# Introduction

## 1.1 The Mystery

When we speak of ourselves as agents with minds, we use the language of *common-sense psychology* (CSP). Within the common-sense psychological project, we attempt to understand people (including ourselves) as agents who act for reasons, who have moral responsibility for our actions, who act freely at least some of the time, who have purposes and goals and projects, who are loci of value, and for whom things have value. We attribute mental states in an effort to understand the point or purpose of behavior; we justify our own actions by referring to the way we believe or desire the world to be, we attribute mental states in the course of allocating moral praise or blame, and so on. CSP will deliver claims like the following: Jackie went to the kitchen in order to get wine. People who prefer beer to wine will typically choose beer if offered both. Tom Seaver threw a curveball because he thought the batter was expecting a fastball.

When we think of human beings as physical objects, we use the language of physical science, and we give causal accounts. The specific terminology we employ depends on the nature of the event we are attempting to explain or predict. If we want to know how fast a person falls if dropped from a plane, we will be concerned with the body's mass and aerodynamic properties. If we want to know why a certain medicine reduces inflammation, our explanation will be couched in terms of chemical properties of the medicine and the human organism. If we want to know how a given subject will react to a particular stimulus, we would have to look deep into the neurological properties of the brain. Although the latter sort of predictions are, in general, quite beyond our current abilities, there seems to be no barrier in principle to neurologically based predictions of human behavior.

Both CSP and physical science have the aim of understanding, explaining, and predicting human behavior. However, talk of mind and agency

does not seem to mesh well with the language of physical science. Notions like action and purpose, the cornerstones of our descriptions of agents, appear to have no role in purely physical descriptions of the world. Normative notions like moral responsibility and criticizable irrationality, also central to the common-sense perspective, likewise do not belong to the conceptual arsenal of physics. We would never say that an asteroid was morally responsible for its motion, even if it crashed into Earth, nor would we say that an elementary particle was being irrational. Moreover, beliefs, desires, and other mental states have the odd feature of being *about* objects and states of affairs. This feature of beliefs and desires—given the label of "intentionality"—is not a property attributed by physics to physical objects. On the other hand, since human beings are physical objects, they must ultimately be appropriate subject matter for physical science, and they must behave in accord with physical laws. Physical science must be able, in principle, to provide detailed causal accounts of all bodily motions. But physical science will presumably not use the language of CSP in describing human beings as physical objects, and hence it will apparently not make the sorts of claims about persons that we are accustomed to making in CSP.

Thus, there is an apparent philosophical mystery. Putting aside substance dualist views according to which we have nonphysical souls, it seems that a human being is, in some sense, nothing more than a very complicated collection of physical particles. (I will discuss the reasons for rejecting substance dualism in the next chapter.) Thus, it seems that human behavior should be within the explanatory domain of physical science. Moreover, there is reason for thinking that, within its explanatory domain, physical science will ultimately provide complete explanations of human behavior. This makes it appear that, if common-sense talk of mind and agency is to have any validity, it must somehow be subsumed within physical science. However, the language of CSP is also quite different from the language of physical science, so it appears unlikely that the claims of CSP can simply be subsumed within the claims of physical science. At the broadest level the problem is this: It seems that mind does not fit with the rest of the physical universe, and yet it seems that mind must fit with the rest of the physical universe. The problem is hardly new; it is the mind-body problem, and it is the problem of discovering the place of rational agency. It is the problem of ascertaining the place of human beings in the natural world and of seeing whether the common-sense facts about mind and agency can be squared with a naturalistic, scientific perspective on the world.

# 1.2 Reformulation of the Mystery

This philosophical mystery has been couched in rather vague terms so far. I have talked metaphorically about whether CSP will *mesh well* with physical science, and have claimed that *in some sense* human beings are nothing over and above collections of physical particles. We feel the mystery strongly, but much of the task of philosophy of mind consists in formulating the problem in a tractable, clear, and precise way. We can, I think, make progress toward this goal by asking the following question: What is the logical relationship between the claims put forward by CSP and the claims that will ultimately be put forward by physical science?

If we consider two propositions p and q, we can classify their possible logical relationships as follows:

- (1) p and q contradict each other.
- (2) One of p and q entails the other.
- (3) p and q are logically independent.

We would have the first of these options if the two propositions were "Only adults are at the beach" and "Hayden is at the beach, and Hayden is not an adult." The second relationship holds between "Water is wet, and water is  $H_2O$ " and " $H_2O$  is wet." The third relationship, that of logical independence holds between "Tom Seaver was the best pitcher the Mets ever had" and "Grass is green."

Applying this to the relationship between CSP and physical science, there are three options.

option 1: CSP and physical science contradict one another.

option 2: Physical science entails CSP, CSP entails physical science, or CSP and physical science are equivalent.

option 3: CSP and physical science are logically independent.

However, we must tread somewhat carefully in characterizing these options, since CSP and physical science are *sets* of propositions, and we need to be explicit about what we mean when we say that one set of propositions bears a certain logical relationship to another. One way to go would be to say that CSP and physical science contradict each other if any proposition or conjunction of propositions within CSP contradicts any proposition or set of propositions within physical science. Similarly, we could say that physical science entails CSP if the conjunction of the propositions of physical science logically implies each of the propositions of CSP. However,

it will be preferable to take a somewhat more relaxed approach. If, for example, only one or two of the claims of CSP conflict with those of science, we would be hesitant to say that, in general, CSP and physical science contradict one another. We would reach this conclusion only if science contradicts some substantial subset of CSP. Similarly, we would say that physical science entails CSP if the propositions of physical science entail the bulk of CSP, without requiring that every single proposition of CSP be entailed by science in this way. Finally, we would say that CSP and physical science are independent if the bulk of CSP is logically independent of science. Of course, since "substantial subset" and "bulk of" are vague, there could be a corresponding degree of vagueness concerning which option holds.

An example involving something other than CSP might help at this point. Instead of common-sense psychology, consider astrology. In the words of one of its proponents, astrology is "the science which treats of the influence upon human character of cosmic forces emanating from celestial bodies" (N. de Vore, past president of the Astrological Research Society, quoted in Culver and Ianna 1988, p. 2). Within this purported science, numerous claims can be made, such as the following:

There is "an observable parallelism between the timing of events in the universe and in the individual consciousness." (Rudhyar, quoted in Culver and Ianna 1988, p. 1)

If Mars is in the first house at the time of your birth, then you are likely to be an aggressive individual with a muscular body.

If the Sun is in the fifth house at the time of your birth, then you are likely to have great creative talents.

Most of us see no great philosophical mystery regarding astrology. However, for the 25 percent of Americans who believe in astrology, there is a mystery, for astrology and science do not seem to mix well. Physical science does not claim that astronomical bodies have any astrological influence on human events and human character. The planet Jupiter, for example, does exert a force on me; however, according to physical science, the force is that due to gravity, and the gravitational attraction between Jupiter and a human body on Earth is minuscule. One might then wonder why physical science—our best attempt to make sense of the nuts and bolts of the universe—does not include astrological claims or principles. One might then also wonder about the nature of astrological forces.

As with CSP and physical science, I suggest that a "philosopher of astrology" could sharpen these worries by asking what the logical relationship is between astrology and physical science. And again we have the same three

options. Option 1 is that astrology and science contradict one another and that at most one is right. Option 2 is that physical science ultimately entails the truths of astrology, or that, in an even more unlikely scenario, all facts of physical science are entailed by the astrological facts about the world. Option 3 is that astrology and physical science are logically independent, meaning that neither contradicts the another neither is entailed by the other. If we can determine which of these relationships holds, and why, we are well on our way to resolving any sense of philosophical mystery about astrology. If, as seems all but certain, astrology and our best physical science contradict one another, then at most one can be right. If we reject astrology on these grounds, we have resolved the mystery by demoting astrology to the realm of falsehood and illusion. On the other hand, if option 2 turns out to be correct, then the principles and methods of astrology follow from the facts of physical science. This would resolve the mystery by vindicating astrology; astrology would be true if science is true. Astrology would in this sense be subsumed within physical science. Finally, if we could produce reasons for maintaining that astrology and science are logically independent of each other, then we would thereby have resolved any tension between astrology and science, and we would have shown that astrology and science can both be true.

Similarly, in the case of CSP and physical science, if we can determine which of the three logical relationships holds, then we will resolve the philosophical mystery. Option 1 gives us the following resolution to the mystery: we are right to feel a tension between claims about the mind and claims about the physical universe, for typical CSP claims about the mind actually contradict what we know or will discover through physical science. Thus, we would be forced to reject either CSP or physical science. Since bets placed against science have a sorry history, we would conclude that the claims of CSP are generally false.

Option 2 is a disjunction of three possibilities; however, only one of them has any plausibility. No one would maintain that all the facts of physical science, including the value of the gravitational constant, are implied by the facts of CSP. Thus, option 2, as we will consider it, is this claim: The facts of CSP are ultimately implied by the facts of physical science. As was noted above, this might be a matter of degree, for it might be that a few of the alleged facts of CSP ultimately contradict physical science and have to be rejected. But so long as the facts of physical science imply the bulk of CSP, we have a resolution to our philosophical mystery. On this view, CSP, despite its apparent differences with the language of physical science, is implied by physical science. If CSP simply follows from the propositions

ultimately put forward by science, then CSP is vindicated. CSP would thus be neatly subsumed within physical science.

Rather than affirm either of the first two options, we might take a more cautious approach. We could say that CSP will either be implied by science or it will contradict science, but that before we can determine which we will have to wait for more scientific results. We might, for example, have to learn more about the brain before we can say with confidence either that CSP fits nicely within physical science or that it contradicts physical science. That is, we might say that either option 1 or option 2 is correct, but only further empirical research in science will be able to determine which. I will refer to this approach as *strong naturalism*.

Finally, there is option 3, according to which CSP and physical science are logically independent and psychological facts about the mind and physical facts about bodies can coexist peacefully. This would resolve the initial mystery by showing that there is no conflict between the two after all—i.e., that the conflict had been merely apparent rather than real. One way to support the third option would be to adopt substance dualism, the claim that minds are immaterial substances; this would make it natural to assume that propositions about mind and propositions about material bodies simply have different subject matters. CSP and physical science would then each be about different things, and would independent for the same reason that "Grass is green" and "Snow is white" are independent.

#### 1.3 Addressing the Reformulated Question

Determining the logical relationship between two individual statements is typically a straightforward affair, although there is no mechanical procedure that is guaranteed to determine, in a finite number of steps, the logical relationship between two arbitrary propositions. It is similarly straightforward, though more complicated, to determine the logical relationship between two well-defined, finite sets of statements. But our problem is more complicated than that. CSP is an ill-defined mass of propositions including individual attributions of mental states to particular agents, explanations of the behavior of individual agents, and generalizations concerning behavior and mental states. Even if we could produce all the generalizations that are held true in CSP, attributions and explanations can arise in an indefinitely large number of circumstances, and thus there is no effective way of specifying all of the propositions of CSP.

The propositions of physical science will likewise include an indefinitely large number of ascriptions, explanations, and generalizations concerning

physical objects. Moreover, we are not merely concerned with the propositions that are *currently* put forward by physical science. As I will explain below, if we were to limit ourselves to present-day science it would be quite clear that CSP and physical science are logically independent of one another. If we want to know how mind and agency fit in the physical universe, we will have to ask about the logical relationship between the propositions of CSP and the actual physical *facts* about the world—i.e., the propositions ultimately put forward by a completed physical science.

Our task of evaluating the logical relationship between CSP and physical science would thus appear to involve some speculation about the progress of science. However, the task is mostly conceptual, for we will be concerned primarily with the differing vocabularies of CSP and physical science rather than the specific claims each makes. As has already been noted, on the face of it the crucial bits of CSP vocabulary do not seem to have any role in physical science. As will be argued in chapter 5, if physical science entails the facts of CSP then the psychological terms in CSP must have necessary and sufficient conditions that can be spelled out purely in the terms of physical science. Similarly, to show that CSP contradicts physical science one must claim that psychological terms in CSP at least have physically specifiable necessary conditions. If in fact there are physically specifiable necessary and sufficient conditions for psychological properties, this shows that either option 1 or option 2 is correct, and which one is correct is determined by the empirical question of whether there is anything that fulfills the necessary and sufficient conditions. But the logically prior question is whether psychological properties have physically specifiable necessary and sufficient conditions, and this is a largely conceptual question that can be addressed in a relatively a priori fashion.

To say that the problem is largely conceptual is not to say that it is easy. In recent decades most philosophers of mind have been, in effect, proponents of option 2. As has already been noted, this view requires that CSP vocabulary be brought into line with that of physical science. However, as will be discussed in part II, there are major obstacles to specifying necessary and sufficient conditions for mental terms.

On the other hand, according to option 3, facts about the mind are independent of physical science, and CSP in some sense has a different subject matter than does physical science. This might seem palatable if one is a substance dualist and thinks of the mind as a kind of nonphysical stuff. However, most philosophers have given up substance dualism and will acknowledge that human beings are composed of the same elementary particles as are the rest of the objects in the universe. This would *seem* to mean

that there are no limits on the scope of the subject matter of physical science; in particular, it is hard to see how CSP could carve out a distinct subject matter if the human beings that are its central concern are composed of physical particles. Moreover, given that human beings are composed of physical particles, it seems that, in principle, physical science must be able to explain and predict human behavior; but that is, in large part, also what CSP does. If CSP and physical science both give explanations for the same events, then it seems that either their answers must be ultimately the same (i.e., option 2) or their answers contradict one another (option 1); it is hard to see how their respective explanations could be logically independent of one another.

## 1.4 Why Care?

Why does it matter which option turns out to be correct? The question, as I have stated it, will not have antecedently concerned the average person. Is it merely a philosopher's puzzle, of no more intrinsic and general interest than the Sunday Times crossword? No. I think that there is a great deal at stake, even for the ordinary person who would not normally pay much attention to philosophical inquiry. In the previous two paragraphs, I quickly sketched reasons for thinking that each of option 2 and option 3 is untenable. If this turns out to be correct, then we are stuck with option 1 the claim that CSP and the ultimate claims put forward by physical science contradict one another. This would be a truly momentous conclusion. The ultimate claims put forward by physical science represent our best and most considered efforts at gaining knowledge. Giving up physical science would undermine any confidence we have in our ability to gain any knowledge about the world. On the other hand, if we are stuck with option 1 and we hold on to science as a reliable means of gaining knowledge, then our only alternative would be to deny the claims of CSP.

The denial of CSP would be an equally radical conclusion. We constantly engage in CSP in our interaction with one another; indeed, seeing one another as agents with minds seems obligatory in more than one sense. First, insofar as we wish to predict and explain the behavior of human beings, we are in practice forced to use the language of CSP. In principle, an omniscient physicist might be able to predict our bodily motions with nearly perfect accuracy based on knowledge of the state of every elementary particle in and around our bodies; in practice, of course, such a task is impossible. If you want to know what I will be doing tomorrow, it will work far better simply to ask me what my intentions are. Second, insofar as we

abandon the CSP approach to another human being, we give up treating that human being as a *person*. To see someone as a person is to see someone as an agent with a mind, and the language of mind and agency is precisely what CSP provides. To give up CSP would be to give up seeing ourselves and our compatriots as persons. It would mean acknowledging that nearly everything we care about (indeed, even the notion of caring itself) has been an illusion. If such a view is coherent at all, it must surely be seen as a desperate last resort, a view to be adopted only if it is conclusively demonstrated that the other two options are untenable.

In fact, one might be convinced that option 1 is so unattractive that *nothing* could ever rationally convince us of its truth; that is, one might argue that we have no theoretical commitments deeper than CSP and the general reliability of science, and hence there is no commitment from which we could launch an argument that would dislodge us from our belief in both CSP and physical science. If this is right, then either option 2 or option 3 *must* be right; i.e., it must be the case that CSP and physical science are not in conflict with one another. I have considerable sympathy for this line, and I firmly agree that it would take powerful considerations to convince us that CSP is wrong; however, it might be a bit hasty to rest on the claim that such considerations will never be put forward.

In any event, even if we could remove option 1 from consideration, there would still be a considerable philosophical mystery; the choice between option 2 and option 3 is not a technical matter of little significance. First, as I will argue much later in the book, other significant philosophical questions can turn on the question of option 2 versus option 3: I will argue that a certain brand of skepticism about moral value becomes much more plausible given option 2. (See chapter 11.) I also think that the problem of free will and determinism likewise looks rather different depending on which option is correct, but a treatment of that issue will have to await another time. Second, there is a genuine intellectual curiosity to be satisfied here, one concerning the nature of the universe and the place of mind and value within it. If option 2 is correct, then mental facts are, in the end, a species of physical fact; if physics can arrive at a grand unified theory of elementary particles and forces, then any facts about the mind will simply be entailed and subsumed by this theory. On the other hand, if option 3 is correct, then there are facts about the world that are logically independent of physical science, even if everything in the world is composed of physical particles. We would have to accept that CSP, in some sense, has a different realm—the realm of mind (or the space of reasons)—than does physical science. Facts about mind and value would not rest on the foundation of the

physical sciences. In particular, the normative notions that are central to CSP (e.g., *rational, irrational,* and *responsible*) would not rest on, or be reducible to, physical facts and laws.

In a related vein, Martin Luther King Jr. once claimed that we needed to rediscover the following principle:

. . . that all reality hinges on moral foundations. In other words, that this is a moral universe and that there are moral laws in the universe just as abiding as the physical laws. (King 1954)

Depending on the details, the proponent of option 3 might accept King's formulation. Regardless of whether one accepts King's claim, option 3 would require a more limited principle, resembling this: This is a rational universe, in the sense that there are normative principles of psychological explanation that are just as abiding as, and are independent of, the physical laws of the universe. This might sound supernatural and mystical to the strong naturalist, and indeed King's own comments came in an overtly religious context. I will argue that the proponent of option 3 need not be wedded to anything recognizably supernatural, but that option 3 still amounts to a very interesting claim about the nature of mind and the nature of the universe. Ultimately, there will also be a lesson here about the foundation-lessness and yet firmness of facts about mind and value.

### 1.5 Preview of the Teleological Realist Answer

I will defend option 3; that is, I will claim that the facts of CSP and physical science are logically independent. Thus, although I agree that human beings are composed of physical particles, I will claim that the facts about the mind are not ultimately a species of physical fact, and that they are not going to be subsumed within physical science.

The nature of psychological explanation will be paramount. Explanations of human behavior do not exhaust the content of CSP, but they do constitute an essential part of CSP. In physical science, explanations are typically causal explanations: We explain an event by giving its antecedent cause. For example, we explain an eclipse of the sun by citing the position of the moon (i.e., the moon's position and size is cited as the cause for the sunlight's not reaching particular portions of Earth), the general warming of the Northern Hemisphere during summer is explained by citing the tilt of Earth's axis as cause, and an increase in global temperatures generally is explained by citing the increase in the amount of certain greenhouse gases in the atmosphere. Most philosophers of mind assume that CSP explana-

tions are causal too; thus, when CSP says, "Vera went to the kitchen because she wanted tea," this is construed as making the claim that Vera's desire for tea caused her going to the kitchen. However, we also know that physical science will presumably produce a perfectly good causal explanation of Vera's movement to the kitchen, an explanation that will cite some brain state. If both common-sense and physical science are offering causal explanations of human behavior, it is hard to avoid seeing them as competitors. Either the common-sense explanation will somehow reduce to the scientific explanation (if, e.g., Vera's desire for tea just *is* the brain state cited in the physical explanation), or the two causal explanations will conflict with each other and at most one can be right. In other words, if common-sense psychological explanation is causal, then it is hard to avoid the conclusion that either the option 1 or option 2 must be correct. (See chapter 12.)

As the crucial part of my defense of the third option, I claim that CSP explanations are not causal; instead, I claim that they are teleological. A teleological explanation explains by citing the purpose or goal of the behavior in question; thus teleological explanations cite a future state of affairs toward which the behavior was directed, rather than an antecedent state that caused the behavior. Many common-sense explanations already have this form—e.g., "Vera went to the kitchen in order to get tea." "Vera went to the kitchen because she wanted tea" can be given an explicitly teleological construal as follows: "Vera went to the kitchen in order to bring it about that her desire for tea was satisfied." According to this account, commonsense explanations of behavior are simply answering a different question than are physical science explanations: in physical science we ask for the antecedent cause of the behavior, whereas when we give psychological explanations we ask for the end at which the behavior was directed. If, as I claim, teleological and causal explanations are independent of one another, then we can begin to see how the common-sense facts about mind and agency can be independent of the physical facts about human beings. I call the view teleological realism, and it is the subject of part III.