation and rock layers, and will use any method that comes to hand to try to find out how these phenomena work. The philosopher of language is interested in meaning, truth, reference and necessity, and analogously should use any epistemic method that comes to hand to try to figure out how these phenomena work in the minds of actual speakers and hearers. What we are interested in is what are the facts which are known; and to a much lesser extent are we interested in the question of how we come to know these facts.

Finally, I think the greatest source of weakness in the philosophy of language is that its currently most influential research project is based on a mistake. I said earlier that Frege was the founder of the philosophy of language, but Frege had a conception of meaning that placed the meanings of words inside the heads of the speakers of a language. Frege was anxious to insist that these meanings were not psychological entities, but he did think that they could be grasped by speakers and hearers of a language. Frege thought that communication in a public language

was possible only because there is an ontologically objective realm of meanings, and that the same meaning can be grasped equally by both speaker and hearer. A number of authors have attacked this conception. They believe that meaning is a matter of causal relations between the utterances of words and objects in the world. So the word 'water', for example, means what it does to me not because I have some mental content associated with that word, but rather because there is a causal chain connecting me to various actual examples of water in the world. This view is called 'externalism', and it is usually opposed to the traditional view, called 'internalism'. Externalism has led to an extensive research project of trying to describe the nature of the causal relations that give rise to meaning. The problem with this research project is that nobody has ever been able to explain, with any plausibility whatever, the nature of these causal chains. The idea that meanings are something external to the mind is widely accepted, but no one has ever been able to give a coherent account of meaning in these terms.

My prediction is that no one will ever be able to give a satisfactory account of meanings as something external to the head, because such external phenomena could not function to relate language to the world in the way that meanings do relate words and reality. What we require in order to resolve the dispute between internalists and externalists is a more sophisticated notion of how the mental contents in speakers' heads serve to relate language in particular, and human agents in general, to the real world of objects and states of affairs.

(d) *The philosophy of society*

It is characteristic of the history of philosophy that new branches of the subject are created in response to intellectual developments both inside and outside of philosophy. Thus, for example, in the early part of the 20th century the philosophy of language in the sense in which we now use that expression, was created largely in response to developments in mathematical logic and work on the foundations of mathematics. A similar evolution has occurred in the philosophy of mind. I would like to propose that in the 21st century we will feel a pressing need for, and should certainly try to develop, what I will call a philosophy of society. It is characteristic of the social parts of philosophy that we tend to construe social philosophy as either a branch of political philosophy, thus the expression 'social and political philosophy', or we tend to construe social philosophy as a study of the philosophy of the social sciences, just as the philosophy of natural sciences is a branch of the philosophy of science. I am proposing that we should have a social philosophy, which stands to social sciences in the same way that the philosophy of mind stands to psychology and cognitive science, or the philosophy of language stands to linguistics. It would deal with more general framework questions. In particular, I think we need much more work on questions of the ontology of social reality. How is it possible that human beings, through their social interactions can create an objective social reality of money, property, marriage, government, wars, games, etc., when such entities in some sense exist only by virtue of a collec-

tive agreement or a belief that they exist? How is it possible that there can be an objective social reality that exists only because we think it exists?

When questions of social ontology have been properly sorted out it seems to me that the questions of social philosophy, namely the nature of explanation in the social sciences and the relation of social philosophy to political philosophy, will naturally fall into place. I attempted to begin this research project in *The construction of social reality* (Searle 1995).

Specifically, I believe that in our study of political and social reality, we need a set of concepts that will enable us to describe political and social reality, so to speak from the 'middle distance'. The problem that we have in attempting to cope with social reality is that our concepts are either immensely abstract, as in traditional political philosophy, for example the concepts of the social contract or the class struggle; or they tend to be essentially journalistic, dealing with day-to-day questions of policy and power relations. Thus we are quite sophisticated in abstract theories of justice, and with developing criteria for assessing the justice or injustice of institutions. Much of the progress in this area is owed to John Rawls, who, as I mentioned, revolutionized the study of political philosophy with his classic work *A theory of justice* (Rawls 1972). But when it comes to political science, the categories traditionally do not rise much above the level of journalism. Therefore, if, for example, you read a work in political science as recent as 20 years old, you will find that much of the discussion is out of date.

What we need, I believe, is to develop a set of categories that would enable us to appraise social reality in a way which would be more abstract than that of day-to-day political journalism, but at the same time would enable us to ask and answer specific questions about specific political realities and institutions in a way that traditional political philosophy was unable to do. Thus, for example, I think the leading political event of the 20th century was the failure of ideologies such as those of Fascism and communism, and in particular the failure of socialism in its different and various forms. The interesting thing from the point of view of the present analysis is that we lack the categories in which to pose and answer questions dealing with the failure of socialism. If by 'socialism' we mean state ownership and control of the basic means of production, then the failure of socialism so defined is the single most important social development of the 20th century. It is an amazing fact that that development remains unanalysed and is seldom discussed by the political and social philosophers of our time.

When I talk of the failure of socialism, I am referring not only to the failure of Marxist socialism, but the failure of democratic socialism as it existed in the countries of Western Europe. The socialist parties of those countries continued to use the vocabulary of socialism, but the belief in the basic mechanism of socialist change, namely the public ownership and control of the means of production, apparently has been quietly abandoned. What is the correct philosophical analysis of this entire phenomenon?

A similar sort of question would be the appraisal of national institutions. So, for example, for most political scientists it would be very difficult to attempt to analyse

the backwardness, corruption and general dreadfulness of the political institutions of several contemporary nation states. Most political scientists, given their commitment to scientific objectivity, and the limited categories at their disposal, cannot even attempt to describe how dreadful many countries are. Many countries have apparently desirable political institutions, such as a written constitution, political parties, free elections, etc., and yet the way they operate is inherently corrupt. We can discuss these institutions at a very abstract level, and Rawls and others have provided us with the tools to do so. But I would like an expanded social philosophy which would provide us with the tools for analysing social institutions as they exist in real societies, in a way that would enable us to make comparative judgements between different countries and larger societies without, at the same time, rising to such a level of abstraction that we cannot make specific value judgements about specific institutional structures. The work of the economist-philosopher Amartya Sen is a step in this direction.

(e) *Ethics and practical reason*

For much of the 20th century the subject of ethics was dominated by a version of the same scepticism that has affected other branches of philosophy for several centuries. Just as the philosophy of language was damaged by the urge to treat the users of language as essentially researchers engaged in an epistemic task of trying to figure out what a speaker of a language means, so ethics was obsessed by the question of objectivity. The principal issue in ethics was about whether or not there could be epistemic objectivity in ethics. The traditional view in analytic philosophy was that ethical objectivity was impossible, that you could not in Hume's phrase derive an 'ought' from an 'is', and consequently ethical statements could not literally be either true or false, but functioned only to express feelings or to try to influence behaviour, etc. The way out of the sterility of these debates is not, I think, to try to show that ethical statements are true or false in a way that, for example, scientific statements are true or false, because there are clearly important differences between the two. The way out of the impasse, I believe, is to see that ethics is really a branch of a much more interesting subject of practical reason and rationality. What is the nature of rationality in general and what is it to act rationally on a reason for an action? This, I believe, is a more fruitful approach than the traditional approach of worrying about the objectivity of ethical statements.

Something like the study of rationality, as a successor to ethics as it was traditionally construed, seems to be already happening. Currently there are, for example, a number of attempts to revive Kant's doctrine of the categorical imperative. Kant thought that the nature of rationality itself set certain formal constraints on what could count as an ethically acceptable reason for an action. I do not believe these efforts will succeed, but much more interesting than their success or failure is the fact that ethics as a substantive branch of philosophy—freed from its epistemic obsession to find a form of objectivity and the inevitable scepticism when the quest for objectivity fails—seems now to have become possible again. I am not sure what the reasons for the change are,

but my impression is that, more than any other single factor, Rawls's work not only revived political philosophy but made substantive ethics seem possible as well.

(f) *The philosophy of science*

In the 20th century, not surprisingly, the philosophy of science shared the epistemic obsession with the rest of philosophy. The chief questions in the philosophy of science, at least for the first half of the century, had to do with the nature of scientific verification, and much effort was devoted to overcoming various sceptical paradoxes, such as the traditional problem of induction. Throughout most of the 20th century the philosophy of science was conditioned by the belief in the distinction between analytic and synthetic propositions. The standard conception of the philosophy of science was that scientists aimed to get synthetic contingent truths in the form of universal scientific laws. These laws stated very general truths about the nature of reality, and the chief issue in the philosophy of science had to do with the nature of their testing and verification. The prevailing orthodoxy, as it developed in the middle decades of the century, was that science proceeded by something called the 'hypothetico-deductive method'. The scientists formed the hypothesis, deduced logical consequences from it, and then tested those consequences in the form of experiments. This conception was articulated, I think more or less independently, by Karl Popper and Carl Gustav Hempel.

Those practising scientists who took an interest in the philosophy of science at all, tended, I think, to admire Popper's views, but much of their admiration was based on a misunderstanding. What I think they admired in Popper was the idea that science proceeds by acts of originality and imagination. The scientist has to form a hypothesis on the basis of his own imagination and guesswork. There is no 'scientific method' for arriving at hypotheses. The procedure of the scientist is then to test the hypothesis by performing experiments and reject those hypotheses that have been refuted.

Most scientists do not, I think, realize how anti-scientific Popper's views actually are. On Popper's conception of science and the activity of scientists, science is not an accumulation of truths about nature, and the scientist does not arrive at truths about nature, rather, all that we have in the sciences are a series of so far unrefuted hypotheses. But the idea that the scientist aims after truth, and that in various sciences we actually have an accumulation of truths, which I think is the presupposition of most actual scientific research, is not something that is consistent with Popper's conception.

The comfortable orthodoxy of science as an accumulation of truths, or even as a gradual progression through the accumulation of so far unrefuted hypotheses, was challenged by the publication of Thomas Kuhn's *Structure of scientific revolutions* in 1962. It is puzzling that Kuhn's book should have had the dramatic effect that it did, because it is not strictly speaking about the philosophy of science, but about the history of science. Kuhn argues that if you look at the actual history of science, you discover that it is not a gradual progressive accumulation of knowledge about the world, but that science is subject to periodic massive revolutions, where entire world views are overthrown when an existing paradigm is overthrown

by a new scientific paradigm. It is characteristic of Kuhn's book that he implies, though as far as I know he does not state explicitly, that the scientist does not give us truths about the world, but gives us a series of ways of solving puzzles, a series of ways of dealing with puzzling problems within a paradigm. And when the paradigm reaches puzzles that it cannot solve, it is overthrown and a new paradigm is erected in its place, which again sets off a new round of puzzle-solving activity. From the point of view of this discussion, the interesting thing about Kuhn's book is that he seems to imply that we are not getting progressively closer to the truth about nature in the natural sciences, we are just getting a series of puzzle-solving mechanisms. The scientist essentially moves from one paradigm to another, for reasons that have nothing to do with giving an accurate description of an independently existing natural reality, but rather for reasons that are in greater or lesser degree irrational. Kuhn's book was not much welcomed by practising scientists, but it had an enormous effect on several humanities disciplines, especially those connected with the study of literature, because it seemed to argue that science gives us no more truth about the real world than do works of literary fiction or literary criticism; that science is essentially an irrational operation where groups of scientists form theories which are more or less arbitrary social constructs, and then abandon these in favour of other theories, which are likewise arbitrary social constructs.

Whatever Kuhn's intentions, I believe that his effect on general culture, though not on the practices of real scientists, has been unfortunate, because it has served to 'demythologize' science, to 'debunk' it, to prove that it is not what ordinary people have supposed it to be. Kuhn paved the way for the even more radical sceptical view of Paul Feyerabend, who argued that as far as giving us truths about the world is concerned, science is no better than witchcraft.

My own view is that these issues are entirely peripheral to what we ought to be worried about in the philosophy of science, and what I hope we will dedicate our efforts to in the 21st century. I think the essential problem is this: 20th-century science has radically challenged a set of very pervasive, powerful philosophical and common sense assumptions about nature, and we simply have not digested the results of these scientific advances. I am thinking especially of quantum mechanics. I think that we can absorb relativity theory more or less comfortably because it can be construed as an extension of our traditional Newtonian conception of the world. We simply have to revise our ideas of space and time, and their relation to such fundamental physical constants as the speed of light. But quantum mechanics really does provide a basic challenge to our world view, and we simply have not yet digested it. I regard it as a scandal that philosophers of science, including physicists with an interest in the philosophy of science, have not so far given us a coherent account of how quantum mechanics fits into our overall conception of the universe, particularly as regards to causation and determinacy.

Most philosophers, like most educated people today, have a conception of causation that is a mixture of common sense and Newtonian mechanics. Philosophers tend to suppose that causal relations are always instances

of strict deterministic causal laws, and that cause and effect relations stand to each other in the kind of simple mechanical relations of gear wheels moving other gear wheels, and other such Newtonian phenomena. We know at some abstract level that that is not right, but we still have not replaced our commonsense conception with a more sophisticated scientific conception. I think that the most exciting task of the 21st-century philosophy of science, and this is something for both scientists and philosophers, would be to give an account of the results of quantum mechanics to a coherent overall world view. I think that in the course of this project we are going to have to revise certain crucial notions, such as the notion of causation; and this revision is going to have very important effects on other questions, such as the questions concerning determinism and free will. This work has already begun, and I hope it will continue successfully in the 21st century.

4. CONCLUSION: OVERCOMING EPISTEMOLOGY

The history of philosophy, as it is described in the standard textbooks, is largely a history of the works of a number of towering geniuses. From Socrates, Plato and Aristotle, to Wittgenstein and Russell, the chief results of philosophy are in the works of its great figures. In that sense there simply are no towering geniuses alive today. This, I believe, is not because we have less talent than our predecessors. On the contrary, I believe that, paradoxically, the reason why there are no recognized geniuses today is simply that there are more good philosophers alive now than there were in the past. Because there is so much talent, and so much good work is being done, it is impossible for a single figure or a few figures to dominate the field in a way that was possible up until the early part of the 20th century. I think there are probably a number of other fields like philosophy in this respect—the apparent shortage of geniuses is the result of a surplus rather than a deficit of talent. But whether or not the phenomenon is general, I am quite confident that this is true of philosophy: the sheer number of hard-working, able, talented figures in the field makes it impossible for a small number of people to be recognized as standing head and shoulders above all the others.

One of the many advantages in having a field which is not dominated by a tiny number of overpowering figures, is that philosophy as a cooperative enterprise seems to be more possible than it has typically been in the past. It is quite possible for people working on a common set of problems to see their enterprise as one of advancing theoretical understanding in a given domain.

In my view, and it has been the theme of this article to expound that view, the biggest single obstacle to progress of a systematic theoretical kind has been the obsession with epistemology. I believe that epistemic problems, 'How is it possible that we can have knowledge at all in the light of the various sceptical paradoxes?', should be regarded in the same way as other such paradoxes have been regarded in the history of philosophy. Zeno's paradoxes about space and time, for example, pose interesting puzzles, but no one supposes that we cannot seriously attempt to cross a room until we have first answered Zeno's scepticism about the possibility of moving through

space. Analogously, I believe, we should have the same attitude towards the paradoxes about the possibility of knowledge that were advanced by sceptical philosophers. That is, these are interesting puzzles, and they provide good five-finger exercises for training young philosophers, but we should not suppose that the possibility of knowledge and understanding rests on our first being able to refute Hume's scepticism. I cannot, of course, predict what is going to happen in the 21st century, but I can express the hope, and I think at this stage in our intellectual history it is a well-founded hope, that with the abandonment of the epistemic bias in the philosophy of language, the philosophy of mind, ethics, political philosophy and the philosophy of science, we may achieve

greater theoretical understanding and more constructive theoretical accounts than we have had at any time in the past history of the subject.

REFERENCES

- Kuhn, K. 1962 *Structure of scientific revolutions*. University of Chicago Press.
Rawls, J. 1972 *A theory of justice*. Cambridge, MA: Harvard University Press.
Searle, J. R. 1980 Minds, brains and programs. *Behav. Brain Sci.* **3**, 417–457.
Searle, J. R. 1995 *The construction of social reality*. New York: The Free Press.