

PHILOSOPHY IN A NEW CENTURY

JOHN R. SEARLE

UNIVERSITY OF CALIFORNIA, BERKELEY

ABSTRACT: The central intellectual fact of the present era is that knowledge grows. This growth of knowledge is quietly transforming philosophy, making it possible to do a new kind of philosophy. With the abandonment of the epistemic bias in the subject, such a philosophy can go far beyond anything imagined by the philosophy of a half century ago. It begins, not with skepticism, but with what we know about the real world. It begins with such facts as those stated by the atomic theory of matter and the evolutionary theory of biology, as well as such “common sense” facts as that we are all conscious, that we all really do have intentional mental states, that we form social groups and create institutional facts. Such a philosophy is theoretical, comprehensive, systematic, and universal in subject matter.



General ruminations on the state and future of philosophy often produce superficiality and intellectual self-indulgence. Furthermore, an arbitrary blip on the calendar, the beginning of a new century, would not seem sufficient, by itself, to override a general presumption against engaging in such ruminations. However, I am going to take the risk of saying some things about the current and future state of philosophy, even though I think it is a serious risk. A number of important overall changes in the subject have occurred in my lifetime and I want to discuss their significance and the possibilities they raise for the future of the subject.

I. PHILOSOPHY AND KNOWLEDGE

The central intellectual fact of the present era is that knowledge grows. It grows daily and cumulatively. We know more than our grandparents did; our children will know more than we do.

We now have a huge accumulation of knowledge which is *certain, objective, and universal*, in a sense of these words that I will shortly explain. This growth of knowledge is quietly producing a transformation of philosophy.

The modern era in philosophy, begun by Descartes, Bacon, and others in the seventeenth century, was based on a premise which has now become obsolete. The premise was that the very existence of knowledge was in question and that therefore the main task of the philosopher was to cope with the problem of skepticism. Descartes saw his task as providing a secure foundation for knowledge, and Locke, in a similar vein, thought of his *Essay* as an investigation into the nature and extent of human knowledge. It seems reasonable that in the seventeenth century those philosophers took epistemology as the central element of the entire philosophical enterprise, because while they were in the midst of a scientific revolution, at the same time the possibility of certain, objective, universal knowledge seemed problematic. It was not at all clear how their various beliefs could be established with certainty, and it was not even clear how they could be made consistent. In particular there was a nagging and pervasive conflict between religious faith and the new scientific discoveries. The result was that we had three and a half centuries in which epistemology was at the center of philosophy.

During much of this period the skeptical paradoxes seemed to lie at the heart of the philosophical enterprise. Unless we can answer the skeptic, it seemed we cannot proceed further in philosophy or, for that matter, in science. For this reason epistemology became the base of any number of philosophical disciplines where it would seem that the epistemological questions are really peripheral. So, for example, in ethics the central question became, "Can there be an objective foundation for our ethical beliefs?" And even in the philosophy of language, many philosophers thought, and some still do, that epistemic questions were central. They take the central question in the philosophy of language to be, "How do we know what another person means when he says something?"

I believe the era of skeptical epistemology is now over. Because of the sheer growth of certain, objective, and universal knowledge, the possibility of knowledge is no longer a central question in philosophy. At present it is psychologically impossible for us to take Descartes's project seriously in the way that he took it: We know too much. This is not to say that there is no room for the traditional epistemic paradoxes, it simply means they no longer lie at the heart of the subject. The question, "How do I know that I am not a brain in a vat, not deceived by an evil demon, not dreaming, hallucinating, etc?"—or, in a more specifically Humean vein, "How do I know I am the same person today that I was yesterday," "How do I know that the sun will rise in the East tomorrow?" "How do I know that there really are such things as causal relations in the world?"—I regard as like Zeno's paradoxes about the reality of space and time. It is an interesting paradox how I can cross the room if first I have to cross half of the room, but before that, half of the half,

and yet before that, half of that half, etc. It seems I would have to traverse an infinite number of spaces before I can even get started and thus it looks like movement is impossible. That is an interesting paradox, and it is a nice exercise for philosophers to resolve the paradox, but no one seriously doubts the existence of space or the possibility of crossing the room because of Zeno's paradoxes. Analogously, I should like to say, no one should doubt the existence of knowledge because of the skeptical paradoxes. These are nice exercises for philosophers, but they do not challenge the existence of objective, universal and certain knowledge.

I realize that there is still a thriving industry of work on traditional skepticism. I am suggesting, however, that the traditional forms of skepticism cannot have the meaning for us that they had for Descartes and his successors. Whether we like it or not, the sheer weight of accumulated knowledge is now so great that we cannot take seriously arguments that attempt to prove that it doesn't exist at all.

One clarification I need to make immediately. When I say that philosophy is no longer about epistemology I mean that the professional paradoxes of epistemology, the skeptical paradoxes, are no longer central to the philosophical enterprise. But in addition to epistemology in this specialized professional sense, there is, so to speak, "real life" epistemology. How do you know that the claims you make are really true? What sorts of evidence, support, argument, and verification can you offer for the various claims you make? Real life epistemology continues as before, indeed, it is as important as ever, because, for example, in the face of competing real life claims about the cause and cure of AIDS, or the rival claims of monetary policy and fiscal policy in managing the economy, it is as important as ever that we insist on adequate tests and verification. So when I say that we are in a post-epistemic era, I mean we are in a post-skeptical era. Traditional philosophical skepticism I regard as now obsolete. But that does not mean we should abandon rational standards for assessing truth claims. On the contrary.

I just said that we have a large and growing body of knowledge which is *certain, objective, and universal*. I emphasize these three traits because they are precisely what is challenged by a certain contemporary form of extreme skepticism sometimes called "post-modernism," with such subsidiary branches as "deconstruction," "post-structuralism," and even some versions of pragmatism. According to this skeptical challenge, it is at best a mistake, and at worst a kind of totalitarian impulse, that leads us to say that we can have certainty, objectivity, and universality. According to this view, we never attain certain, objective, and universal knowledge at all. This is supposedly shown by certain investigations of science, such as those conducted by Thomas Kuhn and Paul Feyerabend that emphasize the irrational elements in the development of scientific theories. On this view, scientists do not attain truth; rather, they rush irrationally from one paradigm to another. Furthermore, the story goes, it is impossible to have objectivity, because all claims to knowledge

are always perspectival; they are always made from a certain subjective point of view. And finally, it is impossible to have universality, because all science is produced in local, historical circumstances and is subject to all of the constraints imposed by those circumstances. I believe that these challenges are without merit, and I want to briefly say why. The main point I want to make is that what is true in the skeptical challenges is in no way inconsistent with certainty, objectivity, and universality.

One of the problems that we have, in coming to terms with the huge growth of knowledge, is to see how all of these features can exist simultaneously. How can knowledge be at the one and the same time *certain* and yet *tentative* and *corrigible*, how can it be totally *objective* and yet always from one *subjective* perspective or another, how can it be absolutely *universal*, and yet the product of *local* circumstances and conditions? Let us go through these in order. The certainty in question derives from the fact that the evidence for the claims in question is so overwhelming, and the claims themselves are so well embedded in a systematic set of inter-related claims, that are all equally well-supported by overwhelming evidence, that it is simply irrational to doubt these truths. At present it is irrational to doubt that the heart pumps blood, that the earth is a satellite of the sun, or that water is made of hydrogen and oxygen. Furthermore, all of these items of knowledge are embedded in very powerful theories, the theories of human and animal physiology, the heliocentric theory of our planetary system and the atomic theory of matter. But at the same time it is always possible that there could be a scientific revolution that will overthrow these whole ways of thinking about things, that we might have a revolution comparable to the way in which the Einsteinian Revolution assimilated Newtonian mechanics as a special case. Nothing in any state of knowledge, however certain, can preclude the possibility of future scientific revolutions. This tentativeness and corrigibility is not a challenge to certainty. On the contrary, at one and the same time, we have to recognize certainty, and yet acknowledge the possibility of future major changes in our theories.

I want to emphasize this point: There is a very large body of knowledge that is known with certainty. You will find it in the university bookstore, in, for example, textbooks on engineering or biology. The sense in which we know with certainty that the heart pumps blood, for example, or that the earth is a satellite of the sun, is that given the overwhelming weight of reasons that support these claims, it is irrational to doubt them. *But certainty does not imply incorrigibility.* It does not imply that we could not conceive of circumstances in which we would be led to abandon these claims. It is a traditional mistake, one I am now trying to overcome, to suppose that certainty implies incorrigibility by any future discovery. We are all brought up to believe that certainty is impossible because claims to knowledge are always tentative and subject to further correction. But this is a mistake. Certainty is not inconsistent with tentativeness and corrigibility. There is no question that we know a great many things with certainty, and yet those things are revisable by future discoveries.

That leads to the second combination of features: how can knowledge at one and the same time be completely objective and yet perspectival, always stated and assessed from one perspective or another? To say that a knowledge claim is epistemically objective is to say that its truth or falsity can be established independently of the feelings, attitudes, prejudices, preferences, and commitments, of investigators. Thus, when I say that “Water is composed of H_2O molecules,” that claim is completely objective. If I say, “Water tastes better than wine”—well, that claim is subjective. It is a matter of opinion. It is characteristic of knowledge claims, of the sort that I have been discussing, that when I say that such knowledge grows cumulatively, the knowledge in question is, in this sense, epistemically objective. But such objectivity does not preclude perspectivity. Knowledge claims are perspectival in the obvious and trivial sense that all claims are perspectival. All representations are from a perspective, from a point of view. So when I say, “Water consists of H_2O molecules,” that is a description at the level of atomic structure. At some other level of description, at the level of sub-atomic physics, for example, we might wish to say that water consists of quarks, muons, and other sundry sub-atomic particles. The point for our present discussion is that the fact that all knowledge claims are perspectival does not preclude epistemic objectivity.

I want to state this point emphatically: All representation of reality, human or otherwise, and *a fortiori* all knowledge of reality, is from a point of view, from a certain perspective. But the perspectival character of representation and knowledge does not imply that the knowledge claims in question are dependent on the preferences, attitudes, prejudices, predilections, of observers. The existence of objectivity is in no way threatened by the perspectival character of knowledge and representation.

Finally, knowledge claims of the sort that I am talking about, where we make claims about how the world works, are universal. What is true in Vladivostok is also true in Pretoria, Paris, and Berkeley. But the fact that we are able to formulate, test, verify, and conclusively establish such claims as certain, universal, and objective, requires a very specific socio-cultural apparatus. It requires an apparatus of trained investigators, and the social cultural conditions necessary for the existence of such training and such investigation. These have developed most strongly in Western Europe and its cultural offshoots in other parts of the world, especially North America, during the past four centuries. There is a trivial and harmless sense in which all knowledge is socially constructed. In this trivial and harmless sense knowledge is expressed in statements, in claims; and these claims have to be formulated, formalized, tested, verified, checked and rechecked. That we are able to do this requires a very specific sort of socio-cultural structure, and in that sense, our knowledge claims are socially constructed. But social construction in this sense is not in any way in conflict with the fact that knowledge so arrived at is universal, objective, and certain.

I want to emphasize this third point just as I did the first two: Knowledge claims are made, tested, and verified by historically situated individuals working against the background of specific cultural practices. In this sense all knowledge claims are socially constructed. But the truth of such claims is not socially constructed. Truth is a matter of objective facts in the world that correspond to our knowledge claims.

So far I have considered three objections to the common sense view that we have a large body of knowledge that is certain, objective, and universal. First, knowledge is always tentative and corrigible, second, it is always stated from a point of view, and third, it has to be arrived at by cooperative human efforts working in particular historically situated social contexts. The chief point I am making is that there is nothing inconsistent between these theses and the claim that knowledge so arrived at is often *certain, objective, and universal*.

If by “modernism” is meant the period of systematic rationality and intelligence that began in the Renaissance and reached a high point of self-conscious articulation in the European Enlightenment, then we are not in a post-modern era. On the contrary, modernism has just begun. We are, however, I believe, in a post-skeptical or post-epistemic era. You will not understand what is happening in our intellectual life if you do not see the exponential growth of knowledge as the central intellectual fact. There is something absurd about the post-modern thinker who buys an airplane ticket on the internet, gets on an airplane, works on his laptop computer in the course of the airplane flight, gets off of the airplane at his destination, takes a taxicab to a lecture hall, and then gives a lecture claiming that somehow or other there is no certain knowledge, that objectivity is in question, and that all claims to truth and knowledge are really only disguised power grabs.

II. THE POST-SKEPTICAL ERA

Assuming that I am right about these features of knowledge and about the fact that knowledge continues to grow, what are the implications for philosophy? What does philosophy look like in a post-epistemic, post-skeptical era? It seems to me that it is now possible to do systematic theoretical philosophy in a way that was generally regarded as out of the question a half a century ago. Paradoxically, one of Wittgenstein’s great contributions to philosophy is one that he himself would reject. Namely, by taking skepticism seriously and attempting to cope with it, Wittgenstein has helped to pave the way for a type of theoretical and systematic philosophizing that he himself, in his later work, abominated and thought impossible. Precisely because we are no longer worried about the traditional skeptical paradoxes and about their implications for the very existence of language, meaning, truth, knowledge, objectivity, certainty, and universality, we can now get on with the task of general theorizing.

The situation is somewhat analogous to the situation in Greece after the transition from the philosophy of Socrates and Plato to the philosophy

of Aristotle. Socrates and Plato took skepticism seriously; Aristotle was a systematic theoretician.

With the possibility of developing general philosophical theories, and the decline of the obsession with skeptical worries, philosophy has eliminated much of its isolation from other disciplines. So, for example, the best philosophers of science are as familiar with the latest research as are specialists in those sciences.

There are a number of topics I could discuss concerning the future of philosophy, but for the sake of brevity, I will confine myself to six subjects.

1. 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 cooperation between scientists and philosophers. There are different versions of the mind-body problem but the one most intensely discussed today is: 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 that all the features of

the world from quarks and electrons to nation states and balance of payments problems are, in their different ways, parts 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 "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 old conception of 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 which 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 epistemic sense of the objective/subjective distinction, 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. 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 which 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 the animal's behavior 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 methods we use to verify 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: In order to find out how the world works, you have to use any weapon that 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 seeing the color red, 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 program 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 color or the shape, or the sound, of an object, we perceive all of these at once 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 color, 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. The crucial problem is not, for example, “How does the brain produce the conscious experience of red?” but rather, “How does the brain produce the unified, subjective conscious field?” We should think of perception not as *creating* consciousness, but as *modifying* a preexisting conscious field. 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 sorts of 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 NCC’s 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 more global features of the brain, such as massive synchronized patterns of neuron firing in the thalamocortical system.¹

2. 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 center 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 fifty 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 perception, intentions, beliefs, and desires. 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 behaviorism 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.

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. The notion "same implemented program" defines an equivalence class that is specified entirely in terms of formal or syntactical processes and is independent of the specific physics of this or that hardware implementation. This principle underlies the famous "multiple realizability" feature of computer programs. The same program can be realized in an indefinite range of hardware. The mind cannot consist in a program or programs, because the syntactical operations of the program are not by themselves sufficient to constitute or to guarantee the presence of semantic contents of actual mental processes. 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 processes of the implemented program do not by themselves have any semantic contents. I demonstrated this years ago with the so-called Chinese Room Argument.²

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 favor of speculative neurobiology.

What is actually happening in cognitive science is a paradigm shift away from the computational model of the mind and toward 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 common sense 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 investigations have made apparent to us. Even in popular speech we now distinguish between short-term and long-term memory, and no doubt as our investigation proceeds, we will have further distinctions. 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 common sense 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.

I give the example of memory as one instance where an ongoing research project raises philosophical questions and has philosophical implications. I could have given other examples about linguistics, rationality, perception and evolution. I see the development of a more sophisticated cognitive science as a continuing source of collaboration between what was traditionally thought of as the two separate realms of "philosophy" and "science."

3. THE PHILOSOPHY OF LANGUAGE

I said that the philosophy of language was the center of philosophy for most of the twentieth century. Indeed, as I remarked, during the first three quarters of the twentieth century, the philosophy of language was taken to be “first philosophy.” But by the end of the century that had changed. Less is happening in the philosophy of language now than in the philosophy of mind, and I believe that the currently most influential research programs have reached a kind of dead end. Why? There are many reasons of which I will mention only two.

First, one of the main research programs 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 behaviorism, led some prominent philosophers to try to give an analysis of meaning according to which the hearer is engaged in the epistemic task of trying to figure out what the speaker means either by looking at his behavior in response to a 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 would thereby analyze meaning. This reoccupation with the epistemic aspect of language use leads to the same confusion between epistemology and ontology that has bedeviled the Western philosophical tradition for over three centuries.

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, 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. Frege was anxious to insist that 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 “internalist” 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.

Frege’s real mistake, and it is one that I repeated, is to suppose that the way in which language relates to reality—the “mode of presentation”—also fixes propositional content. Frege assumed both that sense determines reference and that propositional content consists in sense. But if by the notion of “proposition” we are interested in that which we assess as true or false, then it is not the case that sense is identical with propositional content, because often we are interested in the actual objects that are referred to rather than the mode in which they are referred to. This is especially true of indexicals. We need to separate the question, “How do words relate to the world?” from the question, “How is propositional content determined?” However, the correct observation by the externalists, that the content of a proposition cannot always be specified by what is internal to the mind, does not show that the contents of the mind are insufficient to fix reference. I have discussed these issues in some detail elsewhere and will not repeat the discussion here.³

4. 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 twentieth 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 twenty-first century we will feel a pressing need for, and should certainly try to develop, what I will call a philosophy of society. At present 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. A student today taking a course called “Social Philosophy” is likely either to be studying Rawls on justice (political philosophy) or Hempel on covering law explanations in the social sciences (philosophy of social science). I am proposing that we should have a free standing philosophy of society, 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, games, etc. when such entities, in some sense, exist only in virtue of a collective 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 my book, *The Construction of Social Reality*.⁴

Specifically, I believe that in our study of political and social reality, we need a set of concepts which will enable us to describe political and social reality, so to speak from the “middle distance.” Our problem 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 the developing criteria for assessing the justice or injustice of institutions. Much of the progress in this area is owed to John Rawls, who revolutionized the study of political philosophy with his classic work, *A Theory of Justice*.⁵ 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 twenty 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 which 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 twentieth 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. There are different definitions of “socialism,” but they all have one thing in common: a system can only be socialist if it has public ownership and control of the basic means of production. The failure of socialism so defined is the single most important social development of the twentieth century. It is an amazing fact that that development remains unanalyzed 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 also to the failure of democratic socialism as it existed in the countries of Western Europe. The socialist parties of those countries continue to use some of 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, has been quietly abandoned. What is the correct philosophical analysis of this entire phenomenon?

A similar sort of question would involve the appraisal of national institutions. So, for example, for most political scientists, it would be very difficult to attempt to analyze 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 give us the tools for analyzing social institutions as they exist in real societies in a way that would enable us to make comparative judgments between different countries and larger societies, without rising to such a level of abstraction that we cannot make specific value judgments about specific institutional structures. The work of the economist-philosopher Amartya Sen is a step in this direction.

5. ETHICS AND PRACTICAL REASON

For much of the twentieth century the subject of ethics was dominated by a version of the same skepticism 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 epistemic objectivity. The principal issue in ethics was about whether or not there could be 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 that ethical statements could not literally be either true or false, but functioned only to express feelings or to influence behavior, 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 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 skepticism 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.

6. THE PHILOSOPHY OF SCIENCE

In the twentieth 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 skeptical paradoxes, such as the traditional problem of induction. Throughout most of the twentieth 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." Scientists formed a hypothesis, deduced logical consequences from it, and then tested those consequences in experiments. This conception was articulated, I think more or less independently, by Karl Popper and Carl Gustav Hempel. Those practicing scientists who took an interest in the philosophy of science at all tended 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 of 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 at 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.⁶

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 *The Structure of Scientific Revolutions* in 1962.⁷ It is puzzling that Kuhn's book 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 scientific paradigm is overthrown by a new 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 only 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 practicing scientists, but it had an enormous effect in several humanities disciplines, especially those connected with the study of literature. Kuhn seemed to have refuted the claim that science gives us truths about the world; rather science gives us no more truth about the real world than do works of literary fiction or works of literary criticism. Science is essentially a set of irrational processes whereby groups of scientists form theories which are more or less arbitrary social constructs, and then abandon these in favor of other theories, which are equally 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 skeptical view of Paul Feyerabend, who argued that as far as giving us truths about the world, 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 twenty-first century. I think the essential problem is this: Twentieth 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, not only as regards causation and determinacy but also as regards the ontology of the physical world.

Most philosophers, like most educated people today, have a conception of causation which 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 simple mechanical relations like gear wheels moving gear wheels, and other such Newtonian phenomena. We know at some abstract level that this picture is not right, but we still have not replaced our common sense conception with a more sophisticated scientific conception. I think that working through these issues is one of the most exciting tasks of the twenty-first century philosophy of science. We need to give an account of physical theory, especially quantum theory, that will enable us to assimilate physical results 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.

III. CONCLUSION

The main message I have tried to convey is that it is now possible to do a new kind of philosophy. With the abandonment of the epistemic bias in the subject, such a philosophy can go far beyond anything imagined by the philosophy of a half century ago. It begins, not with skepticism, but with what we all know about the real world. It begins with such facts as those stated by the atomic theory of matter and the evolutionary theory of biology, as well as such “common sense” facts as that we are all conscious, that we all really do have intentional mental states, that we form social groups and create institutional facts. Such a philosophy is theoretical, comprehensive, systematic, and universal in subject matter.

ENDNOTES

This article is a revision of “The Future of Philosophy” which was written for a scientific rather than a philosophical audience and published in a special millennium issue of the *Philosophical Transactions of the Royal Society series B, London 354* (1999), 2069–2080. I am indebted to Dagmar Searle for discussion of all these issues.

1. I have discussed these issues in much greater detail in my, “Consciousness,” *The Annual Review of Neuroscience* 23 (2000): 557–578.
2. John Searle, “Minds, Brains and Programs,” *Behavioral and Brain Sciences* 3 (1980): 417.
3. John Searle, *Intentionality: an Essay on the Philosophy of Mind* (Cambridge: Cambridge University Press, 1983), chaps. 8–9.
4. John Searle, *The Construction of Social Reality* (New York: Free Press, 1995).
5. John Rawls, *A Theory of Justice*, (Cambridge, Mass.: Harvard University Press, 1971).
6. For an interesting criticism of Popper’s views see David Stove, *Against the Idols of the Age* (Somerset, N.J.: Transaction, 1999).
7. Thomas Kuhn, *The Structure of Scientific Revolutions* (Chicago: University of Chicago Press, 1962).