CONSCIOUSNESS AND LANGUAGE

JOHN R. SEARLE

University of California, Berkeley

CONTENTS

INTRODUCTION

The essays collected in this volume were written over a period of two decades. They deal with a wide range of subjects and were intended for a variety of audiences. Despite the variety, there are certain unifying principles that underlie this collection; indeed, I have tried to make a selection that will exhibit a natural progression, as the topics move from consciousness to intentionality to society to language, and finally conclude with several debates about the issues that have preceded. In this introduction I want to try to state some of these unifying principles and offer a brief description (and note I say "description" rather than "summary" or "abstract") of the essays.

There is a single overarching problem that has preoccupied me since I first began work in philosophy almost a half-century ago: How can we have a unified and theoretically satisfactory account of ourselves and of our relations to other people and to the natural world? How can we reconcile our common-sense conception of ourselves as conscious, free, mindful, speech-act performing, rational agents in a world that we believe consists entirely of brute, unconscious, mindless, meaningless, mute physical particles in fields of force? How, in short, can we make our conception of ourselves fully consistent and coherent with the account of the world that we have acquired from the natural sciences, especially physics, chemistry, and biology? The questions that have most preoccupied me – What is a speech act? What is consciousness? What is intentionality? What is society? What is rationality? – have all in one way or another been addressed to this larger problematic. I think this problem – or set of problems – is the most important problem in philosophy, and indeed there is a sense in which, in our particular epoch, it is the only major problem in philosophy.

If one accepts this characterization of my philosophical project then certain other features of my approach to these problems will become apparent, features that are by no means universally shared in the profession. First, philosophical problems of the kind we are dealing with should have clearly stated and definite solutions. There has to be a definite answer to such questions as "What exactly is consciousness and how does it fit in with the rest of the world?"; otherwise we are not making any progress with our task. Though philosophical problems have definite solutions, the solutions can seldom be given as direct answers to philosophical questions. The way I try to proceed is first to analyze the question. Indeed, this is the great lesson of twentieth-century linguistic philosophy: Do not take the questions for granted. Analyze the question before attempting to answer it. I like to proceed by analyzing the question to see whether it rests on a false presupposition, or whether it assimilates the problem at issue to an inappropriate set of paradigms, or whether the terms used in the question are systematically ambiguous. I find that in one way or another, philosophical problems characteristically require dismantling and reconstructing before they can be solved. Once clarity is achieved about exactly what questions are being asked, the answers, or at least the philosophical part of the answers, are often quite clear and simple.

Let me illustrate these points with some examples. Consider the famous "mind-body problem." What exactly is the relationship between consciousness and brain processes? To tackle this question we have to go behind the problem as posed and ask: What presuppositions does this formulation of the problem rest on? This problem resists solution as long as we continue to acept the traditional seventeenth-century vocabulary which presupposes that mental phenomena, naively construed, are in a completely different and separate ontological realm than physical phenomena, naively construed. Once we abandon this vocabulary, and the presuppositions on which it rests, the philosophical problem has a rather simple solution. Once we see that consciousness, with all its inner, qualitative, subjective, touchy-feely qualities, is just an ordinary property of the brain in the way that digestion is a property of the stomach, then the philosophical part of the problem is fairly easy to resolve. But there remains a terribly difficult scientific problem about how it actually works in the brain. I have something to say about both these issues, the philosophical and the neurobiological, in the course of these essays.

This is one pattern of addressing philosophical problems. A clarification of the problem will leave us with residual issues, but these are amenable to solution by scientific, mathematical, or other well-established means. What

I just said about the mind-body problem also applies, for instance, to the free will problem.

If one thinks of the primary task of philosophical analysis as giving a true, complete, and unifying account of the relations of humans and nature, then another distinction between this approach and that of much mainstream philosophy immediately appears. I find that I, personally, cannot take traditional skeptical worries very seriously. I think we made a mistake in taking so seriously, through the twentieth century, the line of skeptical argument and response that was begun by Descartes. I think in the seventeenth century it was reasonable to consider the existence of knowledge as problematic, and to feel that it required a secure foundation. Now it seems to me absurd to try to treat the existence of knowledge as problematic. If one thing is clear about the present intellectual situation it is that knowledge grows daily. The existence of knowledge is not in doubt. But you would be surprised at how much the persistence of skeptical worries, and the consequent epistemic stance, continue to have deleterious effects in philosophy. So, for example, Quine's famous indeterminacy argument, and indeed the whole project to examine meaning using the method of radical translation, is a matter of adopting an epistemic stance where I think it is inappropriate.

The first five essays – "The Problem of Consciousness," "How to Study Consciousness Scientifically," "Consciousness," "Animal Minds," and "Intentionality and Its Place in Nature" – are all concerned in one way or another with situating consciousness in particular, and intentional phenomena in general, within a scientific conception of how the world works. I think of consciousness and intentionality as biological phenomena on all fours with digestion or photosynthesis. If you think of mental phenomena in this way, you will not be tempted to think, for example, that a computational simulation of the mental phenomena is somehow itself a mental phenomenon. Computational simulations of the mind stand to the real mind the way computational simulations of the stomach stand to the real stomach. You can do a simulation of digestion, but your simulation doesn't thereby digest. You can do a simulation of thinking, but your simulation doesn't thereby think.

When I say that consciousness is a biological phenomenon, many people understand that claim to imply that you could not create consciousness artificially by building a machine. But that is exactly the opposite of the implication that I intend. The brain is a machine, and there is no reason in principle why we could not build an artificial brain that thinks and has other mental processes in the way that our brain does. Indeed, there is no reason in principle why we couldn't build an artificial brain out of completely nonbiological materials. If you can build an artificial heart that pumps blood,

why not build an artificial brain that is conscious? However, note that in order to actually produce consciousness, you would have to duplicate the relevant threshold causal powers that the brain has to do it. That is like saying that if you want to build an airplane that can fly, you don't necessarily have to have feathers and flapping wings, but you do have to duplicate the threshold causal powers that birds have to overcome the force of gravity in the earth's atmosphere. These are obvious points, but you would be surprised at how much difficulty I have had at getting them across to the philosophical and cognitive science community over the past two decades.

The first three essays are concerned explicitly with the problem of consciousness. They are arranged in chronological order and exhibit a certain level of development in my thinking about consciousness, together with increasing sophistication in neurobiology about the problem of consciousness. The first essay tries to state exactly what the problem of consciousness is by listing the main features that a theory of consciousness would need to explain. The second essay lists a number of mistakes that we need to avoid if we are to get a scientific account of consciousness, and the third essay reviews some features of the current project of explaining consciousness in neurobiology. In this last and most recent essay, I try to identify the prospects as well as the limitations of current research on consciousness. I distinguish between what I call the "building block model" and the "unified field" conception of consciousness, and make an argument that the unified field conception is more likely to succeed as a neurobiological research project.

This approach to the mind has important implications for the explanatory apparatus to be used in explaining cognitive phenomena. These issues are discussed in the next three essays. The first of these, "Animal Minds," defends the common sense view that animals have minds just as we do – with consciousness, beliefs, and desires, as well as pains and pleasures – but that their mental contents are restricted because animals do not have a language. "Collective Intentions and Actions" extends the account of Intentionality from the individual cases "I think," "I believe," "I intend" to the collective cases "We think," "We believe," "We intend." There is a nontrivial problem about getting the formulation of the conditions of satisfaction right for collective intentionality and much of that essay is devoted to that question. How exactly do we represent the content of my intention to do something, where I have the intention, as an individual, only as a part of our having the intention to do something as a collective? I am playing the violin part of Beethoven's Ninth Symphony and thus making my contribution to our collective activity of playing Beethoven's Ninth Symphony. You are, let us suppose, singing the soprano part and thus making your contribution to

our collective activity. How exactly do we represent the conditions of satisfaction of the individual and the collective intentions, and their relation to each other? I don't know if anybody else cares about that question, but it gave me a lot of headaches, and this essay is an attempt to characterize those relations in a way that recognizes both the irreducibility of collective intentionality and the fact that individual intentionality will be necessary to move my body even in cases where I move my body as part of a collective activity. I am definitely not trying to reduce the collective intentionality to individual intentionality, even though I recognize that all human intentionality is in the brains of human individuals.

The next two essays, 7 and 8, discuss the implications of my overall approach to the mind for psychology and other social sciences. In "The Explanation of Cognition," I explore in detail what I think is the correct explanatory apparatus for a sophisticated cognitive science to use, and in "Intentionalistic Explanations in the Social Sciences," I discuss the implications of my overall account of the mind for the traditional disputes between empiricist and interpretivist approaches to the problem of explanation in the social sciences.

The next two essays, 9 and 10, are concerned with extending my earlier work on speech acts in light of my researches in the philosophy of mind. A nagging dispute in the theory of speech acts is between those authors like Grice who take the individual intentionality of the speaker as the essential analytic device, and those like Austin and Wittgenstein who emphasize the role of convention and social practice. These appear to be inconsistent approaches, but I argue in "Individual Intentionality and Social Phenomena in the Theory of Speech Acts" that it is possible to reconcile these two approaches. Properly construed they are not competing answers to the same question, but noncompeting answers to two different and related questions. Essay number 10, "How Performatives Work," attempts to explain a phenomenon that originally gave rise to the whole subject of speech acts but which, oddly enough, seems to me not to have been satisfactorily explained. Namely, how is it possible that we can perform the act of promising or ordering, or christening, or blessing, and so forth, just by saying that we are doing it? The paradox here is that the whole subject of speech acts grew out of Austin's discovery of performative utterances and his ultimate rejection of the distinction between constatives and performatives. If you reject the distinction between constatives and performatives, as I do, and as Austin did – and without this rejection there is no such thing as a theory of speech acts – then you are still left with a problem. How do you explain the original existence of performative utterances? Essay 10 attempts to answer that question.

"Conversation" is about the possibility of extending my account of speech acts to larger stretches of discourse involving two or more people, to conversations. I reach a somewhat pessimistic conclusion. We will not get an account of conversation comparable in explanatory power to the theory of speech acts.

The last three essays are more argumentative than the earlier material. "Analytic Philosophy and Mental Phenomena" attempts to explain and overcome the puzzling tendency that many analytic philosophers have had to reject mental phenomena, naively construed. There has always been a tradition in analytic philosophy that attempts to get rid of consciousness and intentionality in favor of behaviorism, functionalism, computationalism, or some other version of "materialism." I diagnose and answer what I think are some of the most flagrant versions of this error.

Finally, the last two essays. In my work in both the theory of speech acts and the theory of intentionality I presuppose a form of mental realism. I assume we really do have beliefs and desires and other intentional states, and that we really do mean things by the words that we utter. Our mental states have a more or less definite content, and our utterances have a more or less definite meaning. When I began work on intentionality there were two forms of skepticism that seemed to challenge these assumptions: Quine's skepticism about the indeterminacy of translation and meaning, and Kripke's version of Wittgenstein's private language argument. If Quine's argument were correct, it would apply not only to linguistic meaning but to intentionality generally. Similarly with Kripke's skeptical interpretation of Wittgenstein's famous argument: If there is no definite ascertainable fact of the matter about whether or not I am using a word correctly, then it seems there is nothing for a theory of meaning and intentionality to be about. I could not proceed with my work on meaning and intentionality before I answered these two skeptical arguments, at least to my own satisfaction. I try to answer Quine's indeterminacy argument and to show that it is best construed as a reductio ad absurdum of the behaviorist premises from which it proceeds. In answering the Kripkean form of skepticism I distinguish two lines of argument in his book, and I claim that only the second of these is used by Wittgenstein. In any case I try to deal with both. This last chapter, by the way, has not been previously published because, frankly, I thought too much had been published about this issue already. I wrote it at the time of Kripke's publication but did not attempt to publish it. However, it does seem to fit neatly into the context of the present book, so I have included it, even though it is the only previously unpublished article in the collection.