

1

The Fodorian Project

Introduction

Since the early 1960s Jerry Fodor has been engaged in developing a collection of distinct yet mutually reinforcing positions on a range of related issues and in doing so has set much of the agenda in the philosophy of mind and cognitive science. What binds together the various elements of Fodor's output is that they all count as engagements in a single project. In this chapter I will give an account of Fodor's project and a description and motivation of the commitments that underlie it.

Fodor has two basic commitments: one is to folk psychology and the other is to physicalism. These two commitments generate the task of providing a vindication of folk psychology within a physicalist framework. But what is folk psychology and what is physicalism? And why would anyone devote a career spanning several decades to the project of vindicating folk psychology within a physicalist framework? It is to these questions that I now turn.

Folk Psychology

Folk psychology (otherwise known as common-sense psychology) is a descriptive, explanatory and predictive practice that is central to human life. We engage in this practice whenever we interact with our fellows and we would be at sea in the social world without it. Folk psychology is bound up with a conception of the human

individual according to which we are minded and engage in behaviour that is influenced by our mental states. Here are some examples of typical engagements in folk psychology.

- I hear Edgar utter the English sentence ‘Fang is ferocious and if there is one thing that I don’t like it’s a ferocious dog’. I immediately jump to the conclusion that Edgar believes that Fang is ferocious and that, in all probability, he wants to avoid any contact with Fang.
- Whilst running with Edgar, I hear a gruff bark. Edgar gasps ‘Oh no! It’s Fang’ and suddenly veers off the path and dives behind a nearby bush, where he remains cowering for some time. Later I am asked why Edgar behaved in such a strange manner. I answer by reporting that Edgar, believing Fang to be ferocious, wants to avoid Fang and that, having heard a gruff bark, he believed that Fang was in the vicinity.
- On being asked why I didn’t follow Edgar’s Fang-avoidance tactic, I point out that though I believe Fang to be ferocious and want to avoid close contact with him, I also believe that my running prowess is such that Fang has little chance of catching me.
- On another occasion I am running ahead of Edgar. I see Fang in the distance heading our way and predict that, shortly, Edgar will engage in hiding behaviour.
- I fear that Fang’s menacing behaviour is going to put Edgar off running altogether. So, in order to keep my running companion, I devise a new route lying outside the beast’s territory and tell Edgar about my plans and the Fang-free nature of the new route.

In the above examples I did various things: I described some of my own mental states by means of language; I came to a conclusion about the mental states of one of my fellows on the basis of seeing him behave in a particular way; I explained and predicted the behaviour of one of my fellows on the basis of the mental states that I took him to have; I engaged in behaviour (including linguistic behaviour) with a view to influencing the mental states and subsequent behaviour of one of my fellows. Such activity is the very stuff of folk psychology and it is difficult to see how I – or anyone else – could explain, predict and manipulate the behaviour of my fellow humans as effortlessly and effectively by any other means.

The mental states that figured in the above examples were beliefs and desires. Beliefs and desires (along with intentions, expectations,

hopes, fears etc.) belong to an important category of mental states known collectively as the intentional states.¹ In this context 'intentional' is a technical term and does not mean what it means in everyday discourse. To describe an individual as having an intentional state is not to say that she deliberately has that state. Moreover, although intentions are a species of intentional state, not all intentional states are intentions.

Folk psychology is committed to the existence of a wide range of mental states and events that are to be distinguished from the intentional states. In addition to beliefs, desires and the like, folk psychologists routinely attribute perceptual experiences, sensations, feelings, emotions, moods, personality traits and so on to themselves and to their fellows. Hence, as far as folk psychology is concerned, the mental realm is populated by a heterogeneous collection of phenomena. However, it is arguable that the intentional states occupy a central role in folk-psychological practice so that it is only an exaggeration – as opposed to an outright misrepresentation – to describe folk psychology as being an intentional psychology.

Intentional states have the following important features, all of which are emphasized by Fodor. (See Fodor 1987, ch. 1.)

Intentional states have semantic properties; in particular, they have meaning or content and, correlatively, they can be about particular objects and possible states of affairs and they have truth or satisfaction conditions. Consider the state of believing that Fang is ferocious. It is central to the identity of this state that it has a particular meaning or content: namely, the content expressed by the English sentence 'Fang is ferocious'. Alternatively, one might say that it is central to the identity of this state that it is about Fang, that it represents him as being ferocious and that it is true if and only if Fang is ferocious. Folk psychology distinguishes between intentional states that differ in terms of their content. For example, the belief that Fang is ferocious is a different belief than the belief that Fang enjoys savaging joggers.²

Intentional states can be about objects and phenomena that do not exist. For example, it is possible to have beliefs about Santa Claus, ghosts, unicorns and phlogiston.

Just as one can believe that Fang is ferocious, one can desire that it be the case that Fang is ferocious, intend to make it the case that Fang is ferocious and so on. What distinguishes these intentional states is not their respective contents (for they all have the same content). Rather, what distinguishes them is the relation involved: believing that Fang is ferocious involves standing in the belief

relation to the content expressed by the English sentence 'Fang is ferocious', whereas desiring that it be the case that Fang is ferocious involves standing in the desire relation to the same content. In short, whenever an individual has an intentional state, there is a relation and a content involved and folk psychology classifies or individuates intentional states in terms of both their content and the relation involved. Thus, when engaged in the folk-psychological practice of attributing an intentional state to one of our fellows, we will typically specify both a relation and a content ('I believe that Fang is ferocious', 'He wants it to be the case that Fang is impounded', 'She hopes that Fang does not spot her').

Intentional states are not causally inert; on the contrary, they are very much part of the causal fray. There are three distinct types of causal process in which intentional states figure prominently. First, intentional states are often caused by environmental factors impinging on an individual, as when a display of Fang's ferocious behaviour causes me to believe that Fang is ferocious. Second, intentional states often cause intentional states, as when my belief that Fang is ferocious causes me to want to avoid all contact with Fang. Third, intentional states often cause behaviour, as when Edgar's belief that Fang is in the vicinity causes him to crouch behind the nearest bush. Consequently, folk-psychological explanations of intentional states and behaviour are causal explanations. The manner in which intentional states causally interact with environmental impingements, other intentional states, and behaviour is not random and undisciplined; rather it is regular or law-governed. Hence, there exists a large battery of counterfactual and hypothetical supporting generalizations relating intentional states to one another, to environmental impingements and to behaviour.³

There is a systematic relationship between the causal powers of intentional states and their semantic properties; intentional states tend to cause intentional states and behaviour to which they are semantically related. Suppose that an individual faced by a snarling Fang wanted to scare Fang away and believed that the best way of doing this would be to start shrieking at the top of her voice. Then, typically, this belief and desire pair would interact to cause her to form the intention to start shrieking at the top of her voice and this intention would subsequently cause her to so behave. In this case, an intention and subsequent behaviour are caused by a belief and desire pair to which they are coherently related. This would not be the case were the intention and behaviour caused by a desire to see an aardvark and a belief that Kilimanjaro is the highest mountain

in Africa. The process of reasoning constitutes another graphic case of this general kind. When a subject reasons, she will have a thought that causes her to have another thought that in turn causes her to have a third thought, and so on. Typically, the contents of the members of the resultant chain of thoughts will be related to one another in a way that mirrors the relationship between the propositions of a logically valid argument. As Fodor puts it, 'one of the most striking facts about the cognitive mind as commonsense belief/desire psychology conceives it [is] . . . the frequent similarity between trains of thoughts and *arguments*' (1987: 13).

A natural language such as English is productive in the respect that there are infinitely many distinct sentences of English that a competent speaker of that language is capable of constructing and understanding. A similar point can be made about intentional states: there are infinitely many distinct thoughts that a normal human subject is capable of having. Moreover, it is commonplace for an individual to form an intentional state that she has never had before. For example, the sentence 'a ferocious dog has never been seen swallowing a double-decker bus' is likely to cause most readers to entertain a thought that they have never had before. In thinking such a thought for the first time, an individual will typically deploy concepts that she has a prior grasp of but arrange them in a new way so as to express a content that she has never previously entertained.

The system of intentional states is systematic in the respect that any human subject capable of believing (or desiring, or intending, etc.) that object *a* stands in the relation *R* to object *b* (for any relation *R* and any objects *a* and *b*) is also capable of believing that *b* stands in the relation *R* to *a*. For example, any human subject capable of believing that the ferocious dog savaged the jogger is capable of believing that the jogger savaged the ferocious dog.

It is possible to believe that Mark Twain was the greatest American writer of the nineteenth century without believing that Sam Clemens was the greatest American writer of the nineteenth century despite the fact that Mark Twain and Sam Clemens were one and the same person. Similarly, it is possible to believe that water is wet without believing that H₂O is wet despite the fact that water and H₂O are one and the same substance. Put generally, it is possible to believe (or desire, or intend, etc.) that *a* is *F* without believing (or desiring, or intending, etc.) that *b* is *F* despite the fact that *a* = *b*. How is this possible? Because the belief that Mark Twain was the greatest American writer of the nineteenth century is a different

belief from the belief that Sam Clemens was the greatest American writer of the nineteenth century and an individual can have either of these beliefs without being aware of the fact that Mark Twain and Sam Clemens were one and the same person. This fact has a linguistic consequence: within sentences that ascribe an intentional state to a subject, co-referential expressions are not generally interchangeable *salva veritate* (that is, without a change of truth value). For example, the sentence 'Edgar believes that Mark Twain was the greatest American writer of the nineteenth century' is true. But the sentence generated from that sentence by replacing the expression 'Mark Twain' with the co-referential expression 'Sam Clemens' is false. This feature of sentences that ascribe intentional states to individuals is known as their intensionality.⁴

Sensations such as pains, aches, itches and tickles have a qualitative character, as do experiences or feelings of joy, disappointment, frustration and the like. To use Thomas Nagel's (1974) memorable expression, there is something that it is like to have a sensation or a feeling and what it is like to have a sensation or a feeling of a particular type is different from what it is like to have a sensation or a feeling of any other type. Indeed, it seems plausible to say that sensations and feelings are at least partly individuated in terms of their qualitative character so that a subject's sensation isn't (for example) a pain or her feeling an experience of joy if it doesn't have the appropriate qualitative character. Matters are quite different with respect to intentional states for they do not have any distinctive qualitative character; for example, there is nothing in particular that it is like to believe that Fang is ferocious. That is not to say that experiences with a distinctive qualitative character are never associated with intentional states. It is perfectly possible for Fang's new neighbour to experience a sinking feeling when she comes to believe that Fang is ferocious. But such an experience would appear to be an effect of the acquisition of the belief rather than part of the belief itself, and someone else could acquire just the same belief whilst having a different resultant experience or whilst having no associated experience at all.

How do we engage in folk psychology? How do we so effortlessly and effectively construct intentional explanations and predictions of the intentional states and behaviour of our fellows? Fodor's answer invokes the idea that folk psychology is a theory akin to a scientific theory. (See Fodor 1987, ch. 1.) A scientific theory is a collection of sentences. These sentences contain terms that refer

to observable phenomena. In addition, they contain terms that refer to unobservable entities postulated by the theorist in order to explain the observable phenomena that fall within the domain of the theory. Such terms are known as theoretical terms and familiar examples are 'gene', 'virus', 'electron', 'quark' etc. The sentences that constitute a theory specify generalizations concerning the relations (particularly the causal relations) that are taken to hold between the theoretical entities and observable phenomena referred to by the terms of the theory.

It is widely held that the core sentences of a scientific theory implicitly define the theoretical terms of the theory en masse (Lewis 1970). So, for example, the core sentences of genetic theory implicitly define the term 'gene' or, in other words, specify what it is to be a gene. Consequently, if the core sentences of a theory are largely false – in that the world doesn't contain a collection of entities whose causal relations to one another and to observable phenomena correspond to that described by those sentences – then the entities postulated by the theory have no existence.

Fodor takes folk psychology to be a theory in the respect that it consists of a large collection of sentences that assert the existence of a network of causal generalizations relating mental phenomena, environmental impingements and observable behaviour. The mental terms that appear in these sentences (for example, 'belief', 'desire' etc.) are theoretical terms that refer to unobservable phenomena postulated by the theory in order to explain observable behaviour.

The view that folk psychology is a theory, though not without its detractors,⁵ is widely held in the philosophical community and has been dubbed the theory-theory (Morton 1980). An early and important expression of the theory-theory can be found in the work of Wilfred Sellars (1956). Another important element of the theory-theory is the idea that folk psychologists have knowledge of the generalizations that constitute the theory of folk psychology and that this knowledge is employed in the construction of folk-psychological explanations and predictions. At first appearances, this idea might sound somewhat implausible. After all, few of us would be able to specify more than a small number of banal psychological generalizations. However, consider the case of our linguistic capacities. Few speakers of English are able to specify any of the grammatical rules of that language. Nevertheless, their speech largely consists of grammatically correct sentences and they are pretty effective when it comes to the task of determining

whether any given string of English words constitutes a grammatically correct sentence. How, one might ask, could speakers of English have such capacities unless there was some respect in which they knew the grammatical rules of their language? Chomsky has famously argued (see Chomsky 1986, for example) that the best explanation of such capacities is that speakers of natural languages have a form of unconscious knowledge of the grammar of the language that they speak. Such unconscious knowledge is known as tacit knowledge and is held to be encoded in the brain. Fodor thinks that an analogous claim holds with respect to folk psychology. The best explanation of our ability to explain and predict the behaviour of our fellows is that we have a tacit knowledge of a rich battery of causal generalizations relating intentional states to one another and to behaviour. Moreover, just as Chomsky holds that human subjects have an innate knowledge of universal principles of human language, Fodor hypothesizes that much of our folk-psychological knowledge is innate. (See e.g. Fodor 1978b.)

Fodor cites the following as a paradigmatic example of a folk-psychological generalization: 'If X wants P, and believes that not-P unless Q, then, *ceteris paribus*, X tries to bring it about that Q' (1987: 13). Another prominent example is this: if X intends to do P then, *ceteris paribus*, X will do P. Note that these generalizations quantify over intentional states and behaviours that are semantically related. This reflects the above described fact that the causal powers of intentional states tend to be coherently related to their content. In virtue of this feature of such generalizations, a typical folk-psychological explanation of a subject's having a particular intentional state or behaving in a particular way will appeal to intentional states that are semantically related to the target state or behaviour. Hence, such explanations are not merely causal explanations; in addition, they rationalize or make sense of the subject's intentional states and behaviour. For example, when I explain Edgar's jumping behind a bush in terms of his desire to hide from Fang and his belief that, in the circumstances, the best way to hide from Fang would be to jump behind the bush, I reveal Edgar's behaviour to be sensible or rational in the light of its intentional causes.

What are we to make of the fact that the phrase '*ceteris paribus*' appears in the above examples of folk-psychological generalizations? '*Ceteris paribus*' is a synonym of 'all else equal'. Hence, generalizations that contain *ceteris paribus* clauses are hedged or qualified rather than universal. Fodor accepts that folk-psychological generalizations invariably contain *ceteris paribus* clauses, but he

denies that this tells against their explanatory or predictive value. (See Fodor 1991b.) For, it is typical of the special sciences (that is, the sciences other than physics) that their generalizations are hedged as is indicated by the geological generalization that a meandering river erodes its outer bank, *ceteris paribus*. I will have much more to say about *ceteris paribus* generalizations and how they can be of explanatory value in due course.

Fodor is an enthusiastic champion of folk psychology. For him we really do have beliefs, desires and the like, and these mental states have the properties that folk psychology takes them to have. Moreover, the generalizations of folk psychology are largely true.⁶ This is not to say that folk psychology captures all the facts about our mental lives. No doubt there are mental states and generalizations concerning our mental lives that folk psychology is blind to. And folk psychology doesn't have the power to explain all of the mental phenomena with which we are acquainted. (For example, folk psychology does not seem to provide us with a particularly effective way of explaining perceptual illusion, mental illness, depression or intellectual ability.) However, when it comes to explaining and predicting everyday human behaviour, folk psychology has no peers. Here is a case of Fodor waxing lyrical.

[Folk] psychology works so well it disappears. It's like those mythical Rolls Royce cars whose engines are sealed when they leave the factory; only it's better because they aren't mythical. Someone I don't know phones me at my office in New York from – as it might be – Arizona. 'Would you like to lecture here next Tuesday?' are the words he utters. 'Yes thank you. I'll be at your airport on the 3 p.m. flight' are the words that I reply. That's *all* that happens, but it's more than enough; the rest of the burden of predicting behaviour – of bridging the gap between utterances and actions – is routinely taken up by the theory. And the theory works so well that several days later . . . and several thousand miles away, there I am at the airport and there he is to meet me. Or if I *don't* turn up, it's less likely that the theory failed than that something went wrong with the airline. (1987: 3)

In characterizing folk psychology as a body of true causal generalizations that are of considerable explanatory and predictive power, that appeal to unobservable entities and that contain *ceteris paribus* clauses, Fodor is explicitly comparing folk psychology to a special science. Moreover, he is suggesting that a respectable scientific psychology would, at the very least, bear much in common with folk psychology.

Physicalism

Fodor's second basic commitment is to a metaphysical position known as physicalism. In a nutshell, physicalism is the doctrine that reality is ultimately physical in nature so that the sum totality of physical facts determines the sum totality of facts. But what does that mean? What is the physical and in what respect are physical facts supposed to determine all other facts? A good way of approaching these questions involves reflecting on the structure of science.

There are many distinct sciences including physics, chemistry, biology, geology, neurophysiology and (perhaps) psychology. Each science has its own particular domain of enquiry and, therefore, is concerned with explaining a distinctive range of phenomena. To this end, each science will appeal to a range of objects, states, events and processes that are described and classified in terms of their possession of a distinctive range of properties. Hence, each science has its own proprietary vocabulary. For example, physics talks of atoms, fields, quarks, photons and the like, but not of living organisms, genes, viruses or (biological) reproduction. And just as biology doesn't talk of atoms, fields, quarks and photons, geology doesn't talk of living organisms, genes, viruses or reproduction.⁷

The sciences other than physics are known as the special sciences. It would appear that there is a respect in which all special science phenomena inhabit the physical domain, that is, the domain of physics. Consider the one pound coin located in my pocket. It is identical to, or constituted by, a particular physical object, an object that is nothing more than a collection of physical particles standing in certain physical relations to one another. Hence, the coin can be described in physical terms, that is, in terms of such physical properties as its mass, its micro-physical structure and the like. Such a description would constitute a more basic or fundamental description of the coin's nature than one that talked of its monetary value or its purchasing power. The latter description would specify some of the higher-level properties of the coin. There is nothing special about my coin; all coins are physical objects. But not vice versa, for there are plenty of physical objects that are not identical to, or do not constitute, a coin. Consequently, a complete description of the nature of the world at the physical level would provide a specification of the fundamental nature of all the coins that exist. Moreover, it would specify the fundamental nature of a domain of objects

that was much wider than the domain of coins but one that the domain of coins belongs to or inhabits.

What is true of coins would also appear to be true of biological phenomena. Each and every biological object, state, event or process (be it a living organism, a gene, a viral infection, a mating or whatever) would appear to be identical to, or constituted by, a physical object, state, event or process. Hence, anything that is biological will also have a physical description but not vice versa (as there are plenty of physical things that are not identical to, or do not constitute, anything biological). And a physical description of, say, a living organism will constitute a more basic or fundamental description of that thing than would a biological description. Whenever a biological event causes some other biological event, a physical event that is identical to, or constitutes, the biological cause will set off a causal process that eventuates in a physical event that is identical to, or constitutes, the biological effect. The course that this physical causal process follows will be determined by the laws of physics. Hence, biological causal processes and any biological causal generalizations or laws that there are will be underpinned by the laws of physics.

This point about biology can be generalized to all of the special sciences: each and every special science object, state, event or process is identical to, or constituted by, a physical object, state, event or process; special science causal processes and laws are underpinned by physical causal processes and the laws of physics; and a complete physics would provide the most basic and general description of the scientific domain.⁸

What is the relationship between physical properties, on the one hand, and special science properties, on the other? Given that special science phenomena inhabit the physical world and that that world is at bottom governed by the laws of physics, if there are any special science laws then there must be a systematic and disciplined relationship between physical and special science properties. In the absence of such a relationship, the behaviour of phenomena at the special science level would be entirely random, in which case the special sciences would not be sciences at all.⁹

We can distinguish between two types of disciplined and systematic relationship that might hold between physical and special science properties, both of which would appear to enable physical laws to underpin higher-level special science laws. First, special science properties might be identical to physical properties so that any true statement expressing an identity relationship between a

particular special science property and some physical property will lay bare the ultimate nature or essence of that property. There would appear to be examples of such a relationship: water is H_2O , that is, the property of being (a sample of) water is identical to the property of being a collection of H_2O molecules, and it is often claimed that heat is mean kinetic energy. However, it would appear that such examples of property identity are thin on the ground. Consider an example from outside the scientific realm. It is difficult to envisage a non-physical can-opener; each and every can-opener is a physical thing of some description. However, the property of being a can-opener cannot be identified with any physical property for the simple reason that the class of can-openers constitutes a physically heterogeneous bunch. (Think of how widely can-openers can differ in terms of their shape and size, what they are made of, their principles of operation, and the like.) Suppose that the human species had been particularly unadventurous when it came to developing the can-opener so that all can-openers were made out of the same material to the same design. It would then be true that all the can-openers that existed were physically alike so that the property of being a can-opener was coextensive with a particular (perhaps conjunctive) physical property. However, it would not be true that the property of being a can-opener was identical to that physical property; the fact that these properties were coextensive would be a mere accident.

What is true of can-openers is true of many special science phenomena. It is difficult to envisage how anything could be a heart or a mountain (to pick just two examples) without being a physical thing. But notice how physically different the heart of a frog is from the heart of a human and how much mountains can vary in terms of their shape, size, mass and micro-physical constitution. Moreover, had the world been such that all actual existing hearts were physically similar, one could still coherently imagine the discovery or evolution of a heart that bore no physical similarity to all the other hearts that had previously existed or been known about. What these reflections suggest is that the property of being a heart and that of being a mountain, though perfectly respectable special science properties, cannot be identified with any physical property. In this respect, these properties are hardly atypical special science properties.

The second type of disciplined relationship that might exist between special science and physical properties is that of supervenience: special science properties might supervene upon physical

properties. Supervenience is a relationship that holds between distinct families of properties and is somewhat weaker than that of identity.¹⁰ Consider another example from outside the scientific realm. Fang is unpleasant. But why, one might ask, is Fang unpleasant? This question is ambiguous. On the one hand, it could be a question about the causal origins of Fang's being unpleasant. On this reading, the question would be answered by an appeal to the nature of Fang's upbringing, his history or his genetic make-up. On the other hand, it might be a question concerning what it is about Fang in virtue of which he merits the description 'is unpleasant'. This is not a causal question; rather, it is a question as to which lower-level or more basic properties of Fang are responsible for his being unpleasant. It is a question that is answered by describing the way in which Fang typically behaves; by pointing out his tendency to chase runners, snarl at other dogs, drool and dribble whilst eating and so on. These behavioural properties are such that if you have them you thereby have the property of being unpleasant. In short, these properties are capable of generating and sustaining the property of being unpleasant so that having them is sufficient for having the property of being unpleasant. However, it is not necessary to behave as Fang does in order to be unpleasant. Edgar is thoroughly unpleasant despite the fact that he doesn't chase people, never snarls and has impeccable table manners. Why then, in the non-causal sense, is Edgar unpleasant? What is it about Edgar in virtue of which he can be described as being unpleasant? Again, the answer has to do with how Edgar behaves, but in his case it is a different collection of behavioural properties that generates and sustains his unpleasantness. Edgar has a tendency to ignore strangers, boorishly monopolize conversation and ridicule those who express views that he does not share. That is why he is unpleasant. Hence, there are different collections of behavioural properties that are equally capable of generating the property of being unpleasant. In order to be unpleasant, two individuals need not have the same behavioural properties, but if they have the same behavioural properties then if one is unpleasant so must be the other.

The relationship that the property of being unpleasant bears to lower-level behavioural properties of the sort had by Fang and Edgar is one of supervenience. Put generally, supervenience is a relation of non-causal determination holding between distinct families of properties. B properties supervene upon A properties if and only if the A properties that an object has fix or determine its B properties. If B properties supervene upon A properties then two objects

that agree with respect to their A properties will thereby have just the same B properties. However, as in the above example, it is possible for objects to share a B property without sharing any A properties.

There is a distinction between different types of supervenience relation. For example, there is a distinction between weak supervenience and strong supervenience. If A properties weakly supervene upon B properties then any two objects inhabiting the same possible world that agree in their B properties will thereby agree in their A properties. The existence of such a relation is consistent with the possibility of two objects that inhabit different possible worlds agreeing in their B properties yet having different A properties. If A properties strongly supervene on B properties then any two objects that agree in their B properties will thereby agree in their A properties regardless of whether or not they inhabit the same possible world. In what follows I will talk about supervenience in general and will ignore the question of what specific type of supervenience relation (for example, strong, weak or whatever) must hold between mental and physical properties for physicalism to be true.

Having described the relations of identity and supervenience, we can give a general account of the structure of science in the following terms. Physics is the most basic and general science. Special science phenomena are physical in the respect that each and every special science object, state, event and process is identical to, or constituted by, some physical phenomenon. Special science properties are either identical to or supervene upon physical properties. And the laws of physics underpin all the special science laws. This account of the structure of science, though not mandatory, is certainly the received view in the philosophical community and is endorsed by Fodor.

We are now in a position to characterize physicalism. Physicalism is the view that reality in general – and not just the scientific domain – is at bottom physical in nature so that all properties that are genuinely instantiated ultimately supervene upon physical properties in such a way that the sum totality of physical facts determines the sum totality of facts. It is important to appreciate that the physicalist need not hold that all real properties are scientifically respectable or that science is the only source of knowledge. For, there can be properties that ultimately supervene upon the physical yet do not figure in any laws. The scientist will have no interest in such properties or the facts involving their instantiation. Con-

sider the property of having a degree. This is a perfectly real property that many people have. This property features in a whole battery of rough-and-ready generalizations (people with degrees tend to be of higher than average intelligence; people with degrees tend to earn more than those who have no academic qualifications; people with degrees tend to be over the age of 5), but such generalizations don't have the status of scientific laws. Consequently, the property of having a degree is not one that any science cares about and knowledge of the geographical and social background, the economic well-being and the leisure activities of those people who have degrees does not count as scientific knowledge despite the fact that such knowledge may well be very interesting and important. Moreover, degrees are not physical things; there is no physical thing that can be identified with my degree despite the fact that we talk about degrees as if they were objects (although there is a physical thing that can be identified with my degree certificate). None of this need pose any great problem for the physicalist. For the physicalist might argue that it is physical things (namely human individuals) that have degrees and that they have them in virtue of facts about their history and facts concerning certain social institutions and practices, facts that ultimately hold in virtue of how the world is at the physical level. Therefore, physicalism, though bound up with a healthy respect for science and a commitment to the primacy of physics and the physical, is not a crudely scientific position.

In developing an account of the nature of physicalism I have freely written of physical properties, facts and phenomena. But how, one might ask, are we to understand the term 'physical'? It has been objected that those who call themselves 'physicalists' face major problems when it comes to formulating their position in such a way that it comes out as both plausible and non-trivial (see Crane and Mellor 1990). An easy answer would be to say that physical properties, facts and phenomena are the kind that the science of physics cares about; the kind that are referred to in the statements of the laws of physics and that are appealed to in the descriptions, explanations and predictions proffered by physicists. No doubt, there are many physical properties and phenomena that contemporary physics is ignorant of. And, I dare say, some of the properties and phenomena that contemporary physicists appeal to are as real as phlogiston and caloric fluid. So, perhaps it would be more accurate to characterize physical properties and phenomena as those that would be recognized by some future true and complete science that was a direct descendant of contemporary physics. Of

course, this characterization of physical properties is a little vague until an account is given of the descendancy relation that it appeals to. I am not going to attempt to provide such an account as, for the purposes at hand, it should be clear enough what it is to be a physical property and there is no reason to believe that the required account could not be given.

Physicalism and the Mental

Physicalism implies that human individuals are physically constituted and that they inhabit a world that is, at bottom, physical in nature. Therefore, if we really have mental states then those states must be states of a complex physical system. Moreover, mental properties must either be identical to or supervene upon physical properties so that the mental states that an individual has will be a product of her physical nature and/or the physical relations that she bears to external physical phenomena.

How, one might ask, can the mental be related to the physical in this kind of way? One might ask a similar question about biological or geological phenomena, but such a question doesn't appear to be quite as pressing. No doubt the details are very complicated, but it is not an affront to common sense to assert that biological and geological phenomena are physically constituted and that if something has any biological or geological properties then, ultimately, it has those properties in virtue of the way the world is at the physical level. However, matters are somewhat different with respect to the mental. An important class of mental states, namely the intentional states, have a number of salient characteristics that make it difficult to see how the truth of physicalism could be consistent with the existence of such states. To recap, here are just three of those characteristics. Intentional states have content; they represent particular objects and states of affairs and sometimes represent objects that do not exist and states of affairs that do not hold. Intentional states have causal powers that cohere with their content so that causal processes involving intentional states are typically rational. A human subject is capable of having infinitely many content-distinct intentional states. How, one might reasonably ask, could a finite physical system (in virtue of its physical nature and/or its physical relations to other physical phenomena) have states that have content, be capable of having infinitely many such states and be rational to boot? Throughout the ages, many philosophers have

had strong intuitions to the effect that this question has no positive answer: a mere physical system could not have intentional states. Therefore, the physicalist is faced with a dilemma. Grasping the first horn of this dilemma involves denying the reality of the mental; we do not have beliefs, desires and the like and so we do not instantiate mental properties. Grasping the second horn of the dilemma involves constructing a physicalist theory of mind that shows how a physical system could have intentional states and have such states in virtue of its physical nature and/or physical relations to other physical phenomena. As we have seen, Fodor is an enthusiast of folk psychology and is therefore committed to grasping the second horn of the physicalist dilemma. Indeed, constructing a physicalist account of how we could be minded in the way that folk psychology conceives us to be is the central task of Fodor's philosophical project. In other words, his basic aim has been to vindicate folk psychology within the framework of a physicalist world-view. And given that he conceives of folk psychology as being very much like a scientific theory, he has thereby been engaged in the project of showing how there could be a respectable scientific psychology that is – at the very least – a close relation of folk psychology.

Are mental realists forced to be physicalists? Is there any mileage in the view that mental phenomena are perfectly real yet do not inhabit the physical domain? Physicalists typically appeal to causal considerations when objecting to such a view. A good example of such an argument runs as follows. Mental phenomena often have physical effects, as when a desire for ice-cream along with a belief that there is ice-cream in the freezer causes an individual to open the freezer door and reach inside. Whenever an individual acts, a physical event involving the movement of the individual's body takes place, and if intentional states were not capable of causing such physical events then there would be little point in having a mind. If intentional states did not inhabit the physical realm then it would be a mystery as to how they could cause actions. But there would be no mystery were intentional states physical. Moreover, it is a principle of physics that the physical world is causally closed, that is, that every physical event has a physical cause that is sufficient to determine its occurrence (or determine the probability of its occurring). If intentional states were not physical then every action would have two independent causes sufficient to determine its occurrence. On the one hand, it would have a physical cause, and on the other, it would have a mental cause. In other words, actions would be systematically overdetermined. But is it plausible that

actions are systematically overdetermined? To paraphrase Steven Schiffer (1987), surely God is not such a bad engineer.

Indeed, the idea that actions are systematically overdetermined strikes me as being close to incoherent. It is part of the nature of causation that if A is the cause of B then had A not happened, all else equal, neither would have B. Therefore, for a mental state to be the cause of an action it would have to be the case that if the subject had not had that mental state then, all else equal, she would not have executed the action in question. But if actions are systematically overdetermined then putative mental causes will never satisfy this condition for the simple reason that had a subject not had the mental state that apparently causes her action then she would still have so acted as the result of one of her physical states. This argument applies equally to putative physical causes of actions. Therefore, the overdetermination thesis implies that neither mental states nor physical states are causes of our actions. In short, mental states had better inhabit the physical realm if they are to be the causes of our actions.

A response open to the anti-physicalist would be to deny the principle of the causal closure of the physical realm so that a physical event need not have a sufficient physical cause thereby freeing the way for actions to have sufficient mental causes. In addition, it might be pointed out that the principle of the causal closure of the physical realm is not a law of physics and is therefore hardly sacrosanct. The problem with this line of thought is that it is difficult to see how there could be any laws of physics if the principle did not hold. Suppose that it is a putative law of physics that whenever an event of physical type P occurs (perhaps in certain physically specifiable circumstances) it causes an event of physical type P* to occur. If non-physical events are capable of causing physical events that themselves have no physical cause then what is to stop such a non-physical event from intervening when an event of type P occurs so preventing the occurrence an event of type P*? In principle, such an interference would be perfectly possible and, clearly, could not be ruled out by the laws of physics. Thus, it would be a mere accident that Ps generally caused P*s rather than a law. To avoid denying that there are any laws of physics, the anti-physicalist is forced to postulate the existence of mechanisms that prevent the interference of non-physical phenomena whenever the antecedent of a physical law is satisfied. An example of such a mechanism would be a psychophysical law to the effect that whenever a P occurred it had an effect in the mental realm that (given the laws operative in that

domain) prevented the occurrence of any mental event of a type that was capable of blocking the occurrence of an event of type P*. In short then, if one rejects the principle of the causal closure of the physical domain, one is forced either to deny that there are any physical laws or to commit oneself to the existence of a battery of psychophysical laws. Neither of these options strikes me as very attractive.¹¹

Conclusion

In this chapter I have given an account of the nature of folk psychology as conceived by Fodor along with an account of the metaphysical doctrine of physicalism. Fodor is a great champion of folk psychology. For him, the entities that it postulates are real and the generalizations that it takes to hold are largely true and are of considerable explanatory and predictive power. Fodor is also a physicalist. Thus, he faces the problem of squaring his enthusiasm for folk psychology with his physicalism. In other words, he faces the problem of explaining in physicalist terms how we could be as folk psychology represents us as being; that is, of explaining in physicalist terms how we could have intentional states that are governed by the generalizations of folk psychology. Solving this problem, and so vindicating folk psychology within a physicalist framework, has been Fodor's major philosophical project.