

‘The Program Model, Difference-makers, and the Exclusion Problem’, in *Making a Difference*, ed Helen Beebee, Christopher

12

The Program Model, Difference-makers, and the Exclusion Problem

Philip Pettit

In ‘Non-reductive Physicalism and the Limits of the Exclusion Principle’, Christian List and Peter Menzies (2009) make a significant contribution to our understanding of causation within a multi-level structure.¹ Their concern in particular is with the sort of structure exemplified by the different sciences: at one level, physics, and at higher levels, special sciences like biology, psychology, and social science. They aim to vindicate the viability of a non-reductive physicalism: that is, an approach that takes all that there is to be composed of physical material, and to be determined by physical laws, but still recognizes a certain autonomy in the causal factors invoked by the special sciences.²

The picture developed in their paper, which highlights the notion of difference-making, compares and contrasts in a number of ways with a picture of the multi-level causal structure illustrated by the different sciences that Frank Jackson and I dubbed ‘the program model’.³ In this short piece I explore the relation between the two pictures and their divergent implications for how far the special sciences—for example, biology, psychology, and social science—can claim to provide causal explanations that are autonomous in relation to physics. There is more to List and Menzies’ paper—and of course more to the overall metaphysics of causation with

¹ I am indebted to Helen Beebee, Frank Jackson, Graham Macdonald, Chrys Mantzavinos, and especially Christian List and Peter Menzies for helpful comments. I can only marvel at Peter’s courage in providing detailed comments despite the struggle with his final illness. In this, as in many other respects, he remains a model for us all.

² There is an issue as to how ‘physical’ is to be understood that I ignore here. For a line on that problem see Pettit 1993b, 1994, 1995. See too Jackson 1998.

³ The papers we jointly wrote on this topic are collected in parts 1 and 2 of Jackson, Pettit, and Smith 2004. I wrote independently on the program model in a number of places, particularly in Pettit 1993a, 2007a.

which they are individually and jointly associated—than a comparison with the program model can bring out. But for reasons of space and focus I ignore those other elements in their theory.

The chapter is in four sections. In the first, I introduce the notions of programming and difference-making, and relate them to one another. In the second, I look at the connections made with causation in the rival pictures associated with those concepts; I link the program model with a hierarchical picture in which lower causal levels have a certain priority, and the difference-making model with an anti-hierarchical picture in which this is not so. In the third section I turn to the different perspectives that the pictures offer on the exclusion problem that Jaegwon Kim has raised for non-reductive physicalism (see e.g. Kim 1998). And then in the final section I argue that, despite the appearance of conflict, the difference-making picture converges in important ways with the program model.

1 Programmers and Difference-makers

The concern in this chapter is with the sort of multi-level causal structure illustrated by almost all accounts in the structure of the different sciences. Within this structure, roughly described, social properties are taken to be multiply realizable at the psychological level, psychological properties at the biological or neural, biological properties at the chemical, chemical properties at the physical or micro-physical, and so on. The list is open, because there are also liable to be levels within the different sciences, especially the microphysical. Indeed on some theoretical accounts, there may be microphysical levels without end; the structure may continue downwards indefinitely.

List and Menzies explicitly or implicitly endorse three special framework assumptions governing this sort of multi-level causal structure. First, the realization assumption: properties invoked at higher levels are each taken to be capable of being realized in any of an open range of ways at lower levels. Second, the realization-invariance assumption: properties invoked in explanation at a higher level usually bear relations to other properties at that level that are invariant over, or insensitive to, how the properties are realized at lower levels. And third, the distinct-existences assumption: those higher-level properties—or, more strictly, their instances—are distinct existences, to introduce a Humean expression, from the properties or property-instances to which they bear such relations.

List and Menzies endorse the first, realization assumption insofar as they present various examples as exemplifying the realization of higher-level properties at a lower level. The property of redness is realized by one or another shade of redness, they say; the property of a magnitude of at least N milligrams realized by any magnitude at or above N ; and, assuming that physicalism holds, the property of intending something is realized by one or another neural profile in the agent. While higher-level and lower-level properties may belong to different scientific levels of realization, as in this

last psychological example, the first two examples show that this is not essential.⁴ The realization illustrated in all these cases, it should be noted, entails the supervenience of the higher-level property on how things are at the lower level. It means, in other words, that if a higher-level property is realized by a lower-level property, then there can be no change in the higher level without a change at the lower level. Thus a picture will remain red so long as it remains a certain shade of red; an object will continue to weigh at least N milligrams so long as it continues to have a particular weight at or above N milligrams; and an agent will continue to intend something so long as he or she remains in a suitable neural state.

List and Menzies endorse the second, realization-invariance assumption insofar as they show that it is only when very special conditions are satisfied that the realization-invariance fails.⁵ In this discussion we restrict attention to the standard cases of multi-level structures where they are not satisfied and realization-invariance holds. In these cases the higher-level properties invoked—say, in the special sciences—bear relations to other properties that are ‘invariant under changes to the way in which the higher-level properties are physically realized’ (499). Thus the relation between an intention and an action that plays an explanatory role in such a case will be invariant under changes in how the intention and the action are realized, whether in the same agent at different times or in different agents.

The third, distinct-existences assumption is designed to ensure that the significant relations between properties at a higher level are not the sorts of relations that obtain between, say, a disposition and the manifestation of that disposition. The disposition of solubility in water—the property of having a property that disposes a bearer under normal conditions to dissolve in water—is significantly related, by some criteria of significance, to the dissolving of the bearer in water. But it is not a distinct existence from the dissolving and not capable, for example, of bearing a causal relationship to it: that is, a contingent relationship that might not have obtained, even under presumptively suitable conditions. While List and Menzies do not explicitly endorse this assumption, the examples on which they focus systematically avoid cases where it would not apply and the spirit of interpretative charity requires us to ascribe it to them.⁶

⁴ In taking this line, they ignore some subtleties in the position of Jaegwon Kim, the target of much of their discussion. See Kim 1998: 83–7. I also ignore those subtleties in this chapter. It is worth noting that most of the cases relevant to the sciences are of a kind with the psychological example, where the realization claim is a posteriori; in the other examples the relationship is a priori knowable.

⁵ The conditions that falsify the assumption—and falsify ‘the revised exclusion principle’—are those that figure on the right-hand side of ‘the compatibility result’ that List and Menzies (2009: 491) establish. They assume that those conditions are special and, in our usage, non-standard. Let M be a higher-level property—say, an intention—and B a piece of behavior that M explains. Under these conditions, as they say, ‘the causal system is very sensitive in the sense that small perturbations in the way in which M is realized will result in the absence of B ’ (492–3). My thanks to Christian List for a useful exchange on this point.

⁶ Without the assumption, as I argued in an earlier version of the chapter, they would be committed to the utterly implausible view that dispositions can be difference-makers and causes and can therefore

In our discussion of the program model, or at least of the program model as it applies in the sciences, Frank Jackson and I began from similar assumptions and I shall take these as given in what follows.⁷ I shall also take as given the simplifying assumptions made by List and Menzies in neglecting issues of overdetermination, in avoiding the complexities raised by probabilistic as distinct from deterministic causation, and in assuming that properties—or, strictly, the instances of properties⁸—can figure as the terms of causal relations.

I now try to provide a sketch of the connection between the notions of difference-making and programming, although one that abstracts from the competing ways in which they are linked with the notion of causation; I turn to that topic in the second section. List and Menzies rely on the framework assumptions to identify difference-makers in standard multi-level structures and they then argue, as we shall see in the next section, that difference-makers have a unique title within those structures to being characterized as causes. Jackson and I rely on the assumptions to introduce the wider category of programmers and this enables us to take a more generous view in identifying causes.

List and Menzies introduce the idea of a difference-maker as follows, where ‘*F*’ and ‘*G*’ are properties:

The presence of *F* makes a difference to the presence of *G* in the actual situation just in case (i) if any relevantly similar possible situation instantiates *F*, it instantiates *G*; and (ii) if any relevantly similar possible situation instantiates not-*F*, it instantiates not-*G*. (2009: 482)

One of the most important contributions of their paper is the way in which they elaborate this account in more formal, counterfactual terms, employing a revised and independently attractive version of David Lewis’s semantics for counterfactuals. But for purposes of this chapter we can rely entirely on the more elementary account just given.

exclude the causal role of their realizers: the disposers at the lower level. I have been persuaded, especially in light of correspondence with Menzies, that this would be an unfair criticism and that they should be taken to endorse the distinct-existences assumption.

⁷ While we discussed the application of the model to cases of dispositions, where the distinct-existences assumption does not apply, we used that application only to illustrate how the programming relationship need not add to our causal understanding. Notice that the distinct-existences assumption does apply, however, with functional as distinct from dispositional states. Unlike dispositions these are not characterized by just one manifesting connection but by the fulfillment of a number of conditions—usually, typical rather than invariable conditions. When the existence of the state requires a number of conditions to obtain, then even if they invariably include the connection with the effect to be explained, invoking the state in the explanation directs us to a connection between the other conditions required for the state to obtain and that effect. This would seem to hold whether we interpret functional states as ‘role-states’ or ‘realizer-states’ (Jackson and Pettit 1988), although my preference is to take them to refer to realizer-states.

⁸ List and Menzies (2009: 475, fn. 2) hold that ‘strictly properties are not causes of other properties; it is instances of properties that cause instances of other properties’. Anticipating later discussion, I think that according to the program model, the instance of one programming property leads to the instance of another property but does so in virtue of being an instance of the appropriate property: in particular, a property such that any instance would have led to an instance of the other property.

According to that account, the instantiation or presence of the property, *F*, makes a difference that consists in the instantiation or presence of the property, *G*, just in case two conditions hold over a certain range of possible situations: first, the presence of *F* necessitates the presence of *G* and, second, the absence of *F* necessitates the absence of *G*.⁹ The range of situations over which the *F*–*G* relationship has to hold includes the actual situation but it also includes all relevantly similar situations: all situations that vary from the actual one only in irrelevant ways.

What are to count as relevant, and what as irrelevant, variations? List and Menzies agree with James Woodward (2005) that which situations count as similar, which variations as relevant, is ‘a context-sensitive matter, and disciplinary-specific rules may play an important role’ (2009: 499). But in standard cases within the special sciences, there is a natural interpretation of relevant similarity. On that interpretation a property, *F*, will count as a higher-level, special-science difference-maker for a property, *G*, just in case the *F*–*G* relationship holds in all possible situations—subject, no doubt, to certain constraints—that differ at most from the actual one in how the presence of the *F*-property or the absence of the *F*-property is realized. Thus the property of intending to move my hand thus and so will make a suitable difference to my moving my hand in that manner just insofar as its relation to the movement of the hand satisfies suitable conditions. The presence of the intention will lead to the performance of the action and the absence of the intention will lead to the failure of the performance, regardless of how the presence or absence of the intention is neurally or otherwise realized.

How does the notion of a programming property relate to that of a difference-maker, in particular a higher-level difference-maker of the kind that features in the special sciences? It is defined on the assumption that the hierarchical structure at issue is of a standard kind that satisfies the assumption of realization-invariance. A property, *F*, programs for a property, *G*, just in case the first of the two conditions relevant to difference-making holds in such a structure: the presence of *F* suitably necessitates the presence of *G* but the absence of *F* does not necessitate the absence of *G*. On this picture, the *F*-property is multiply realizable at lower levels; the actual instantiation of the *F*-property, involving a realizer property, *RF*, leads to the instantiation of the *G*-property, necessitating a realizer of that property, *RG*; and the instantiation of the *F*-property, no matter by what realizer, would also lead to the instantiation of the *G*-property. In short, there is a realization-invariant linkage between the presence of the *F*-property and the presence of the *G*-property, though there need be no such linkage between the absence of the *F*-property and the absence of the *G*-property.¹⁰

⁹ The language of necessitation is appropriate, given that probabilistic linkages are put aside for reasons of simplicity; see above.

¹⁰ For convenience of presentation I do not distinguish between the case where a higher-level programming property is co-instantiated with the lower-level realizer that leads to the programmed effect, so that

Difference-makers constitute a specific variety of programmers in standard cases.¹¹ They are super-programmers in the sense that not only does presence program for presence, it is also the case that absence programs for absence. This difference means that while many different properties may program for the instantiation of a property, *G*, only one single property can play the super-programming role in a standard case: that is, a case where realization-insensitivity holds.

Consider the psychological example again. The presence of the intention may program for my moving my hand in a certain way, regardless of how it is neurally or chemically realized. And, assuming that there is no still higher-level programming property the intention realizes, its absence will program for the absence of that movement: I would not have moved my hand if I had not had an intention to do so. And so in that sense the intention may super-program for the movement of my hand and count as a difference-maker of the movement.

But while the intention serves as the only super-programmer of the movement, the presence of the neural realizer—say, *NI*—also programs for the movement, leading to that movement invariantly over how it is chemically realized, whether by *C1* or *C2* . . . or *Cn*. It will not super-program for the movement, however, and not count as a difference-maker, since its absence would not program for the absence of the movement. By hypothesis, the presence of the intention programs for the action, regardless of whether it is realized by *NI* or by some other neural state, *N2* or . . . *Nn*. But if the presence of the intention in the absence of *NI* ensures the presence of the action, then the absence of *NI* cannot program for the absence of the action. And so *NI* cannot count as a super-programmer or difference-maker. That role will be reserved for the intention alone.

To sum up, then, given a multi-level structure of the kind illustrated in standard cases, both sides to the debate can agree that there are higher-level properties that program for other properties, in the sense that their presence necessitates the presence of those other properties in a realization-invariant manner. And both sides can agree equally that in those cases there is only one property that super-programs for the associated property and counts as a difference-maker. This will be the programmer of which it is true, not only that its presence programs for the presence of the associated property, but also that its absence programs for the absence of that property.

their instances are one and the same event, and the case where the lower-level property-instance that leads to the effect is only a part of the lower-level realizer of the programming property. For an example of the latter scenario, think of how the boiling of water in a closed flask may program for the collapse of the flask, where the lower-level event that leads to the collapse is not the realizer of the boiling—the aggregate of water molecules, moving at a certain rate—but a part of that realizer: the particular molecule, as it might be, that cracks a molecular bond and precipitates the collapse. I discuss this case in Pettit (2007a: 222–4), responding to Macdonald and Macdonald (2007).

¹¹ We might define difference-making in terms of programming, were we happy to enforce a restriction to standard cases. On this account the presence of *F* must program for the presence of *G* and the absence of *F* must program for the absence of *G*.

2 Causation in the Rival Models

I have made no mention of causation up to this point, identifying difference-makers and programmers by their functional roles in a standard multi-level structure. But the most striking difference between the approaches appears in how they relate their preferred functional relationships to the notion of causation.

List and Menzies argue for construing causation, in whatever context, as a matter of difference-making. They endorse Stephen Yablo's (1992) proportionality constraint on causation according to which 'causes make a difference to their effects' (2009: 481). In Yablo's illustration, when a pigeon is conditioned to peck on a red target, there is a proportionality between redness and pecking that does not hold, for example, between crimson and pecking: the presence of red occasions pecking, its absence occasions non-pecking; the presence of crimson occasions pecking but its absence does not occasion non-pecking—at least not if the target is some other shade of red. Yablo argues that it is natural in this case to identify the redness as the cause of the pecking and to deny crimson or any particular shade of red this status. List and Menzies accept Yablo's proportionality constraint and they tailor difference-makers to meet a generalized version of that constraint, counting as proportional causes.

In equating causation with difference-making, they suggest that they are endorsing a conception that is 'common to several different theories of causation—for example, counterfactual, interventionist and contrastive ones' (2009: 476). They maintain that the idea of difference-making motivates various theories, that it is implicit in the presumptively causal pursuit of 'manipulation and control', that it figures in causal explanation, and that it is the sort of relationship that causal experiments particularly target (489–90). Thus they have no hesitation in affirming that all difference-makers are causes: the 'notion of difference-making is clearly a causal one' (489).

Do they also think that all causes are difference-makers? It appears so. The only alternative candidates they consider would conform to the 'conception of causation as generation, or effective production' (489), which Jaegwon Kim (1998) and others endorse. But they complain about 'the opacity of the notion of productive or generative causation' (490), arguing that no one has provided a satisfactory analysis of a producer or generator. And they conclude that 'unless a better explication can be given of causation as production, this notion can hardly play a significant role in the debate about mental causation' (488)—or presumably in any debate about causation across the sciences, basic and special.

We saw in section 1 that as between candidates at different scientific levels in standard cases there can only be one super-programmer for a given effect. The main result of identifying difference-makers and causes is that this ensures among investigations at different levels that for a given effect, say the instantiation of a property like *G*, there is only one cause to be found. The cause may be the instantiation of a relatively high or a relatively low level property: it may be the instantiation of the property *F*, or of a lower-level property *RF*, which realizes *F*, or of the property that realizes that property in turn, *RRF*, and so on.

Once we have determined which is the operative, difference-making cause, we will have reached the end of our inquiry and there will be no reason—or at least no reason in causal inquiry proper—to pay attention to any other factor. Thus, once we know that an agent's intention is the difference-maker for a given action, we know all that there is to know about the cause of the action. In a certain sense our causal understanding will not be enhanced by learning about the neural nuts and bolts whereby the intention is realized and the realizer linked to the neural realizer of the action. Looking into the brain may reveal interesting detail about the way the intention works but, according to this approach, it will not unearth anything that might be thought to compete with the intention for a causal role.

But can this be what List and Menzies mean? They say that while neural investigation will not direct us to the cause of the action, it may reveal what they occasionally describe as a factor that is causally sufficient for the behavioral effect.¹² Thus they say, for example, that in the sort of case already discussed the 'intention *I1* to reach for a specific target is the cause of the reaching action *A1* even though it supervenes on the neural property *N11*, which is causally sufficient for the action' (488–9). What do they mean here by 'causally sufficient'? They cannot mean that the neural property is causally sufficient in the sense of generating or producing the behavior, since that would be to take the very line they condemn in Kim. But how then are we to interpret them?

I see two alternative ways of construing what they have in mind. One interpretation would equate a causally sufficient condition for an effect with a condition that gives us evidence that the event has a related difference-making cause. The condition might itself be the cause of the effect, as would happen if the neural property operated in the absence of any intention; or it might be something distinct from the cause, as in the case actually envisaged where the intention is the cause. The other interpretation of what counts as a causally sufficient condition is a factor that programs for the effect without necessarily super-programming for it and without necessarily constituting a difference-maker—that is, a cause in the proper sense.

I shall return at the end of the chapter to the question of how we should take what they say on this matter, as it turns out to be of decisive significance in relating their approach to the program model. But for the moment we may note that the two interpretations have rival attractions. The appeal of the evidential model is that it would restrict causal relationships to difference-makers, thereby supporting the rhetoric of the position they maintain on the exclusion principle, as we shall see in the next section; it enables them to present that position as turning the tables on Kim's exclusion claim, not just denying it. The appeal of the programming model is that it would make intuitive sense of the notion of causal sufficiency, allowing that a

¹² They use this language at various points in the text, and use it in a way that condones the notion of causal sufficiency. See for example, pp. 479–80, 488–9, and 499.

property is causally sufficient for another when its presence necessitates the presence of the other but its absence does not necessitate the absence of the other.

List and Menzies, as their focus on difference-makers entails, are anti-hierarchical about the different causal levels explored in the different sciences—and explored, as indeed they may be, within any single science. When it comes to determining the cause of a given effect, so their line suggests, we should be open to the discovery that the cause lies at a higher or a lower level; we should treat higher- and lower-level properties as equally eligible candidates for the role of difference-making cause. And if we make a choice in favor of higher-level candidates, then that should not leave us with any misgivings. ‘In such cases we have good reason to believe in the causal autonomy of higher-level properties’ (499). Higher-level causes are just as capable of pulling causal weight, on this approach, as their lower-level counterparts and there is no ground for discriminating in any way between them.

Where List and Menzies equate difference-making with causation, and reject any hierarchy of causal levels, the program model proposes a less committal view. Without making any assumptions as to what factors should count as causes, it holds that when a higher-level property like F is causally relevant to another higher-level property, G , then however such relevance is interpreted, it depends on the causal relevance of the lower-level realizer of F , RF , to the lower-level realizer of G , RG . Thus if the instantiation of F counts as the cause of the instantiation of G , under whatever account of causation we accept, the causal role it plays is derivative from the causal role played by the instantiation of RF in relation to the instantiation of RG . And what is true of the F – G relationship vis-à-vis the RF – RG relationship, is true of the RF – RG relationship vis-à-vis any still lower relationship, RRF – RRG ; and so on. There is a natural hierarchy among causal levels, and lower-level factors enjoy causal relevance in a more basic or less dependent sense; they enjoy a greater degree of efficacy, as we may wish to put it.

This hierarchical picture may be developed in any of a number of ways, depending on how causes are construed, but the program model does not commit us to any construal in particular; it merely requires that we can characterize a property that programs for another property as being causally relevant to that property. For all that the model assumes, then, one possibility is that programming properties all count as causes, in which case lower-level programmers represent more efficacious or less dependent causes. Another is that there are producers, as Kim assumes, and that when F programs for G —and when it is causally relevant to G —that is because any realizer of F would produce a realizer of G ; or because any realizer of the realizer of F would produce a realizer of the realizer of T ; or something on that pattern.¹³ And yet

¹³ But suppose that there is no bottom level, only ‘an infinite progression of levels downwards’ in the program architecture. In that case, assuming that we continue to think of causation as production, it will be an error to think that there is any causation in the actual world. Thus ‘the program story will have a different significance, bearing on relations between equally inefficacious levels’ (Jackson and Pettit 1990a: 116).

a third possibility is that we identify causes by interests unrelated to programming or production; this will be fine so long as we concede that the causal relevance among higher-level properties is still dependent on causal relevance among lower-level.

When Jackson and I introduced the program model, we did so with a view to replacing the less demanding theory that a higher-level property has causal relevance to an effect just insofar as the actual instantiation of a realizer at some lower level causally produces a realizer of the effect, however causal production is understood. The fact that a higher-level property is actually realized by a productive property does not ensure that that property is causally relevant on the program model to an effect of the realizer. It will be causally relevant only if the instantiation of any realizer of the higher-level property would have done equally well in giving rise to the effect.

Consider an example that Jackson and I gave in our initial presentation of the program model; while it does not involve different sciences, it certainly involves differences of level (Jackson and Pettit 1988: 395). Fred coughs during a musical performance and the conductor is annoyed. Does the higher-level, existential property of the situation, which consists in somebody's coughing—a property realized by Fred's coughing—count as causally relevant? Yes, if somebody's coughing programs for the conductor's irritation: that is, if the conductor would be annoyed by anyone's coughing during the performance. No, if somebody's coughing does not program for the effect: that is, if it is only Fred—perhaps an old adversary—who is liable to annoy the maestro by coughing.

Jaegwon Kim (1984) had defended the weaker view of causal relevance that we sought to replace although, as we shall see, he does not give much importance to relevance in that sense. On this view one higher-level property is causally relevant to another—it counts as a 'supervenient cause'—just in case the lower-level realizer of the property is causally related in a proper sense to the lower-level realizer of the other. As he puts it:

the general schema for reducing a macrocausal relation between two events, x 's having F and y 's having G , where F and G are macroproperties, is this: x 's having F supervenes on x 's having $m(F)$, y 's having G supervenes on y 's having $m(G)$, where $m(F)$ and $m(G)$ are microproperties relative to F and G , and there is an appropriate causal connection between x 's having $m(F)$ and y 's having $m(G)$. (Kim 1984: 262)

On Kim's criterion, a higher-level property can be causally relevant to a certain effect without programming for it; thus somebody's coughing will be causally relevant to the conductor's irritation in both of the cases we distinguished.¹⁴

¹⁴ It is worth nothing that when Kim (1998) discusses program explanation, he suggests misleadingly that it is not really distinct from his supervenient causation—defined, as we saw, in Kim (1984)—by introducing a new clause, which I put in italics in the following quotation: 'Suppose that F has P_1, \dots, P_n as its supervenience bases. This means that whenever F is instantiated, one of its base properties, P_i , is instantiated. *Suppose further that each of these is causally productive of an effect G .* We can then say that F is

Jackson and I spelt out the notion of causal relevance associated with the program model, starting from David Lewis's (1986) suggestion that causally explaining something involves giving information on the causal history of the effect.¹⁵ We say that invoking a programming property like *F* in causally explaining the instantiation of *G* gives potentially important modal information on the causal history of *G* (pace Kim 1998: 75–7). It tells us that not only did the realizer of *F* give rise to *G*—this is what we would learn on being told about supervenient causation—what also holds is that any realizer of *F* would have done so as well (Jackson and Pettit 1990a, 1992a). And in doing this, it gives us information about the causal history of the instantiation of *G* that may be useful for a variety of purposes, ranging from the predictive to the manipulative to the interactive.

Such information is bound to be particularly useful in the special sciences and in any investigations that abstract from the detail of physical realization. It will enable us to ignore the trees and orientate by the forest: to abstract from neurons and see other people or animals from within the intentional stance, for example Dennett (1987); to abstract from individual actors and recognize the regularities whereby rising unemployment leads to rising crime or rising urbanization to declining religious practice (Jackson and Pettit 1992b); or to abstract from the contributions of individual members in order to make sense of a corporation or church or state in terms of its corporate attitudes (List and Pettit 2011). More generally it will enable us to achieve goals that might not be feasible on the basis of a regular, physical understanding of what is going on: say, an understanding in terms of microphysical properties.¹⁶

3 The Challenge to the Special Sciences

The explicit motive for which List and Menzies develop their conception of causation as difference-making is to deal with a challenge that Jaegwon Kim (1998) has posed for all those who hold that non-physical properties that are realized by physical properties, as physicalism implies, can nevertheless be causally significant. Kim's

a "supervenient cause" of *G*. That was the idea of supervenient causation. But it is clear that this nicely fits the model of program explanation' (Kim 1998: 74).

¹⁵ This theory of explanation naturally supports the view that useful explanation involves giving useful information: that is, information that is useful within the context of a certain inquiry.

¹⁶ We defined programming by the fact that the presence of a property necessitates the presence of another property without its being required that the absence of the property necessitates the absence of that other property. But what to say in response to the observation that a property may necessitate an effect in that sense and be wholly irrelevant to it: say, in a classic example, that were a man to take a female contraceptive pill that would necessitate his not getting pregnant as a result of intercourse? We should say that in examples of the application of the program model, an implicit assumption is that the effect to be explained is not independently necessitated. The man's not getting pregnant is necessitated, of course, by his being male. And it is necessitated independently, not just in the way in which the action necessitated by a neural state, *NI*, is necessitated by the intention this realizes. Thanks to Peter Menzies for raising this issue.

claim, which turns on his so-called exclusion principle, is that the causal efficacy of physical properties excludes the possibility that any properties they realize should enjoy a similar efficacy (Kim 1998: 37–8).

As we saw earlier, realization entails supervenience in the sense that if a property F is realized by a property, RF , then there can be no change in F without a change in RF . According to Kim's exclusion principle, however, if RF is 'causally sufficient' for some effect, G , then no distinct property that supervenes on RF —and so nothing like F —can play a causal role in relation to the effect, G ; being a supervenient cause in the earlier terminology does not count as playing such a role. And that would mean, of course, that many factors we frequently invoke in higher-level causal explanation are not causally relevant in any significant sense; the work that they are assumed to do is actually done by their lower-level physical realizers.

List and Menzies argue that by most accounts in the literature the truth or falsity of the exclusion principle 'can be settled by an investigation of the concept of causation' (2009: 476); that given its opacity, Kim cannot rely on 'the notion of productive or generative causation' (490) in articulating the principle; and that the vague reference to causally sufficient conditions in the statement of the principle should be replaced by a reference to 'causation in a more adequate sense, understood as difference-making' (490). They propose a revised exclusion principle, defined in their preferred terms. This rejects Kim's idea that as between two competing properties for causing an effect, one of which supervenes on the other, the more basic property excludes the causal role of the other. What it says rather is that one or the other may count as the cause of the effect, but in standard cases not both: 'For all distinct properties F and F^* such that F^* supervenes on F , F and F^* do not both cause a property G ' (490).

Revising the exclusion principle in difference-making terms, List and Menzies argue that it holds in any standard cases, for example in the special sciences, where various properties are multiply realizable and the relations between certain higher-level properties are realization-invariant. According to this revised principle, there is no particular threat of exclusion from below. On the contrary, the principle is as liable to establish the exclusionary causal status of higher-level properties as it is to establish the exclusionary causal status of lower-level properties.

Without going into details, the argument for this result is straightforward. For any effect, G , in the standard sort of system envisaged—for example in a neurally realized psychological system—we can only expect to find one difference-making cause, as we saw earlier. Thus if F satisfies the difference-making conditions for G —if its presence programs for G , its absence for not- G —then that means that no other property can play the same role and the instantiation of no other property can count as a cause of G . The difference-making, causal status of F excludes the possibility that any other property should enjoy that same status. But this result may hold, whether F is a higher-level, special-science property or a lower-level, physical property. If G involves an action, for example, then the exclusionary cause of G may be the

intention to perform that action, not the particular neural realizer of that intention. And so, as List and Menzies say, the revised exclusion principle ‘can actually support, rather than undermine, the causal autonomy of special-science properties’ (2009: 476). It gives us ground, not just for denying the universality of Kim’s bottom-up exclusion, but for turning the tables and arguing that top-down exclusion is just as possible as exclusion of the bottom-up kind.

This explains how the difference-making model can respond to the challenge raised on the basis of Kim’s exclusion principle. How by comparison does the program model handle that challenge? Admitting a hierarchy of causal levels, with the instantiations of properties at lower levels being able to claim a less dependent or more efficacious form of causal relevance, this model is exposed to Kim’s challenge in a way in which the difference-making approach is not. For it may well seem that if the instantiation of a lower-level property, *RF*, is causally relevant to the instantiation of *G*—it leads to *RG*, which realizes *G*—then there is a question as to why we should want to invoke a higher-level property that it realizes, *F*, in giving a causal account of the presence of *G*. After all, by the program account, any purportedly causal relation between property-instances, *F* and *G*, derives from the more basic causal relation between the instantiation of *RF* and the instantiation of *RG*. There is a hierarchy among causal levels, according to the model, not the relative equality that List and Menzies allow.

But while the program model is more sympathetic than the difference-making approach to Kim’s premises, acknowledging that lower-level causes have a certain priority in relation to higher-level, it provides a base of resistance to the lesson that he would draw. It insists that as between two higher-level properties, say *F* and *F**, each of which supervenes on a lower-level property *RF*, one of those properties may be causally relevant to an effect, *G*, that is brought about by the instantiation of *RF*, while the other is not; one of them may program for *G*, while the other fails to do so. Thus it offers a ground for thinking that some higher-level properties can be causally relevant in a special way that makes them invulnerable to Kim’s challenge.

Reusing an example that Jackson and I took from Peter Menzies himself, suppose that someone is electrocuted and dies because of standing on a metal ladder that touches some power lines (Jackson and Pettit 1990b). The lower-level property that causes the electrocution and death is the nature of the cloud of free electrons that permeates the ladder. So what is the causal status of the higher-level properties that it realizes? Those realized properties include the electrical conductivity of the ladder, its thermal conductivity, and its visual opacity. Among those properties we would all surely agree that the electrical conductivity is causally relevant to the calamity but that the other properties are not. And that, equally surely, is because only the electrical conductivity programs for the electrocution and death. No matter how the electrical conductivity is realized, the realizer is going to lead to the same calamity. But this is not so with either of the other properties. Even if a wooden or

plastic ladder were thermally conductive in an equal measure, it would not have led to the electrocution.¹⁷

Time to sum up. List and Menzies adopt Kim's exclusion principle in a revised form and resist his challenge by showing that in that version the principle is as likely to support top-down exclusion as exclusion of a bottom-up variety. Jackson and I reject Kim's exclusion principle, at least insofar as it is meant to bear on causal relevance. We resist his challenge by arguing that higher-level properties that program for the effects of lower-level realizers have a distinctive form of causal relevance and cannot be dismissed in a full causal story about those effects.

4 Critical Appraisal

The brisk account given of the difference-making model is inadequate to the richness of List and Menzies' analysis, ignoring many of the issues they cover and neglecting the formal account that they offer of difference-making. Still, I hope that it may provide a reasonable basis for gauging the relative merits of the model, in particular its merits as a representation of higher-level causality and as a response to the challenge that Kim poses for the special sciences. In this final section I argue that whatever its merits in other respects—say, in analysis of what it is for something to count as a cause—its merits as an account of higher-level causality and as a response to Kim's challenge are not greater than the merits of the program account.

The main benefit of the difference-making model, at least on the face of it, is that it provides a sharp, well-defined response to the exclusion problem raised by Kim. It replaces the intuitive notion of productive causation with the well-articulated concept of difference-making. It offers an account of the exclusion principle that makes use of this concept. And then, turning the tables on Kim, it purports to show not only that bottom-up exclusion is not universal, but also that top-down exclusion is just as likely. Thus it seems to provide a powerful vindication of explanation at higher levels in a multi-level architecture and of the claims of the special sciences.

I am prepared to concede that the response to Kim's challenge is appealing for the way it turns the tables on Kim and that the response supported by the program model looks weaker in comparison. The program model response emphasizes that higher-level properties can have a causal relevance that goes beyond the relevance they would have under the supervenient-causation model that Kim (1984) introduces. But it leaves a grain of challenge in place, acknowledging that the lower-level properties that realize a higher-level programmer are more basic in an intuitive sense and can reasonably be held to do more basic causal work.

¹⁷ These observations argue against Kim's (1998: 70) suggestion, in criticism of Jackson and me, that the information available from a program explanation is not going to be of much value in understanding the causal history of an event.

But the benefit that the difference-making model generates in handling Kim's challenge comes at a serious cost. It appears to support an implausible principle of symmetry across the levels in a multi-level structure like that which is associated with the different sciences. If the lower-level realizer of a difference-maker does not count as a cause of the effect brought about—the difference made—then it is on a par in this respect with any still higher-level property that the difference-maker may happen to realize. And if being a difference-making cause is the only causal feature recognized, as appears to be the case in the account given by List and Menzies, then that means that the lower-level realizer is as causally irrelevant to the effect as any higher-level property that the difference-maker may itself happen to realize. There is an apparent symmetry between how higher-level causes relate to lower-level realizers, on the one side, and on the other how they relate to higher-level properties that they themselves realize. That symmetry, which the program model would reject in favor of a hierarchy of causal levels, is deeply counterintuitive.

Go back to the example in which a certain intention is the difference-maker for an action of moving my hands; its presence necessitates the movement, its absence the absence of movement. The neural realizer of the intention that causes the action is as causally irrelevant, according to the symmetry claim, as is a still higher-level property that the intention realizes: say, the existential property of there being some intention or other present—maybe this, maybe that—which is realized by the presence of the intention actually held. But that seems downright absurd. Intuitively, the intention depends on being realized by the neural state for making the difference ascribed in a way in which it does not depend on realizing the higher-level existential property. And that dependence undermines the symmetry claim entailed.

This objection to the difference-making approach only applies, of course, if the approach really is committed to the symmetry claim. The commitment is secure on the assumption, as I put it above, that being a difference-making cause is the only causal feature recognized by List and Menzies. It is because that particular feature is equally lacking in the property that realizes a difference-maker and in a property that the difference-maker realizes that the approach supports the symmetry claim. But could List and Menzies reject the assumption and ascribe a causal role to the realizer of a difference-maker that a property realized by the difference-maker is bound to lack?

This possibility takes us back to the interpretation of the claim that they make in saying that the neural realizer of the intention that causes an action, being the relevant difference-maker, is 'causally sufficient for the action' (488–9). I said that there were two possible interpretations of what they mean. On one account, they mean just that the realizer is evidence of the difference-making role of the intention. On the other, they mean that the realizer programs for the action; its presence necessitates the presence of the action even if its absence does not necessitate the absence of the action.

The attraction of the evidential interpretation is that it would keep alive the anti-hierarchical aspect of the view List and Menzies claim to defend. Thus it would make

it possible to set out a position, as they advertise, that turns the tables on Kim. Specifically, it would enable them to uphold their claim that just as lower-level difference-makers do not have their causal contribution excluded by higher-level properties that they realize, so higher-level difference-makers do not have their causal contribution excluded by lower-level properties that realize them.

But the price of maintaining such an anti-hierarchical view, as we can now see, is that it would commit List and Menzies to an utterly implausible symmetry claim. And in this light, it may seem much more appealing to interpret their talk of causal sufficiency in the other manner. On this interpretation, the realizers of difference-makers program for the effects of which the difference-makers count properly as the causes. They do one part of the two-part job that distinguishes difference-makers: their presence necessitates the presence of the effects but their absence does not necessitate the absence of the effects.

But this interpretation of the difference-making model means that the story told by List and Menzies is not so different, after all, from the story that the program model supports. It posits fundamentally the same ontology and differs only in introducing the novel concept of a difference-maker to that ontology.¹⁸

The difference between the approaches comes out in how they define various causal expressions in relation to the elements in this ontology. The program model, as we saw, is meant to be fairly ecumenical, allowing the notion of a causal relation to be attached to the programming relation, or just to the relation between properties at a presumptively basic level, or whatever. The only conceptual or semantic commitment it enforces is that there is a notion of causal relevance associated with any programming property; in particular, a notion of causal relevance that is more demanding than that of being a supervenient cause.

The difference-making model, so it now transpires, has to acknowledge this notion of causal relevance, since doing so is required to break the symmetry problem. This notion is what makes it possible to hold that the realizers of difference-makers are causally relevant in a sense that distinguishes them from properties realized by difference-makers. The novelty of the approach is to replace the ecumenism that the program model allows in what we are prepared to describe as causes and to insist that the unique deservers of that ascription are difference-makers or super-programmers: that is, properties of which it is true that their presence programs for the presence of the effects and their absence programs for the absence of the effects.

List and Menzies argue, in commenting on the proportionality of causes, that their conceptual proposal picks up important aspects of our causal talk. But it does not

¹⁸ A multi-level causal structure may bottom out, as we saw earlier, or it may be held to progress indefinitely downwards. Both the program and difference-making models can live with either ontology. They can each acknowledge the dependence of higher-level relations of necessitation on lower-level relations and can live with the idea either that this dependence bottoms out in a form of necessitation that does not depend on anything lower or that it progresses indefinitely downwards.

make a great difference whether or not we go along with the proposal. For once we see that they also have to acknowledge that every higher-level causal property—say, every difference-maker—depends for having an effect on the causal relevance of the realizers of the higher-level property, it is no longer clear that their picture is substantively different from that which is projected in the program model.

Consider their position, for example, in relation to Kim's exclusion challenge. By making room for the notion of causal relevance introduced in the program model, they are forced to acknowledge in tandem with defenders of that model that a grain of Kim's challenge remains in place; they can no longer claim to turn the tables on him. If one property makes a difference to another only in virtue of the causal relevance of the realizer of the first to the realizer of the second, doesn't that mean that there is something of substance in the thought that higher-level causes depend for their application on the causal work done by lower-level realizers?

List and Menzies may resist the suggestion, arguing that on the exclusion principle which the concept of a difference-maker enables them to define, there is top-down exclusion as well as bottom-up exclusion. But why should Kim be disturbed? He may argue for bottom-up exclusion solely on the basis that all higher-level causes depend for their relevance on the causal relevance of lower-level realizers. And then, assuming that there is a lowest level, he can invoke that dependence to urge that we should posit causal relevance—real, basic causal relevance, as it were—only at that level.

The right response to that position, it seems to me, is to say that the dependence on the causal relevance of lower-level realizers is no reason to reject causal talk at higher levels, whether this talk be vindicated by reference to programming or super-programming relationships. But this response does not depend essentially on adopting the difference-making view. It is a response of the kind that I associated earlier with the program model.

These observations suggest—to my mind happily—that the difference-making story introduced by List and Menzies does not break in any fundamental way with the program model that Jackson and I developed. Their main contribution consists, first, in the argument that we should treat all difference-makers and only difference-makers as causes; and second, in the application of this argument to standard cases of multi-level causal structures. But that conceptual contribution does not recast the ontology presupposed in the program model and does not provide any novel grounds for resisting Kim's challenge.

Should we be hospitable, however, to the conceptual innovation that List and Menzies introduce: that is, to the proposal to equate causes with difference-makers, at least under the framework assumptions introduced earlier? This is a large issue that I do not want to address in the present context. But without passing a final judgment on the idea, let me conclude by noting that there is at least one sort of difficulty that, despite appearances, it is capable of avoiding.

The difficulty can be nicely illustrated by a problem raised for the view of group agents and their fitness to be held responsible that Christian List and I defended in an

independent work, where we endorsed the program model (List and Pettit 2011: chapter 7).¹⁹ We argued that when someone who acts with authority in the name of a group agent does something culpable then, assuming a suitable organization in the group and a suitable psychology in the individual, both the group agent and the individual actor are fit to be held responsible for the act. In virtue of its organization the group programs for the act and, in virtue of his or her compliance with the group's requirements, the individual enacts that program. But given the difference-making concept of cause, how can both the group and the individual be held responsible for what is done? It appears that, depending on which is the difference-maker, it is only the group or only the individual that can count as the cause of an action and that they cannot be equally fit to be held responsible for it.²⁰

The answer to this question, and the solution to the parallel questions that other cases may raise, is that the group can be a cause in virtue of being a difference-maker for one property of the action, and the individual a cause in virtue of being a difference-maker for another. The group causes the type of action involved to be more or less robust or resilient—it ensures that there is some individual there to act in its name—and the individual causes that action-type to be instantiated in one of many possible ways. Thus the group is a difference-maker for the resilience of the action-type, and the individual is a difference-maker for the instantiation of that type: specifically, for its instantiation by that individual's hands.²¹

The difference-making story may not transform the ontology of the program model, then, and may not turn the tables on Kim's exclusion challenge. But our concluding observations show that it fits quite naturally with everything supposed in the model. It represents a potentially interesting proposal for disciplining the causal talk that we pursue, under the guidance of that model, in the context of multi-level causal structures.

References

- Dennett, D. 1987. *The Intentional Stance*. Cambridge, MA: MIT Press.
Jackson, F. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.

¹⁹ See too Pettit 2007c, 2009.

²⁰ The difficulty described in this paragraph closely reflects an argument presented by Anders Strand in a workshop on List and Pettit (2011) at the University of Oslo, May 2011.

²¹ I argue in *The Common Mind* (Pettit 1993a) that the resilience of a phenomenon is an important explanandum in social science, on a par with its emergence or survival. See too Pettit 2007b. ~~I say that the individual is a difference-maker for the instantiation of the action-type at his or her own hands. That is because the individual is unlikely to be a difference-maker for the action, period; others may be ready to play the required part, should the original person refuse to act. My preference is to say that the individual enactor of something for which the group agent programs is individually responsible—excuses aside—in virtue of producing the action: implementing the program.~~

- Jackson, F., and Pettit, P. 1988. 'Functionalism and Broad Content', *Mind*, 97: 381–400. Reprinted in Jackson, Pettit, and Smith 2004.
- Jackson, F., and Pettit, P. 1990a. 'Program Explanation: A General Perspective', *Analysis*, 50: 107–17. Reprinted in Jackson, Pettit, and Smith 2004.
- Jackson, F., and Pettit, P. 1990b. 'Causation in the Philosophy of Mind', *Philosophy and Phenomenological Research*, 50: 195–214. Reprinted in Jackson, Pettit, and Smith 2004.
- Jackson, F., and Pettit, P. 1992a. 'In Defence of Explanatory Ecumenism', *Economics and Philosophy*, 8: 1–21. Reprinted in Jackson, Pettit, and Smith 2004.
- Jackson, F., and Pettit, P. 1992b. 'Structural Explanation and Social Theory', in *Reductionism and Anti-reductionism*, ed. D. Charles and K. Lennon. Oxford: Oxford University Press. Reprinted in Jackson, Pettit, and Smith 2004.
- Jackson, F., Pettit, P., and Smith, M. 2004. *Mind, Morality, and Explanation: Selected Collaborations*. Oxford: Oxford University Press.
- Kim, J. 1984. 'Epiphenomenal and Supervenient Causation', *Midwest Studies in Philosophy*, 9: 257–70.
- Kim, J. 1998. *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*. Cambridge, MA: MIT Press.
- Lewis, D. 1986. *Philosophical Papers, Vol. 2*. Oxford: Oxford University Press.
- List, C., and Menzies, M. 2009. 'Non-reductive Physicalism and the Limits of the Exclusion Principle', *Journal of Philosophy*, 106: 475–502.
- List, C., and Pettit, P. 2011. *Group Agency: The Possibility, Design and Status of Corporate Agents*. Oxford: Oxford University Press.
- Macdonald, C., and Macdonald, G. 2007. 'Mental Causation on the Program Model', in *Common Minds: Themes from the Philosophy of Philip Pettit*, ed. G. Brennan, R. E. Goodin, F. Jackson, and M. Smith. Oxford: Oxford University Press, 1–27.
- Pettit, P. 1993a. *The Common Mind: An Essay on Psychology, Society and Politics*, paperback edition 1996. New York: Oxford University Press.
- Pettit, P. 1993b. 'A Definition of Physicalism', *Analysis*, 53: 213–23.
- Pettit, P. 1994. 'Microphysicalism without Contingent Micro-macro Laws', *Analysis*, 54: 253–7.
- Pettit, P. 1995. 'Microphysicalism, Dottism, and Reduction', *Analysis*, 55: 141–6.
- Pettit, P. 2007a. 'Joining the Dots', in *Common Minds: Themes from the Philosophy of Philip Pettit*, ed. G. Brennan, R. E. Goodin, F. Jackson, and M. Smith. Oxford: Oxford University Press, 215–344.
- Pettit, P. 2007b. 'Resilience as an Explanandum of Social Theory', in *Contingency*, ed. I. Shapiro and S. Bedi. New York: New York University Press.
- Pettit, P. 2007c. 'Responsibility Incorporated', *Ethics*, 117: 171–201.
- Pettit, P. 2009. 'Corporate Responsibility Revisited', *Rechtsphilosophie & Rechtstheorie*, 38 (Special Issue: Philip Pettit and the Incorporation of Responsibility): 159–76.
- Woodward, J. 2005. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.
- Yablo, S. 1992. 'Mental Causation', *The Philosophical Review*, 101: 245–80.