The problem of higher-level causality raised by Pettit is particularly difficult at the psychological level. Common sense explains behaviour by invoking beliefs and desires. More specifically, it sees the *content* of an agent's mental states as causing his or her behaviour. However, content does not fit easily in a physicalist view of the world. Jacob, in 'The role of content in the explanation of behaviour', reviews the main proposals developed in contemporary philosophy of mind regarding the causal role of content. He shows that all these proposals have deep problems. A clear implication of Jacob's discussions is that we do not really know to what extent and by what means common-sense psychological understanding ('theory of mind') succeeds in being genuine understanding.

D.S.

13

Causality at higher levels

PHILIP PETTIT

There are two different problems to which causality gives rise from the point of view of the psychological and social sciences, and indeed from the point of view of common sense—from the point of view, for example, of the commonplace psychology and sociology that we practise in everyday life. One is the epistemological problem as to how causality is understood and detected within these higher-level forms of inquiry. The other is the more fundamental ontological problem as to whether there is really anything there at all for these forms of research to unearth—whether there is really such a thing as higher-level causality. This chapter deals with the ontological issue and bears only indirectly on the epistemological question.

The ontological problem of concern is not motivated by a general scepticism about causality—by a scepticism about whether there are causal relations to be uncovered by any sciences, natural or psychological or social. Rather, it is motivated by a physicalist outlook on the world. Suppose that we are physicalists and believe that the empirical world moves to a causal rhythm that physics is in the best position to identify.* Does that mean that we have to condemn the psychological and social sciences, and indeed natural sciences like chemistry and biology, to a secondary role? Does it mean that such special sciences, as they are sometimes called, are not in a position to identify any causal dynamics in the empirical world? Does it mean that the factors that they identify are only of ephiphenomenal significance? (See Block (1990), Blackburn (1991), Jacob (1991–2), Macdonald (1992), Pettit (1992), and Yablo (1992).) That is the question that is addressed in this chapter.

The chapter is in three sections. The first formulates the physicalist

* It should be noticed that the physicalist may believe this even if he thinks, as many do, that the category of causality is not strictly needed in physics. Whatever the provenance of this category, it is assumed here that the physicalist will think that it applies in the microphysical realm, if he thinks that it applies anywhere. Certainly, it would seem to be applicable with entities at the level of electrons and protons.

doctrine against the background of which the question arises.* The architecture of instrumental control—control at different levels—that this physicalism allows is discussed in the second section. The third section examines how far this architecture of control is mirrored in an architecture of causality: how far causality, like control, can be found at higher levels. It is argued that we can reasonably countenance higher-level as well as lower-level causality and that the special sciences are not at any particular disadvantage in the exploration of causal matters.

PHYSICALISM

There are two pairs of claims that the physicalist envisaged will make. They give expression, respectively, to the following two ideas: first, the empirical world is constituted out of materials that physics is in the best position to identify; second, the empirical world is governed by forces or regularites that physics is best equipped to describe. The two pairs of claims articulate a picture of the world as at once physically constituted and physically governed (cf. Crane and Mellor 1989).

Claim 1. There are microphysical entities

A. There is an empirical world of the sort that physics posits.

B. Different kinds of thing in the empirical world share (subatomic) levels of composition of the kind that physics—specifically, microphysics—posits: there is a realm of smaller and simpler, microphysical entities.

The first part of this claim gives expression to a realist view of physics, under which the physicist is in the business of telling us about an objective world and the business is a potentially successful enterprise—there really is a world there for the physicist to chart. The second part of the claim directs us to the fact that physics, unlike the other sciences, has comprehensive as well as realist ambitions (on this matter see Papineau (1990) and Crane (1991)). In the world that physics posits, things of different kinds—things of all kinds, if the next claim is sound—share certain (subatomic) levels of composition, and microphysics is identified as the discipline that covers everything in that realm, at whatever level of composition. Physics, then, is potentially more encompassing than the other sciences. It does not confine itself to particular ranges of empirical reality, unlike psychology or biology, nor does it confine itself, in the manner of chemistry, to the study of empirical reality above a certain (atomic) level of complexity.

* This section reproduces material in Pettit (1993*b*). See also Pettit (forthcoming).

In endorsing this first claim, the physicalist can remain relatively uncommitted on a variety of troublesome issues.

1. He may be more or less sanguine about the accuracy of actual physics, or even about the propriety of its methods; he may be more or less optimistic about how far actual physics is on the right track. He has to suppose that microphysics is directed at a real target—a realm of smaller and simpler entities—for the discipline is identified by this orientation (rather than by its claims or methods). But he need not think that its sights are well set or that it has achieved any particular degree of accuracy in hitting that target; he can even admit that microphysics may be forced to countenance entities that by present intuitions are not of an intuitively 'physical' character.
2. He can remain neutral on the proper way to formulate the claims of physics, in particular microphysics: he may think of it as the study of small particles, for example, or he may endorse a field-style formulation, under which small particles correspond to property-instances with thin world-lines.
3. Since he understands microphysics as the physics of the subatomic (or whatever) realm, he can leave open the matter of how small or how simple microphysical entities become at lower levels of composition within that realm; he can leave open, as some physicists wish to do, even the question as to whether there is any bottom level of smallest or simplest grain.
4. Consistently with defending a certain 'non-emergentist' view on the relationship between the microphysical realm and other levels—such a non-emergentism is defined in the further claims below—he may think of the microphysical realm in a non-atomist way. He may believe that certain relational microphysical properties—apart from spatiotemporal properties—are in some way fundamental.*
5. Consistently with defending such a non-emergentist view of the relationship between the microphysical realm and other levels, he may even concede that that view need not hold of the relationship between

* This microphysical non-atomism is inconsistent with the 'Humean' picture that attracts Lewis (1986, pp. ix–x): 'We have geometry: a system of external relations of spatiotemporal distance between points. Maybe points of spacetime itself, maybe point-sized bits of matter or aether or fields, maybe both. And at those points we have local qualities: perfectly natural intrinsic properties which need nothing bigger than a point at which to be instantiated. For short: we have an arrangement of qualities. And that is all'. Lewis articulates the view that all that there is apart from this Humean configuration supervenes on that configuration; this amounts to a microphysicalism akin to the sort defined here. The present point, roughly, is that the microphysicalist can keep the supervenience aspect of Lewis's picture—can maintain his microphysicalism—while dropping the Humean one.

different levels, assuming that there are some, within the microphysical realm.

Claim 2. Microphysical entities constitute everything

A. Everything in the empirical world is composed in some way—composed without remainder—out of (subatomic) entities of the kind that microphysics posits, or is itself uncomposed and microphysical.

B. The composition involved is conservative or non-creative in this sense: absent the introduction of a new source of higher-level laws or forces, two microphysically composed entities cannot differ intrinsically without some difference of a microphysical kind—without some difference in the character or configuration of their microphysical components.*

Part A of the claim is left indeterminate to the extent that no specification is offered of the sort of composition required. The mode of composition may involve any of a variety of relationships, and any of a number of mixes among those varieties: for example it may involve the relationships of identity, member to set, part to whole, token to type, realizer to role, and so on. Part B offers the one constraint that the composition must satisfy. It must be the case, under the proviso about novel laws and forces, that if there is an intrinsic difference between two things, or between the same thing at different times, then there is a microphysical difference between them. There is no macrophysical difference without a microphysical difference. As it is often put, the macrophysical supervenes on the microphysical, or at least it does so subject to the proviso given.

The supervenience claim leaves a number of matters indeterminate, but that is not relevant for our purposes. The point is that the physicalist must believe in a conservative sort of microphysical composition that makes any failure of supervenience problematic—a conservative sort of composition that will force us to explain any failure of supervenience by reference to an independent source of novel laws or forces. We shall be returning to the possibility that composition is attended by the appearance of such novel laws or forces in discussing the second pair of claims associated with physicalism.

The main object of this chapter is to define physicalism, not to defend it. But it may be worthwhile asking who is likely to reject this second claim: who, that is, apart from those who interpret physics in such a non-realist way that they cannot endorse any variant of the first claim. One opponent will be the dualist who thinks that there are things in the empirical world, say Cartesian minds, that are not in any way composed out of

* My thanks to Peter Smith for a helpful exchange on the content of this clause.

microphysical entities. Another will be the person, perhaps difficult to imagine, who accepts microphysical composition but thinks that the composition involved is not necessarily conservative; it allows, without further need of explanation, that two entities that are composed in the same way, and of the same materials, may yet different intrinsically from one another.

Claim 3. There are microphysical regularities

A. Microphysical entities are subject to certain law-like regularities by virtue of their microphysical properties and relations.

B. The laws at work in the microphysical realm do not obtain because they are required to obtain by the obtaining of certain laws at a macrolevel—perhaps the same laws (e.g. the same conservation laws), perhaps different ones. The microphysical laws, as we may say, are primitive.

Part A is not as strong as it may at first seem. It is silent on exactly how laws should be understood and on whether they may be probabilistic as well as deterministic. It allows that certain microphysical laws may ultimately apply by virtue of certain relations among microphysical entities, not by virtue of their atomistic properties. It states not that the behaviour of microphysical entities is governed entirely by microphysical laws, but only that there are some microphysical laws that play a governing part. Finally, part A does not say that the microphysical laws all bear on entities of the same grain; the laws involved may include laws that apply, as it were, at different levels of grain, provided that those levels are all subatomic. Part B gives some edge to the claim. The laws envisaged in the microphysical realm obtain for independent reasons, whatever those may be, not because they are necessary for the operation of macrolevel laws. There is no top-down push at the origin of the laws; they are as primitive as laws can be.

Who might reject this third claim? The rejection of part A involves the denial that there are laws in operation in the microphysical realm. Someone may deny that there are microphysical laws on the ground that there are no laws whatsoever; this would undermine physicalism in the precise sense defined here, but it would be consistent with a variant formulation that refers to forces or whatever in place of laws. Alternatively, someone may deny that there are microphysical laws while asserting that there are laws of a macrolevel kind; this will constitute a much more challenging attack on physicalism in my sense. The rejection of part B would go to matters of deeper metaphysics. Someone who thinks that there are microphysical laws but who holds that they are not primitive is going to have to tell a story whereby macrolevel regularities propagate their requirements downwards and establish an order of microphysical law. This is not a common picture but, presumably, is a possible one.

Claim 4. Microphysical regularities govern everything

If there are macrolevel laws, as there surely are, then the following holds.

A. They do not complement microlevel laws, taking up some degree of slack left by those laws.

B. They are not independent of microlevel laws; they do not have the potential to conflict with them and they do not serve to reinforce them, representing an extra booster for sequences of events that are established in accordance with those laws.

This fourth claim is relevant in the debate between physicalists and anti-physicalists, because if there are any macrolaws, then those laws must ultimately put constraints on the behaviour of microphysical entities.* They govern things that are composed out of microphysical entities and, if they are to dictate what happens to such things, they must impose constraints on what happens to the microphysical components of those things. However, if macrolaws constrain the behaviour of microphysical entities, then there is a question about how they relate to the microphysical laws which, by claim 3, are relevant to the behaviour of those entities. That is the question addressed in claim 4.

One sort of physicalist may deny that there are any macrolevel laws—that there are any laws, for example, that apply by virtue of the satisfaction of certain chemical or biological, psychological, or social conditions. This line is not going to be attractive to many, at least not under a generous conception of laws. The more usual sort of physicalist will admit that there are macrolevel laws, as there are laws in the microphysical realm, but will argue, by defending claim 4, that those laws do not represent a regime that is independent of the microphysical order. The macrolaws do not complement microphysical laws, filling in gaps which those laws leave, and the macrolaws do not obtain on an independent basis. The satisfaction of a macro-antecedent cannot require a result that conflicts with microphysical laws. If the satisfaction of a macro-antecedent requires a result that is already guaranteed by microphysical laws, it does not represent an extra determinant of that result. Macrolevel regularities, however objective and however worthy of notice, are fixed in place by the regime which the microphysical laws establish.†

* Crane and Mellor (1989, p. 190) illustrate the point nicely with reference to Boyle's (macrolevel) law. Suppose that the volume of a gas sample is suddenly halved. 'If the gas is ideal', they say, 'Boyle's law entails that when its pressure settles down again it will be twice what it was. That law does not dictate all the interim behaviour of the sample's molecules—except that it must be such as will eventually double the sample's pressure'.

† The point applies to the case where the same law, for example the same conservation law, is involved, now at a microlevel, now at a macrolevel. The point then is that that law will not have to do double duty.

This physicalist denies that the laws that appear at the macrolevel are of the novel kind, or are associated with the novel forces mentioned in the proviso governing the second claim. He holds that in worlds like ours that proviso is fulfilled, as a contingent matter (Lewis 1983), and so he defends a non-provisional but contingent supervenience of the macrophysical on the microphysical.* His overall position can be expressed as follows: once the microphysical conditions and the microphysical laws have been fixed, then all the crucial features of a world like ours will have been fixed, namely all the other laws that obtain at the world, all the conditions—all the initial conditions—that engage those laws, and all the things that happen in accordance with the laws. The position on laws is put in place by claim 4, where the position on initial conditions is entailed by claim 2. The position on laws amounts to a sort of nomological fundamentalism. Not only is the empirical world microphysically constituted, the empirical world is also microphysically governed.

Why should we be disposed to take the physicalism defined by our four claims as a backgound for discussion? There are three points that should be made. One is the familiar consideration that, even if it offends against democratic instincts, the sort of doctrine in question displays an attractive economy and simplicity (Pettit 1993*a*, Chapter 3). For example the doctrine explains why laws at different levels work so smoothly and systematically in tandem; it does not have to appeal to any happy coincidence of effect or any pre-established harmony (Pettit forthcoming). The second and third considerations are not so commonly mentioned.

The second is that the physicalism defined here is a concrete but fairly cautious version of an abstract and plausible claim (cf. Lewis, forthcoming). The abstract claim is that the various kinds of things in the world are composed of less varied items, that this composition establishes a hierarchy of different levels of thing, and that there is some less-than-highest level of composition such that if we fix how things are governed from there down, then we shall have fixed how things at every level are governed. This claim is concretized in the physicalism defined here. The composing items are said to be the kinds of thing that microphysics aspires to identify and the level at which we go to microphysics—the subatomic level—is said to be one such that if we fix how things are governed from

* Crane and Mellor (1989, p. 205) offer an argument against supervenience, with reference to the psychological and the physical, but one that leaves me unmoved. They admit that 'token thoughts and sensations are only supposed to supervene on simultaneous tokens of non-mental properties' but maintain, none the less, that this supervenience conflicts with the possibility of non-mental indeterministic causation of mental events, in particular, the causation by the same non-mental antecedents of these mental events here and of those, different, mental events there. I see no difficulty. Past non-mental causes can give rise indeterministically to different mental events, consistently with supervenience, provided that they do so, as all physicalists will surely say, through giving rise indeterministically to different non-mental subveners.

there down, then we will have fixed how everything is governed. Both these claims are cautious, abstracting as they do from microphysical detail, and to that extent they should not be found excessively controversial (cf. Smart 1978; Smith 1992, pp. 25–6).

The third point is of a different and perhaps more compelling character. Many philosophical projects attempt to vindicate commonplace discourses about colour or value or mentality or, as in the present case, the causality countenanced in common sense and in the special sciences. It is good practice with such projects to try to vindicate the target discourse under hard rather than easy assumptions—under assumptions that would make vindication more rather than less difficult. The assumptions associated with physicalism, as it is defined here, satisfy that desideratum. They represent a worst-case, or at least quite a bad-case, scenario from the point of view of relevant discourse-saving projects. It is going to be extremely difficult to argue for the reality of higher-level causality if this sort of physicalism holds and, equally, it is going to be really difficult to deal with a large range of philosophical questions, for example with questions bearing on the standing of secondary and evaluative properties, on the reality of mind, and on the validity of modal claims. Of course, the fact that the physicalist picture is a worst-case scenario does not give us reason to believe it, but it may give us reason to carry on as if it were sound; it may give us reason to treat physicalism as a standard working hypothesis.

PHYSICALISM AND THE ARCHITECTURE OF CONTROL

Let us say that a property or a group of properties represents an instrumental control for some type of event or condition just in case it would make sense, in principle, to realize that property or group of properties with a view to realizing the result-type in question. Let us also say that a given instance of a property was a control for some instance of a result-type just in case it would have made sense, in principle, to have put it in place in order to promote that result. The control represents a way in which the associated event or condition can be promoted or could have been promoted. At one limit, it may necessitate the occurrence of the result; at another, it may facilitate the occurrence in some minimal measure. (See Menzies and Price (1993) for the development of such ideas.)

The notion of an instrumental control is weaker than the notion of a causal control. Any cause will represent a control for the sort of event that it generates, but factors that are not causally operative may also represent instrumental controls. For example a standby cause that does not actually do any causal work in generating the relevant effect represents a control for that result: it is the sort of thing that one may put in place with a view, as we say, to promoting the likelihood of the result. A non-causal factor that could not possibly do any causal work may also represent a control for a certain event or condition: in the old example, Socrates' dying is a control for Xanthippe's becoming a widow, although it does not cause her to become a widow. The notion of a control is introduced before coming to causes proper, as it may enable us to gain a perspective on the issue that causes raise within a physicalist picture. In this section we consider how far physicalism allows for higher-level controls, and in the next section we go on to see how far it allows for higher-level causes.

The notion of a control connects intimately with the notion of something that makes a certain event or condition more probable or, at the least, that is fitted by its intrinsic character to make it more probable; it need not make it more probable in the actual world, where another factor may already have raised the probability beyond improvement (in the limit, raised it to unity), but it is, roughly, of a sort that raises the probability of such a result in various other surroundings, whether in the actual world or in worlds like the actual world (cf. Lewis 1986, pp. 206 ff). The control of any result is, as we may put it, a *pro tanto* probabilifier of that result. However, if we describe controls as *pro tanto* probabilifiers we should be careful to note that they are not probabilifiers of just any old kind; they are probabilifiers that represent potential means for promoting the associated type of result. Something may probabilify a result, say because it represents evidence that the result has ensued, without being a potential means for promoting the result (Menzies and Price 1993).

It is assumed that the definition of a control means that generally, when we have a law that relates the satisfaction of some antecedent to the satisfaction of a consequent, the antecedent-property is a control for the consequent.* Under the physicalist picture described in the previous section there are laws that govern the world at different levels. There are laws that relate microphysical conditions to microphysical results, and there are laws that relate conditions involving macrophysical configurations—chemical or biological, psychological or social-scientific—to macrophysical results. This implies that there are controls—controlling properties—to be found at different levels in the empirical world: there are micro-controls and, at different levels of composition, there are a variety of macro-controls.

However, the physicalist picture does not allow us to think of these controls as governing different areas. It does not allow us, for example, to think that the microcontrols rule in the microdomain and that, quite independently, the different macrocontrols rule in the different domains

* I am grateful to Frank Döring for drawing my attention to an incautious formulation of this point.

to which they are relevant; it does not even allow us to think that macro-controls complement their microphysical counterparts, taking up some slack that they leave. Every consequent that is governed by any macro-control involves a configuration of microentities which subveniently fixes the intrinsic character of the consequent. Every such consequent is controlled in the maximum measure available by the microproperties that relate to the configuration in question; it is controlled by the micro-antecedents whose satisfaction leads, by microphysical laws, to the appearance of the configuration.

The fact that everything falls under microphysical control may seem to raise a problem. Under the physicalist picture the microphysical laws, and therefore the microphysical controls, are primitive. Again, under the physicalist picture macrophysical laws, and therefore macrophysical controls, are not independent of their microphysical counterparts; they do not represent potentially conflicting or actually reinforcing influences. However, if everything that is controlled for by macroproperties is also controlled for by microproperties, and if the control of the microproperties is more primitively sourced, does not this mean that there really are not any macrophysical controls at all? Does it not mean that the microphysical properties have a monopoly in the control of what happens in the empirical world?

Not at all. To control a type of result is just to be a factor that might in principle be manipulated with a view to promoting it or to be a factor that might have been manipulated with a view to promoting it in some instance. There is no problem with the idea that one and the same instance of a result-type may be subject simultaneously to many different controls. That is just to say that a number of different factors might in principle have been manipulated with a view to promoting the result-type in that instance. There is no more difficulty in this idea than there is in the notion that a number of different factors may each have served to raise the probability of a certain event's occurring. In principle, then, there is nothing incoherent in the thought that microproperties and macroproperties may serve simultaneously as controls of a given event. But can we go further than defending the coherence in principle of the idea? Can we show, more concretely, how such multiple control can obtain in a world that satisfies the physicalist picture?

In explaining how factors at different levels can be causally relevant to one and the same event, Jackson and Pettit (1988, 1990) have introduced what is described as the program model of multiple relevance; we return to this in the next section. A simple analogue of that model serves to provide an account of how micro- and macroproperties can simultaneously enjoy a controlling role in relation to a given type of event. It suggests an architecture under which such simultaneous control is unproblematic.

The physicalist has no problem with the idea that certain microproperties control for a given event-type E. Suppose then that such an event comes about and comes about, as we may say, under the control of a micro-property or group of properties. How can a macroproperty H (for higher-level) serve at the same time in the role of control for E? The program model suguests that it will do so only in the case when the following three conditions are realized.

1. The realization of H involves the realization of certain microlevel properties—perhaps these, perhaps those.
2. However H is realized at the microlevel more or less, the realizing properties serve as controls for an E-type event.
3. H is realized and E occurs.

The best way to make the projected architecture vivid may be by the use of an analogy. Suppose that we are dealing with a toy world in which certain small balls correspond to microphysical entities and in which macrophysical properties are represented by the shapes which various sets of these balls are capable of assuming on a flat surface. Take the shape-property of being triangular. There will be many different ways in which a set of balls may assume a triangular shape: different numbers of balls may be involved, different sets of balls may be involved, and the same set of balls may be distributed over different places. Suppose now that no matter how a triangular shape is realized, the balls which realize it remain in position for a characteristic period; whatever the mechanics involved, the ball-level laws which apply in each of these cases ensure a certain stability. The toy world envisaged presents itself like a randomly shifting kaleidoscope except in the event that certain of the balls assume a triangular shape; in that event, we find that the balls come to display a certain temporary stability.

Consider now what we should say about the factors that control for the appearance of a period of stability. We may certainly say that that phenomenon is controlled for by the simultaneous satisfaction at a certain time of the relevant ball-level laws—by the satisfaction on the part of the balls involved of certain ball-level properties, intrinsic or relational. But we can also say in such a case that the phenomenon is controlled for by the satisfaction of the triangular property on the part of the aggregate of balls. We can do so, because the program architecture applies. The realization of the triangular property involves the realization of certain ball-level properties. No matter how the ball-level properties realize the triangular property, the realizing configuration of properties controls for a certain temporary stability, in the manner appropriate to a potential means; in fact, by the story we have told, it makes a certain temporary stability certain.

The lesson of the analogy should be clear. Even if we go along with the physicalist picture of the last section, we can think of the history of the empirical world as unfolding under the simultaneous control of micro-physical and macrophysical properties. The microphysical properties may represent more basic controls—they are presupposed by their higher-level counterparts and they are more primitively sourced—but they do not deprive macrophysical properties of the capacity to serve a similar controlling role.

There is an objection. Is the scenario of multiple-level control consistent with the supervenience which the physicalist picture asserts? This supervenience means that if we have put all microphysical controls into place, then we will also have fixed all macrophysical controls in position. At first that may seem to suggest that the only controls that really exist are the microphysical ones. But the suggestion does not survive reflection. Supervenience is a relationship in the same family as entailment; it directs us to what must also be realized, give or take certain constraints, in the event of realizing a certain condition. It is clear that if one state of affairs obtains and if it entails another, then the other also obtains. It should be clear, in parallel, that if certain microphysical controls exist and if they superveniently underpin certain other macrophysical controls, then the macrophysical controls exist also. Supervenience of the kind envisaged is not an eliminativist relationship; it serves to conserve that which it supports, not to subvert it.

PHYSICALISM AND THE ARCHITECTURE OF CAUSALITY

The fact that the physicalist world allows an architecture of lower- and higher-level controls should not really be surprising. The notion of a control carries no explicit ontological burden; it does not explicitly require the presence of any particular sort of substantive relationship. There is nothing astounding in the claim that, though the world is microphysically constituted and governed, we can still identify macrophysical conditions that serve as potential means for realizing various sorts of results.

Controls are ontologically light to the extent that they do not necessarily count as causes. A control is required only to make something instrumentally more probable, in the *pro tanto* sense, than it would otherwise have been. That condition can be satisfied, as we mentioned, by factors that have no causal effect. Consider the standby cause of something: the factor that does not itself cause the event but that is there to play a causal role if the actual cause fails. Consider, for example, the cancer which ensures that someone who has been killed in a car accident would have died before the year's end. Such a standby cause can count, by our definition, as a control for the event, since it represents a means of promoting the person's death before the end of the year; yet, by hypothesis, the control is not itself causally active in generating the event—it represents an inert control, as we might put it.*

Again, consider the sort of factor that instrumentally facilitates or necessitates something but which does so in what our intutions represent as an essentially non-causal way. This factor is not contingently inert, in the manner of a standby cause, it is necessarily of this inert character. The presence of antibodies in John's bloodstream ensures that he is immune. Mary's attendance ensures that there will be someone at the lecture. John's death ensures that Jane has become a widow. In each of these cases we find a controlling property but we do not find anything that looks intuitively like a causal connection. The controlling property constitutes, or helps to constitute, the type of result in question; it does not lead to that result via a causal process.

The contrast between controls and causes leads us to ask whether the physicalist world of the first section is compatible, not just with an architecture involving lower- and higher-level controls, but also with an architecture that allows lower- and higher-level causal, i.e. causally active, controls. Can we believe that the world is microphysically constituted and microphysically governed and still hold that there are macrophysical as well as micro-physical causes? In particular, can we hold that causally active controls are to be found equally at microphysical and macrophysical levels?

If a controlling property is to be causally relevant to the event that it controls, then what extra condition must it satisfy? The property must not be a standby factor, in the manner of the cancer. Equally, it must not be the sort of non-causal constitutive factor that appears in our other examples. But what sort of condition must it satisfy on the positive side if it is to count as causally relevant to the result for which it controls? The controlling property, or rather the relevant instance of the controlling property, must relate to the event in question in the manner illustrated when I push this table and it moves, when the wind bends this branch and it breaks, or when the fire touches this paper and it burns. But what manner of connection do these paradigms illustrate?

Philosophers divide, notoriously, on this question. We can distinguish three broadly different approaches. Some philosophers think that it is possible to give an account of the connection in more or less formal mode—in terms of laws or conditions or counterfactuals or the like. Others go to the opposite material-mode extreme and hold that the connection

* Notice in this connection that, though the criminal law is introduced as a control on deviancy, it is commonly recognized that most individuals are going to conform for reasons that have nothing to do with the law.

should be taken as a primitive irreducible exercise of power on the part of certain property-instances—as the display of a productive 'oomph' or 'bif', whether as a singular or a systematic matter, on the part of those properties.* Others again assume an intermediate position, maintaining that the causal connection should be identified, not by immediate reference to the material mode, but by reference just to certain paradigms—by reference to that relationship which those paradigms are taken to exemplify. This third approach may take the relationship in question to be of a more or less immediately salient kind. Or it may allow that the relationship is to be illuminated in science, in the way in which science illuminates the kind—H_2O—that samples of water exemplify; thus it may allow the possibility that the relationship in question boils down to the sort of energy transfer that is sometimes taken to be involved in paradigm causal connections.

However the causal connection is to be understood, the question before us is whether we can countenance a scenario under which causal controls are to be found at higher as well as lower levels. Causal controls are motor controls, as we may put it—motor controls as distinct from the inert controls represented by standby causes and by essentially non-causal factors. They are properties which control for the types of event to which they are relevant and they are properties whose instances relate to the events on which they bear in the productive manner in which my pushing relates to the movement of the table.†

The problem that we face may be usefully presented with the help of a tree diagram. We have seen that controls are a species of *pro tanto* probabilifiers, and that causal or motor controls are a species of controls. The question before us is whether motor controls can be found at higher levels as well as within the microphysical domain.

* If a certain property-type displays the required efficacy as a singular matter—if singularism holds—then that means that it is a primitive fact about each instance of the property—a fact not replicated necessarily in otherwise indiscernible worlds—that the instance is efficacious. It will not be because it is an instance of that property—of that property rather than of any of the apparently irrelevant properties that it co-instantiates—that it is efficacious; its being efficacious is an independent fact.

† In focusing the discussion on motor controls, I seek to bypass a question that interests some philosophers. This, in our jargon, is the question of whether the word 'cause' refers to any property-instance that plays a motor role or whether it only properly refers to a property-instance that satisfies a double condition: the instance plays a motor role and, moreoever, the property plays a controlling role. Suppose that two properties, F and G, are co-instantiated and that their co-instance is causally relevant to some result, E. Suppose further that the F-property controls for that result and the G-property does not. The G-property may be wholly irrelevant to the type of result in question; it may even be the property of being reported in yesterday's newspaper. Is it reasonable to describe the G as the cause of E? Or should we reserve that description only for the F? I avoid this question, while making it clear that I am only interested in properties like the F-property—in properties thai are at once controls and, in their instances, motors.

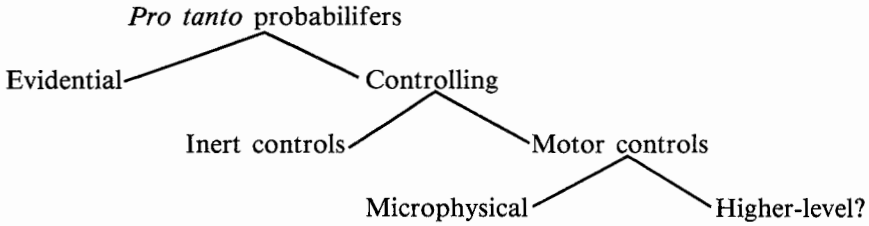

How do we approach our problem? We know from the previous section that there is no difficulty in countenancing higher-level properties as controls, even in a physicalist world. We shall identify conditions under which such higher-level controls will certainly count as 'programmers' of events that are microphysically produced; this is done by reference to the program model. Then we shall explore two questions: Does the enjoyment of such a programming status mean that the controls involved are more than inert controls? In particular, does it mean that they are motor controls in the same sense as microphysical antecedents?

We assume that the physicalist, whatever his conception of the causal connection, will have no hesitation in taking some microphysical properties as motor controls. The program model mentioned earlier identifies conditions under which it is appropriate to say that a higher-level property 'programs' for a result that issues at the same time from a microphysical motor control or, more generally, from any lower-level motor control. A higher-level property H will program for a microphysically produced result E just in case the following three conditions are fulfilled (the last two conditions are stronger than those involved in the architecture of control, because of the difference made by the motor requirement).

1. The realization of H involves the realization of certain microlevel properties, perhaps these, perhaps those.
2. However H is realized at the microlevel—more or less—the realizing properties are apt to be motor controls responsible for an E-type event.
3. The microlevel properties that actually realize H are motor controls that are responsible for E in the manner envisaged in condition 2.*

These conditions can be illustrated by reference back to the example from the toy world. Suppose that the ball-level properties that engage with the relevant laws count in any instance as motor controls of the stability that follows on their realization. The triangular shape must be said to

* The phrase 'in the manner envisaged in condition 2' is necessary to avoid certain deviant chain cases where roughly speaking, the realizing properties are responsible for the right sort of event but in the wrong sort of manner. I am grateful for a discussion on this point with Daniel Andler, Jean-Pierre Dupuy, Pierre Jacob, and Francois Recanati.

program for the appearance of that stability on any occasion where it materializes in the appropriate way. It meets the required conditions. First, the realization of the shape always involves the realization of certain ball-level properties. Second, however the shape is realized, the realizing configuration of ball-level properties tends to produce an ensuring period of stability in the motor-control way. Third, the actual configuration that realizes the shape is a motor control of the actual stability that ensues.

The sort of relationship envisaged between macrophysical programmers and microphysical motor controls is a relationship that can obtain between properties at any two levels of composition.* Assume as an illustration that the molecular structure of this eraser is a motor control relevant to the fact that it bends under a certain pressure. The malleability of the eraser programs for that bending because the instantiation of the malleability involves the instantiation of certain molecular-structural properties. The sorts of properties associated with instantiations of malleability are apt to motor-control the sort of bending effect in question. The molecular-structural property associated with the actual instantiation of malleability is a motor control of the bending.†

Other examples of the program model become salient as we recognize suitably corresponding relations across levels in different cases. The squareness of the peg stops it going into the round hole. The redness of the rag enrages the bull. The rise in unemployment raises the level of crime. The rational appeal of the option leads the agent to select it. In every case the instantiation of the higher-level property involves the instantiation of motor-control properties at the lower level. The squareness of the peg involves the sort of molecular contact which blocks the peg going through the round hole, the redness of the rag involves the sort of physical stimulation which provokes the bull, the rise in unemployment involves a shift in motives and opportunities that is likely to increase aggregate crime, the rationality or perceived rationality of an action involves the sort of neural disposition which produces suitable behaviour, and so on across a great variety of possible cases. The harmony of levels holds for different reasons

* If a microphysical property produces something in the manner of a motor control, then, whatever relevance a macrophysical property has to the result, it meets two conditions. First it presupposes the microphysical influence. Second, it is not the relevance of an earlier cause of the microphysical factor or of a cause with which the microphysical factor needs to join forces; the macrophysical property is causally insulated from the microphysical. I assume here that such conditions will be fulfilled by any families of properties that belong, as I put it, at higher and lower levels of composition. Elsewhere I have referred to such conditions in defining what it is for two families of properties to belong to respectively higher and lower levels (Pettit 1992, 1993*a*, Chapter 3).

† The example is rather trivial because knowledge that the molecular structure combined in the ordinary way with a certain outside pressure to cause the bending yields knowledge that the malleability was also causally relevant. However, not all illustrations of the program model have this feature.

in the different cases, but the fact that it obtains shows how the program model may apply in any of the examples.

Suppose that the program model is satisfied by certain macrophysical controls relative to the microphysical motor controls that are responsible for a given type of event. Does that mean that the macrophysical controls are more than controls in the instrumental sense introduced in the last section? In particular, does it mean that they amount to motor controls of the event in question—to controls of the same causal status, if not of the same hierarchical standing, as the microphysical factors?

First question

The answer to the first question is certainly positive. There is more involved in a higher-level property's programming for an event than there is in its just controlling for it. Take our earlier example of a standby cause. Suppose that someone has cancer and that this higher-level property means that there will be some specific organic failure or vulnerability—perhaps this, perhaps that—which will lead to death, in the motor-control way, within the year. Suppose that the person does indeed die within the year but dies as a result of injury in a car accident. The cancer controlled for that event, but it does not count as a programmer of the event. The reason is that death does not come about at the level of specific organic breakdown in the manner associated with cancer; the breakdown that leads to death would have been just as likely, or so we assume, in the absence of the cancer.

This observation already teaches us an important lesson about higher-level programmers. Whatever views we come to hold about their motor-control status, we must at least recognize that there is more involved in a higher-level property's programming for something than there is in its controlling for it. Even if the programming property is not a motor control, it is a property which is relevant only in the case where the lower-level motor control is of a certain kind. The cancer would count as a programmer of death, not just as a controller, only if death came about in a certain way at the level of organ failure. Whether or not a programming property is to count as a motor control—we come to that question in a moment—it requires the presence of a certain sort of motor control; it puts constraints on the sort of motor control that is in operation. Thus the programming property is not just a control of an inert kind.

Second question

Can higher-level programming properties generally count as motor controls in the full sense? With this question, we come to the heart of the matter that concerns us. It is granted for sure that microphysical properties have

a certain hierarchical priority in the motor control of how the world goes. But the question is whether the factors that count as microphysically productive are motor controls in a sense that allows us to think of investigating similar motor controls—similar productive properties—in the special sciences and in common-sense inquiries. Motor controls are full-blooded causes, so we assume, and the issue is whether such investigations can generally aspire to identify full-blooded causes in the same sense.

My answer to the question is that it all depends. It depends, in particular, on how exactly we choose to conceptualize the motor connection that distinguishes motor controls from inert controls: from standby causes, for example, and from controlling factors of a non-causal kind. As mentioned earlier, while everyone agrees that my pushing causally moves this table, and while the wind's blowing causally bends the branch, the connection involved is variously understood. Whether programmers are to be taken as motor controls depends on which construal is adopted.

The first approach to the question about the motor connection assumes that the connection lends itself to a more or less formal-mode analysis, say by reference to laws or counterfactuals. Under this approach there can be no difficulty, in principle, with the idea that higher-level programming properties count as motor controls. For, as it is possible for such properties to do as well as microphysical properties in fulfilling the control conditions, so it is going to be possible for them to fulfill most of the conditions that are likely to be invoked in explicating the motor connection.

The approach that stands at the opposite extreme from this first approach would say that the motor connection is irreducible—that it involves the exercise of a notion of power that we have no independent means of elucidating; it is a connection in which the instantiation of the relevant property contributes an 'oomph' or a 'bif' that gives rise to the event produced. As there is no difficulty in recognizing programming properties as motor controls under the purely formal-mode analysis of the motor connection, so there is no possibility of doing so under this extreme version of the material-mode approach.

A little reflection makes clear why. It is granted under the physicalist picture that everything is governed in a primitive way by microphysical laws. This must mean that the motor controls to which those laws direct us can and do enjoy the irreducible 'oomph' that is characteristic of the motor connection. It is granted equally under the physicalist picture that higher-level laws are not on the same primitive footing as microphysical. This must mean that they do not direct us to rival motor controls. If the higher-level laws did direct us to rival motor controls, then there would be a potential for conflict between the different powers involved and, in cases of harmony, the higher-level powers would have to be seen as reinforcing the microphysical powers.

As an aside, it should be noticed that the fact that higher-level laws do not direct us to independent motor controls need not mean that every particular macrolevel control that we invoke is bound to be something less than a motor control. It may happen in this description that a macrolevel control is just the same particular—the same property-instance—as the corresponding microlevel control. This will mean that it has a dual aspect. As an abstract property it will count as a higher-level control, on a par with other macrophysical factors. As a concrete property-instance it will count as a motor control, on a par with microphysical factors (Macdonald and Macdonald 1986).*

What of the remaining sort of approach to the motor connection—the intermediate approach that directs us to the material mode but in a way that defers to certain paradigms, and perhaps to what science is capable of revealing in those paradigms? This approach may say that the motor connection is that sort of relationship—one that has a natural salience for us and one whose features should be available to reflection—involved in certain paradigms. Or it may say that it is that sort of relationship which science is likely to uncover or define—say, a relationship involving a transfer of energy or, at a certain level, a locality of contact—in those paradigms. The approach assumes that the paradigm cases point us to a definite kind and the connection is picked out rigidly as that relationship which unifies those cases in the actual world.

This intermediate approach points in the same way as the first on the question before us. For all that the approach says on the motor connection, it is possible for higher-level programmers in general to satisfy the connection. It is possible for a higher-level programmer to control for a given event, for example, and to control for it in a manner that involves even something apparently so fitted for microphysical causes as a transfer of energy. It is possible for the redness of the rag to control for the bull's anger in such a manner or for the rise in unemployment to control in that way for an increase in the level of crime. The realization of the higher-level property leads to the event and leads to it in a way that involves a transfer of energy; the realization of the higher-level property does not merely connect with the result in the fashion of an inert control. Therefore when a macrocontrol and a microcontrol relate in the programming way, they can both count as motor controls for the event that ensues. They can each involve a transfer of energy—the same transfer of energy, no doubt—and

* Under the singularism mentioned in the footnote to p. 412, the fact that the property-instance is an instance of the higher-level property will be just as relevant—in fact, just as irrelevant—to its efficacy as the fact that it is an instance of a microphysical property or group of properties. But there will still be an asymmetry between microphysical and macrophysical properties, since many macrophysical controllers—those which cannot be co-instantiated with microphysical controls—will be systematically deprived of the status of motor controls.

they can each be hailed, at different levels in the program architecture, as things that served to bring the event about.

It has already been stated that it is possible in principle for either of the first two approaches to the motor connection to allow higher-level programmers as motor controls. Not every version of those approaches will do so, of course, but this is not the place to debate specific proposals; I am content to defend the more abstract point. In conclusion, however, I would like to buttress the abstract point by mentioning some reasons why we should avoid the third approach and go along with one of the first two. For the record, my own preference is for the intermediate approach, but I shall not defend that preference in the present context.*

A first point to notice about the third, irreducibilist, approach is that it is hostage to empirical theory in a way in which the first two are not. Some physicists speculate that there may be no bottom microphysical level, as I mentioned earlier in passing, and it is interesting that this would plausibly mean, on the irreducibilist approach, that there are not really any motor connections to be found in the world. It would seem to be an arbitrary choice on the part of the irreducibilist to pick out one level in an infinitely extended progression of levels downward and to say that there, at that level, is where we find the 'oompz' that is allegedly characteristic of motor connections. The absence of a bottom level would not make any problem, of course, for the first, formal-mode, approach and neither need it cause a difficulty with the intermediate approach; the feature that is characteristic of the motor connection can be quite compatible, for all we know, with an infinite progression downwards.

A second point against the irreducibilist approach is that it would force us to give substantive significance to choices of representation that appear to be more or less arbitrary. Do we say that the malleability of the eraser is just the molecular structure—and ultimately, we presume, the microphysical configuration—that leads to bending under suitable pressure? Or do we say that it is the higher-order state of having such a molecular (and microphysical) structure that leads to bending under that pressure? Intuitively this is a free and arbitrary choice of presentation—a matter of book-keeping, not of discovery. (Jackson and Pettit 1988). However, under the irreducibilist approach it is a matter of the deepest ontological significance. Under the first approach the malleability can count as a motor control, since its instance will just be the same thing as the instance of the microphysical control. Under the second approach it cannot; its instance does not coincide with the instance of the microphysical control.

* For the record, I tend to think of motor controls as controls that are analogous to the controls represented by our actions, as when I push and the table moves, or I blow and the candle goes out.

Any approach that gives such significance to ontological book-keeping ought to be suspect.

A third and related point against the irreducibilist approach is that it forces us to recognize a difference where intuitively there is none. Determinable properties are co-instantiated, of necessity, with their more determinate versions: the redness of this rag is co-instantiated with the particular shade of red, and let us suppose with the particular microphysical realizer of that shade, which disturbs the bull. With such a determinable property, then, we must regard it as a motor control of anything for which it programs. But there is no intuitive difference between the way that the redness programs for the bull's reaction and the way that other sorts of property have programming effects—the way that the squareness of the peg programs for its being prevented from going through the round hole or the way that the rise in unemployment programs for the increase in crime. On the third, irreducibility, approach there ought to be a glaring metaphysical chasm dividing these cases. However, there is no obvious gap to be found there.

The upshot of such considerations, then, is that things look good for the special sciences, and indeed for common sense. Under the most pessimistic scenario, that associated with the irreducibilist picture of the causal connection, the special sciences are generally going to identify factors that count as more than just instrumental controls—factors that count as programmers.* Under more likely scenarios the special sciences are going to count as uncovering real full-blooded causes—causes in the same motor sense as that which applies in microphysics. It is true that the causes uncovered in microphysics are going to come lower in the hierarchical programming architecture, and it is true that once the microphysical motor controls are in place, then everything else is in place also. But these considerations are quite compatible with the thought that higher-level controls are real and important and that higher-level controls are of the motor kind.

If there is any lingering resistance to this idea, then I suspect that it will come from the thought that in principle microphysics could do all that the special sciences do in the recognition of higher-level causes: it could, in principle, identify the laws and the controls associated with those sciences by identifying the different ways in which the laws or controls may be instantiated and by recognizing that each mode of instantiation leads to a certain sort of effect. But this possibility is entirely abstract—it supposes an ideal epistemology that we can never approximate—and in any case the

* Non-programming controls may also be worth investigating. My picture of rational choice theory (Pettit 1993*a*, Chapter 5) suggests that outside market-like contexts such theory involves the identification of standby causes that help to stabilize certain patterns of behaviour—that render them resilient—without actually causing people generally to produce the behaviour.

lesson that it teaches sounds a positive note. What it means is that when we involve ourselves in the special sciences then we track the laws and controls—the real and important laws and controls—that an ideal microphysics would address in a different idiom. What it means is that research in the special sciences—research in chemistry and biology, psychology and social science, even research of a common-sense type—can be seen as microphysics by other means.

ACKNOWLEDGEMENTS

This chapter was written during 2 months in Paris, where I held visiting appointments, first at the Ecole des Hautes Etudes en Sciences Sociales (CEMS centre), and then at the Ecole Polytechnique (CREA centre). I am grateful to both institutions for their support. I benefited greatly from the stimulus I received in thinking about the topic of the Chapter from exchanges with a number of colleagues there and from comments on an earlier draft by Pascal Engel, Pierre Jacob, and, especially, Frank Doring. I owe a particular debt of gratitude to John Broome, Frank Jackson, and Peter Menzies for discussions of certain points.

REFERENCES

Blackburn, S. (1991). Losing your mind. In *The future of folk psychology* (ed. J. Greenwood), pp. 196–225. Oxford University Press, New York.

Block, N. (1990). Can the mind change the world? In *Meaning and method: essays in honour of Hilary Putnam* (ed. G. Boolos), pp. 137–70. Cambridge University Press.

Crane, T. (1991). 'Way Indeed? Papineau on supervenience and the completeness of Payous. *Analysis* **50**, 32–7.

Crane, T. and Mellor, D. H. (1989). There is no question of physicalism. *Mind* **99**, 185–206.

Jackson, F. (1994). Armchair metaphysics. In *Philosophy in Mind* (ed. J. O'Leary-Hawthorne and M. Michael). Kluwer, Dordrecht.

Jackson, F. and Pettit, P. (1988). Functionalism and broad content. *Mind* **97**, 381–400.

Jackson, F. and Pettit, P. (1990). Program explanation: a general perspective. *Analysis* **50**, 107–17.

Jacob, P. (1991–2). Externalism and mental causation. Proceedings of the Aristotelian Society **92**, 203–19.

Lewis, D. (1983). New work of a theory of universals, *Australasian Journal of Philosophy* **61**, 343–77.

Lewis, D. (1986). *Philosopical papers*, Vol. 2. Oxford University Press.

Lewis, D. (forthcoming). Reduction of mind. In *A companion to philosophy of mind* (ed. S. Guttenplan). Blackwell, Oxford.

Macdonald, G. (1992). The nature of naturalism. *Proceedings of the Aristotelian Society* (Suppl.), **66**, 225–44.

Macdonald, C. and Macdonald, G. (1986). Mental causes and expl anation of action. *Philosophical Review* **36**, 145–58.

Menzies, P. and Price, H. (1993). Causality as a secondary quality. *British Journal for the Philosophy of Science* **44**, 187–203.

Papineau (1990). Why supervenience? *Analysis* **50**, 66–71.

Pettit, P. (1992). The nature of naturalism. *Proceedings of the Aristotelian Society* (Suppl.), **66**, 245–66.

Pettit, P. (1993*a*). *The common mind: an essay on psychology, society, and politics* Oxford University Press, New York.

Pettit, P. (1993*b*). A definition of physicalism. *Analysis* **53**, 213–33.

Pettit, P. (forthcoming). Microphysicalism without contingent micro–macro laws, *Analysis* **54**.

Smart, J. J. C. (1978). The content of physicalism. *Philosophical Quarterly* **28**, 339–41.

Smith, P. (1992). Modest reductions and the unity of science. In *Reduction, explanation, and realism* (ed. D. Charles and K. Lennon). Oxford University Press.

Yablo, S. (1992). Mental causation. *Philosophical Review* **101**, 245–80.