Is There High-Level Causation?

Luke Glynn

Abstract

The discovery of high-level causal relations seems a central activity of the special sciences. Those same sciences are less successful in formulating strict laws. If causation must be underwritten by strict laws, we are faced with a puzzle (first noticed by Donald Davidson), which might be dubbed the 'no strict laws' problem for high-level causation. Attempts have been made to dissolve this problem by showing that leading theories of causation do not in fact require that causation be underwritten by strict laws. But this conclusion has been too hastily drawn. Philosophers have tended to equate non-strict laws with *ceteris paribus* laws. I argue that there is another category of non-strict law that has often not been properly distinguished: namely, (what I will call) *minutiae rectus* laws. If, as it appears, many special science laws are *minutiae rectus* laws, then this poses a problem for their ability to underwrite causal relations in a way that their typically *ceteris paribus* nature does not. I argue that the best prospect for resolving the resurgent 'no strict laws' problem is to argue that special science laws are in fact typically *probabilistic* (and thus able to support probabilistic causation), rather than being *minutiae rectus* laws.

1. Introduction

The search for and discovery of causes seems a central activity of the high-level sciences, as well as of our pre- and proto-scientific attempts to understand and influence the world.[1] For example, we may think, on occasion, that an ice cube's melting was an effect of its being placed in hot water, that the genetic characteristics of the offspring were caused by those of the parents, that the Fed's lowering of interest rates was a cause of the rise in inflation, that a person's suffering lung cancer may have been caused by her smoking, that the warm sea-surface temperatures were a (contributing) cause of the hurricane, and that my desire for coffee and my belief that the cafe sells it were causes of my going to the cafe.

Examples like these seem to furnish a strong presumption in favor of the following proposition:

(i) There are many genuine instances of high-level causation.

A second proposition which we also seemingly have good reason to believe is:

(ii) There are few if any strict high-level laws.

This second proposition is supported by the observation that the high-level sciences (not to mention our folk theories) rarely if ever deliver strict laws, where a strict law is understood in something like Davidson's sense of

> "[...] a generalization that [is] not only 'law-like' and true, but [is] as deterministic as nature can be found to be, [is] free from caveats and *ceteris paribus* clauses; that [can], therefore, be viewed as treating the universe as a closed system." (1993, p. 8)

While this characterization of strict laws is less than fully perspicuous, more shall be said about the precise respects in which high-level laws lack strictness in Section 5 and Section 6.

There is a third proposition to which philosophers have sometimes subscribed that appears to be in tension with the conjunction of (i) and (ii). This is that:

(iii) All genuine causal relations must be underwritten by strict laws.

Davidson (1970), who was the first to note the tension between (i)-(iii),[2] did not provide an argument for the truth of (iii).[3] But it is certainly not without some *prima facie* plausibility. This plausibility derives from the fact that many popular philosophical theories of causation analyze that relation ultimately in terms of laws, and seem to require that the laws in question be *strict*. For example, sophisticated *regularity* theories analyze causation in terms of relations of nomic sufficiency. And it is, on the face of it, difficult to see how non-strict laws, which may admit of exceptions, could ground such relations. *Counterfactual* theories, on the other hand, analyze causation in terms of counterfactual dependence, with laws standardly taken to be (among) the truth-makers for the relevant counterfactual conditionals (see, e.g., Lewis 1973a, 1973b, 1979). But one might wonder how generalizations that admit of exceptions are able to support counterfactuals. *Probabilistic* theories, by contrast, analyze causation in terms of patterns of probabilistic dependence, with the probabilities in question usually given an objective chance interpretation. And it is often taken to be a platitude about objective chances that they derive from strict, but probabilistic laws of nature (see Schaffer 2003, pp. 36-7; 2007, p. 126). Finally, *process* theories typically cash out the notion of a causal process in terms of relations of nomic sufficiency (see Ney 2009) or counterfactual dependence (Salmon 1984).[4] If these latter relations require strict laws to underwrite them, then causal processes do so too.

So we have at least some reason to believe each of principles (i)-(iii), but there is a tension between them. This should lead us to feel some discomfort. The tension is not yet an outright inconsistency. But this owes partly to the imprecision of the principles as stated. As shall be seen in Section 2, once these principles are given plausible precisifications, they are brought closer to inconsistency. It thus seems that we have a paradox on our hands.[5] The remainder of this essay will then be devoted to examining whether this paradox can be dissolved. Existing attempts to do so (discussed in Section 5) have largely comprised attempts to show that leading philosophical theories of causation do not after all require causation to be underwritten by strict laws, but that *ceteris paribus* laws will do. Proposition (iii) is consequently rejected as false. I will argue that this is too quick. Specifically, I will argue (Section 6) that there is another category of non-strict law that has often not been properly distinguished: namely (what I will call) *minutiae rectus* laws. I will argue that if, as it appears, many special science laws are *minutiae rectus* laws – often in addition to being *ceteris paribus* laws – then this poses a problem for their ability to underwrite causal relations in a way that their including *ceteris paribus* hedges does not. I will argue (Section 7) that the best prospect for completing the solution to the 'no strict laws' problem for high-level causation would be to establish that special science laws are in fact *probabilistic* laws (and thus able to support probabilistic causal relations), rather than being *minutiae rectus* laws. I will examine the plausibility of such a solution (Section 8).

Before proceeding, it is worth nothing that the 'no strict laws' problem for high-level causation, with which this paper is concerned, is not the only problem for high-level causation discussed in the literature. Another is the so-called 'exclusion' problem. Like the no strict laws problem, the exclusion problem was originally raised in connection with mental causation,[6] but generalizes to other cases of high-level causation.[7] By focusing on the no strict laws problem, I do not wish to downplay the significance of other obstacles to a philosophical vindication of high-level causation, but just to recognize that they can't all be tackled at once – at least not in a single paper! I also think that the problem is sufficiently independent from others to make this piecemeal approach worthwhile. For example, someone showing that high-level causation would, if it existed, not be *redundant* (thus solving the exclusion problem) would not thereby have shown either that there exist strict high-level laws to underwrite it, or that it does not stand in need of such laws (thus solving the no strict laws problem).[8] Nor would a solution to the no strict laws problem *ipso facto* constitute a solution to the exclusion problem.

2. The 'No Strict Laws' Problem Formulated More Precisely

In the previous section, an apparent problem concerning high-level causation was identified. This is that the supposition that (i) *there are many genuine instances of high-level causation*, is seemingly in tension with the suppositions (which we also have at least some *prima facie* reason to believe) that (ii) *there are few if any strict high-level laws*, and that (iii) *all genuine causal relations must be underwritten by strict laws*. The purpose of this section is to show that the tension is not relieved, but rather sharpened, when (i)-(iii) are precisified in plausible ways. The threat of paradox is therefore genuine.

In order to see how (i) can be rendered more precise, note that there appears these days to be a consensus among philosophers that a proper accommodation of high-level causation must allow for the causal efficacy of high-level properties, as opposed to merely the events or objects that exemplify those properties. For instance it is argued that, in order to avoid rendering mentality epiphenomenal, it must be in virtue of their *mental* properties (such as *being a pain*) that mental events have certain of their effects (such as *anxiety*, *fear*, or *wincing*).[9] Arguments along these lines have been deployed to show the inadequacy of Davidson's (1970, 1993) proposed solution to the 'no strict laws' problem as it applies to the mental. Roughly, his solution (known as 'Anomalous Monism') was to identify mental events with physical events, thus showing that in spite of the lack of strict psychological (and psycho-physical) laws, mental events can enter into causal relations (in virtue of instantiating strict physical laws). The problem is that, on such an account, it seems as though mental properties are epiphenomenal, since it is not *in virtue* of having mental properties (but rather in virtue of having physical properties) that mental events instantiate strict (physical) laws.[10]

A converse argument also seems to apply. It might reasonably be said that, to do justice to our intuitions about the causal interactions entered into by the mental, it must be allowed that mental properties can themselves be causal consequences. Thus one might insist that it was the *fearfulness* of a certain mental event (and not, or not only, its (say) *being an agitation of the amygdalae*) that was a consequence of the *painfulness* of one of its antecedents.

Both arguments apply to high-level properties more generally. For example, it appears that the same intuition (and scientific theory) that takes it to be the case that the placing of an ice cube in hot water was a cause of its melting demands, more specifically, that it was in virtue of the water's *hotness* that the melting occurred. Again, it seems specifically that it was the latter event's *being a melting* (and not, or not only, its *being an event comprising so-and-so very specific changes to the properties of such-and-such particular quarks, leptons, and bosons*) that was a causal consequence of the hotness.

Such arguments suggest that the reasons for believing (i) are reasons, more specifically, for believing (1):

(1) It is true, for many interpretations of X, Y, x, and y such that X and Y are high-level properties and x and y are individuals (objects or events) possessing those properties, that x's being X was a cause of y's being Y.[11]

Moreover, (ii) and (iii) might reasonably be precisified as (2) and (3) respectively:

(2) There are few, if any, interpretations of X and Y such that X and Y are high-level properties and there is a strict law relating X and Y.

(3) If x's being X is a cause of y's being Y, then there must be a strict law relating X and Y.

If (i)-(iii) were in tension with one another, then (1)-(3) appear outright inconsistent. Moreover, I take it that (2) is a relatively uncontentious precisification of (ii). It merely incorporates the assumption that the defining feature of a high-level law is that it relates high-level properties. As shall be seen in Section 4 and Section 5, this seems to be the standard conception of a high-level law and the one that philosophers have been operating with when they have doubted the existence of *strict* high-level laws.

Why should we accept (3) as a precisification of (iii)? The main reason is that, to the extent that leading accounts of causation are committed to (iii), they seem committed more specifically to (3). For example, if a lawful regularity account of causation is correct then, unless there is a strict law (capable grounding relations of nomic sufficiency) relating *X* and *Y*, it will not be specifically *x*'s being *X* that is causally efficacious in bringing about *y*'s being *Y*. Suppose, on the other hand, that a counterfactual account of causation is correct. Then the causal efficacy of *x*'s being *X* for *y*'s being *Y* is grounded in counterfactuals like *'If x had been X, then y would have been Y'* and (if there's no pre-emption) *'If x hadn't been X, then y wouldn't have been Y'*. But, in order to support such counterfactuals, it seems *prima facie* plausible that there must be strict, non-exceptionless laws relating the presence and absence of *X* to the presence and absence of *Y*. Again, on a probabilistic analysis, it appears plausible that, unless there is a strict probabilistic law relating *X* to *Y*,[12] then the *X-ness* of *x* will not be probabilistically and therefore causally relevant to the *Y-ness* of *y*.[13] Finally, insofar as process theories appeal to nomic, counterfactual, or probabilistic relations in cashing out the notion of a causal process, it seems that the foregoing considerations ought also to lead their adherents to endorse (3).[14]

It was suggested in Section 1 that there are *prima facie* reasons for believing (i)-(iii), which are in tension with one another. It has now been seen that (i)-(iii) can be precisified as the apparently outright inconsistent (1)-(3) without diminishing our justification for believing these principles. We are thus threatened with paradox. In Section 5, the most common way of resolving this paradox will be discussed. This involves rejecting (iii) (and its precisification (3)) by arguing that causal relations needn't be underwritten by strict laws after all. Specifically, it is argued that *ceteris paribus* laws, which are non-strict and liable to admit of exceptions, and which figure heavily in the high-level sciences, can underwrite causal relations. It is thus concluded that the 'no strict laws' problem for high-level causation is a pseudo-problem. In Section 6, I will argue that this conclusion has been too hastily drawn. I will argue that there is another category of non-strict law – *minutiae rectus* laws – that has not been properly distinguished. If, as it appears, many special science laws are *minutiae rectus* laws, then this poses a problem for their ability to underwrite causal relations in a way that their typically *ceteris paribus* nature does not. In Section 7 and Section 8, I will examine the question of whether special science laws might turn out simply to be *probabilistic* laws (and thus able to support probabilistic causal relations), rather than being *minutiae rectus* laws.

In order to explicate the two respects in which high-level laws appear to be non-strict – namely, by being *ceteris paribus* laws and by being *minutiae rectus* laws – I shall (in Section 4) outline a simple model of how we might (as physicalists who wish to make room for high-level causation) hope that strict high-level laws can be derived from strict microphysical laws. The respects in which high-level laws appear to be non-strict will then be illustrated (in Section 5 and Section 6) with reference to ways in which this simple model breaks down. Before doing this, however, it is important to note that both the model itself and the subsequent explication of its shortcomings presuppose a certain widely-held view of the relationship between the *low-level* properties of concern to fundamental physics and the *high-level* properties of concern to the special sciences, as well as to macrophysical sciences like thermodynamics.[15] It is the task of the next section to get clearer about the nature of this view.

3. High- and Low-Level

On a widely held view, high-level properties *supervene* on, but are *multiply realizable* by, basic physical properties. This view has been described as the "reigning orthodoxy" when it comes to the relationship between basic physical properties and *mental* properties (Yablo 1992, p. 254). And, when philosophers have turned their attention to high-level properties more generally, they have found it no less natural a view to take (see Kim 1979, p. 39; 1984, pp. 261-2; and Block 2003, esp. p. 142). Indeed, Kim (1998, p. 38) suggests that *supervenience* can "usefully be thought of as defining minimal physicalism" (cp. also Lewis 1983, pp. 361-4); he notes that even an emergentist should accept it. On the other hand, Putnam (1967, pp. 44-5), observing the diversity of actual and possible physical systems that may have mentality, has argued that *multiple realizability* is an overwhelmingly plausible thesis concerning *mental* properties. Fodor (1974) has influentially extended the thesis to the properties of concern to the special sciences more generally.

To say that high-level properties *supervene* on fundamental physical properties is to say that what high-level properties an object has is determined by what fundamental physical properties it has. More precisely, let H be the set of all high-level properties and let P be the set of all fundamental physical properties. Then the properties in H supervene on those in P just in case it is true that:

(**S**) Necessarily, for each property h_i in H, if any object x has h_i, then there is some property p_j in P such that x has p_j and, necessarily, anything that has p_j has h_i.[16]

The thesis that high-level properties are *multiply realizable* by fundamental physical ones is plausibly captured by the conjunction of (**S**) with (**M**) (see Yablo 1992, p. 256):

(**M**) For each property h_i in H it is true that, for each property p_j in P that necessitates h_i, possibly something has h_i but not p_j.[17]

The conjunction (**M**) + (**S**) characterizes a view of high-level properties as *asymmetrically necessitated* by basic physical properties (Yablo 1992, p. 256). Indeed, the conjunction (**M**) + (**S**) might be seen as giving a precise meaning to the notion of one set of properties being *higher-level* than another. The standard view is then that the properties of concern to the special sciences and to (certain)[18] macrophysical sciences bear the relation characterized by (**M**) + (**S**) to those of basic physics.

It is against the backdrop of this standard view that contemporary discussion of high-level causation has largely taken place. I shall assume it to be correct in subsequent sections, where I shall outline a simple model of how strict high-level laws might be thought derivable from strict fundamental physical laws, and then discuss two ways in which this simple model appears to break down. This will help to illustrate the two respects in which high-level laws appear typically to be non-strict: namely, in respect of being *ceteris paribus* laws, and in respect of being *minutiae rectus* laws. But before turning to these tasks, it will be helpful to conclude the present section by illustrating the relation of asymmetric necessitation characterized by (**M**) + (**S**) with reference to thermodynamics, which will be discussed in some detail in sections 6-8.

A *thermodynamically isolated system* is a region of space the boundaries of which are not crossed by matter or energy. The *thermodynamic state* of such a system is specified by partitioning the system into small (but macroscopic) spatial sub-regions and specifying the values taken by various macro-variables – temperature, pressure, mass density, chemical composition, etc. – in each of the sub-regions. The *microstate* of the system, on the other hand, is specified by giving the position and momentum of each of the N molecules that it comprises. This involves giving six coordinates – three spatial and three momentum – for each of the N particles. In other words, it involves specifying the location of the system in *6N-dimensional phase space*. A given thermodynamic state is compossible with infinitely many such microstates: it is associated, not with a point, but with a *region* of phase space. The system's property of being in the macrostate (that is, the thermodynamic state) in question supervenes (in sense (**S**)) upon the set of properties each consisting in its being at a certain point in the associated region of phase space. But it is also multiply realizable by them (in sense (**M**)) since, for any given point in that region, the system could be in the macrostate in question without being at that particular point in phase space.[19]

4. A Simple Model

The standard assumption that the properties of concern to the special sciences and to certain branches of macrophysics are multiply realizable by fundamental physical properties helps to explain why the causal efficacy of such properties can't immediately be underwritten by their figuring in strict fundamental physical laws. Only perfectly natural microphysical properties figure in such laws.[20] Yet it follows directly from their multiple realizability that such properties are not perfectly natural (see Lewis 1983, p. 357): in the vocabulary of basic physics they are characterizable only disjunctively.

One might, however, think that strict high-level laws can be derived from strict basic physical laws in accordance with something like the following simple model. If adequate, the model would show that high-level causation can be underwritten *indirectly* by strict basic physical laws (via the derivability from them of strict high-level laws).[21] In Section 5 and Section 6, it will be shown that the model is flawed. But it nevertheless serves a useful heuristic function, since the respects in which high-level laws appear to be non-strict – namely by being *ceteris paribus* and *minutiae rectus* laws – can usefully be explicated with reference to the model's failings.

The model is as follows. Let A, B, C, and D be high-level properties. For simplicity, suppose that each of these properties is realizable by just two (incompatible) microphysical properties. Specifically, suppose that A is realizable by the incompatible microphysical properties α_1 and α_2, B by β_1 and β_2, C by γ_1 and γ_2, and D by δ_1 and δ_2. Suppose that, as well as being (at least nomically) incompatible with one another, each of the αs is (at least nomically) incompatible with each of the βs. And suppose that each of the γs is incompatible with each of the δs. The realizability relations that obtain are encoded in the following set of bridge laws:[22]

$$\forall x(Ax \leftrightarrow \alpha_1 x \vee \alpha_2 x) \qquad (B1)$$

$$\forall x(Bx \leftrightarrow \beta_1 x \vee \beta_2 x) \qquad (B2)$$

$$\forall x(Cx \leftrightarrow \gamma_1 x \vee \gamma_2 x) \qquad (B3)$$

$$\forall x(Dx \leftrightarrow \delta_1 x \vee \delta_2 x) \qquad (B4)$$

Suppose, moreover, that the following microphysical laws obtain:[23]

$$\forall x(\alpha_1 x \rightarrow \gamma_1 x) \qquad (L1)$$

$$\forall x(\alpha_2 x \rightarrow \gamma_2 x) \qquad (L2)$$

$$\forall x(\beta_1 x \rightarrow \delta_1 x) \qquad (L3)$$

$$\forall x(\beta_2 x \rightarrow \delta_2 x) \qquad (L4)$$

It follows, on the assumption that the property of lawhood is preserved under logical consequence,[24] that the following two strict high-level laws obtain:[25]

$$\forall x(Ax \rightarrow Cx) \qquad (H1)$$

$$\forall x(Bx \rightarrow Dx) \qquad (H2)$$

The nomic relations that obtain between *A* and *C* and their realizers are illustrated in Figure 1 (an isomorphic diagram could be drawn for *B* and *D* and *their* realizers). In the diagram, single-headed arrows represent 'dynamic' laws relating properties of a single level, while double-headed arrows represent 'bridge' laws relating properties belonging to different levels (for a similar representation, see Fodor 1974, p. 109).

Figure 1.

Where strict high-level laws are derivable as per this simple model, high-level causal relations are readily accommodated. For suppose that some individual *x*, which is *A*, is also *C*. Then, since *x*'s being *A* was nomically sufficient for its being *C*, a *nomic regularity* account can plausibly allow that *x*'s being *A* was a cause of its being *C*. Suppose moreover that, if *x* hadn't been *A*, it would have been *B*. Then it would also have been *D* and not *C* (*C* and *D* are lawfully incompatible in virtue of the lawful incompatibility of their realizers). A *counterfactual* account can therefore plausibly also accommodate the causal efficacy of x's being A for its being C. So, if the simple model provides an accurate representation of reality, then it seems that there may be high-level causal relations underwritten by strict high-level laws that derive from the strict microphysical laws.

5. *Ceteris Paribus* Laws

There is, however, good reason to think that there must be something wrong with the simple model. For it is a common observation that the high-level sciences rarely (if ever) deliver strict laws. For one thing, it is commonly observed that special science laws are typically hedged by *ceteris paribus* clauses specifying that the generalization holds only under normal, or even ideal, conditions.[26] Such laws admit of exceptions where these *ceteris paribus* clauses aren't satisfied. Davidson gives the following explanation of why *psychological* laws, unlike fundamental physical laws, typically contain *ceteris paribus* clauses:

> "Physical theory promises to provide a comprehensive closed system guaranteed to yield a standardized, unique description of every physical event couched in a vocabulary amenable to law. [...] It is not plausible that mental concepts alone can provide such a framework, simply because the mental does not [...] constitute a closed system. Too much happens to affect the mental that is not itself a systematic part of the mental." (Davidson 1970, p. 99; see also his 1974, p. 43)

Davidson thinks that the point generalizes to at least most of the special sciences, as well as to our practical lore (1970, p. 94; 1993, p. 9). In general, the special sciences and our folk theories cannot be expected to yield strict, non-*ceteris paribus* laws because, unlike basic physics, they are not fully comprehensive in their subject matter.[27]

The reason a science's failure to characterize a comprehensive closed system makes for laws that hold only *ceteris paribus* is that a lack of comprehensiveness means that the interference of factors from outside that science's subject matter is always possible. Its laws therefore hold only where there is no such interference, or where the interference is not such as to make a difference.

As Fodor observes (following Davidson 1970, p. 94), the fact that the special sciences typically don't characterize comprehensive closed systems means that the *ceteris paribus* clauses in their laws are ineliminable:

> "Special science laws are unstrict not just de facto, but in principle. Specifically, they

are characteristically '*heteronomic*': You can't convert them into strict laws by elaborating their antecedents. One reason why this is so is that special science laws typically fail in limiting conditions, or in conditions where the idealizations presupposed by the science aren't approximated; and, generally speaking, you have to go outside the vocabulary of the science to say what these conditions are. Old rivers meander, but not when somebody builds a levee. Notice that 'levee' is not a *geological* term. (Neither, for that matter, is 'somebody.')" (Fodor 1989, p. 69n)[28]

The fact that the special sciences are not fully comprehensive in their subject matter means, not only is there always a possibility of interference from outside the systems they seek to characterize, but also that these influences cannot be captured within the vocabulary of the special science in question. Their laws therefore contain, and contain necessarily, *ceteris paribus* hedges which are satisfied only where the interference is not such as to make a difference.

It thus appears that something must be wrong with the simple model described in the previous section. The high-level laws derived using that model did not contain *ceteris paribus* hedges. They said that all *A*s are *C*s and that all *B*s are *D*s, not *ceteris paribus*, but *simpliciter*. Yet high-level laws are typically not like that.

What has gone wrong, I suggest, is the following. The relations of nomic sufficiency that physics uncovers are not *local*.[29] When some microphysical property *m* is instantiated throughout space-time region *r*, typically nothing less than the total microphysical state obtaining on the region of a space-like hypersurface that bisects *r*'s past light cone (or rather: that bisects the past light cones of each of the space-time points in *r*) is nomically sufficient for *m*, since only such a comprehensive state of affairs is such as to exclude any possible interfering factors. But where *A* and *C* are ordinary, locally-instantiated, high-level properties (e.g. the property of being a young, non-meandering river, and that of being an old, meandering river), their microphysical realizers are presumably more localized than microstates of large regions of space-like hypersurfaces. Consequently, the micro-realizers α_1, α_2, etc., of *A* will not alone be nomically sufficient for the micro-realizers γ_1, γ_2, etc., of *C*. Generalizations like (*L*1) and (*L*2) will therefore not be strict laws (as was supposed in outlining the simple model) but will themselves hold at best *ceteris paribus*, in the absence of interference from microphysical states of affairs not mentioned in their antecedents. The possible interfering microphysical states may realize further interesting macroscopic properties (such as the property of being a person building a levee). In any case, the *ceteris paribus* nature of microphysical generalizations like (*L*1) and (*L*2) infects high-level generalizations like (*H*1) that are derived from them. Since the former hold only *ceteris paribus*, so too do the latter.[30]

The special sciences not being comprehensive, the possible interfering factors (even where they realize interesting macroscopic properties) need not all be characterizable in the same special science vocabulary that includes *A* and *C*. Thus, as Fodor notes, in attempting to formulate a strict law about old rivers meandering, we would have to include in the antecedent a negative condition specifying the absence of people building levees. But this would involve drawing upon non-geological concepts. By contrast, being comprehensive, physics can hope to frame laws that include reference (in the physical vocabulary) to all potentially relevant factors (cp. Schurz 2002, pp. 369-70). The *ceteris paribus* nature of microphysical generalizations like (*L*1) and (*L*2) is therefore remediable, since their antecedents can be elaborated in the requisite manner without going outside the physical vocabulary.[31]

So there are good reasons for thinking that special science laws are in principle non-strict, in the sense that they contain ineliminable *ceteris paribus* clauses, and admit of exceptions when these clauses aren't satisfied. If strict laws are needed to underwrite causal relations, it seems that we are in danger of being forced to conclude that there is no genuine high-level causation, which is surely an unacceptable result. There is therefore a problem of explaining how there can be high-level causation even though it appears that there aren't strict high-level laws to underwrite it.

By far the most common response to this problem has been to deny that causal relations must in fact be underwritten by strict laws (i.e. to deny proposition (iii) from Section 1). The idea has been to argue that, while high-level laws typically contain *ceteris paribus* clauses, and are thus non-strict, this does *not* (*contra* Davidson) stand in the way of their ability to underwrite causal relations. The point is one that has been argued for by LePore and Loewer (1987) and by Fodor (1989). A related argument has been advanced by Hitchcock and Woodward (2003a, 2003b) and by Woodward (2003). I shall only briefly outline their arguments here. The idea, in each case, has been to argue that, contrary to initial appearances, plausible philosophical accounts of the nature of causation allow that *ceteris paribus* laws are sufficient to underwrite causal relations.

Fodor (1989) focuses upon the ability of *ceteris paribus* laws to furnish the relations of robust sufficiency appealed to in nomic regularity accounts of causation. In the following passage, he makes the point with reference to a mental predicate *M* and a behavioral predicate *B* (though he takes it to apply to the non-psychological special sciences too):

> "The first – and crucial – step in getting what a robust construal of the causal responsibility of the mental requires is to square the idea that Ms are nomologically sufficient for Bs with the fact that psychological laws are hedged. How can you have it *both* that special laws only necessitate their consequents ceteris paribus *and* that we must get Bs *whenever* we get Ms. Answer: you can't. But what you can have is just as good: viz., that if it's a law that M → B ceteris paribus, then it follows that you get Bs whenever you get Ms *and* the ceteris paribus conditions are satisfied." (p. 73)

Fodor's point is that, while a law specifying that all *M*s are *B*s *ceteris paribus* doesn't give us an unconditional nomic sufficiency of something's being *M* for its being *B*, it does give us a nomic sufficiency of the former for the latter *when the* ceteris paribus *clause is satisfied*. It may not be that we know, or can list, all of the conditions that must hold for it to be satisfied[32] (and certainly we may not be able to list them all without going outside the special science vocabulary in question). But it will suffice that there is some fact of the matter about when these conditions do indeed all hold.[33] According to Fodor, this is what

> "[...] captures the difference between the (substantive) claim that Fs cause Gs ceteris paribus, and the (empty) claim that Fs cause Gs except when they don't." (*Op. cit.* p. 73)

If Fodor's reasoning is correct, then *ceteris paribus* laws can ground the relations of nomic sufficiency in terms of which regularity theorists typically analyze causation.

LePore and Loewer (1987, pp. 640-2) argue – drawing upon a point made by Lewis (1973b, pp. 563-4) – that *ceteris paribus* laws can support counterfactuals. Suppose, for example, that there is a *ceteris paribus* law *L* that associates properties in the set $\{P_1, \ldots, P_n\}$ with those in the set $\{Q_1, \ldots, Q_n\}$ according to the following pattern: $P_1x \rightarrow Q_1x$, $P_2x \rightarrow Q_2x$, ..., and $P_nx \rightarrow Q_nx$. Suppose, moreover, that the *ceteris paribus* conditions associated with *L* are satisfied on some occasion for some object, *a*. And suppose that the fact that they are satisfied is counterfactually independent of which of P_1, ..., or P_n is instantiated by *a* (in the sense that they would be satisfied no matter which of P_1, ..., or P_n it instantiated). Then it is true that $P_1a \mathrel{\Box\!\!\rightarrow} Q_1a$, $P_2a \mathrel{\Box\!\!\rightarrow} Q_2a$, ..., and $P_na \mathrel{\Box\!\!\rightarrow} Q_na$. Likewise, if *a* actually instantiates P_i and Q_i and it is true that, if *a* had not instantiated P_i, then it would have instantiated some other P ($\neq P_i$), then (if the *Q*s are pairwise incompatible) it will also be true that $\sim P_ia \mathrel{\Box\!\!\rightarrow}$

$\sim Q_i a$. Consequently, *ceteris paribus* laws can (at least when these various assumptions are met) support the relations of counterfactual dependence in terms of which the counterfactualist analyzes causation.

Hitchcock and Woodward (2003a, 2003b) and Woodward (2003) also argue that non-strict laws are able to support the counterfactuals relevant to causation. They develop the point in the context of the recent tradition of attempts to analyze causation using so-called 'structural equations models'.[34] Structural equations express functional dependencies between variables. An example given by Hitchcock and Woodward (2003a, pp. 4-5) is of a linear regression equation relating the amount of water (X_1) and fertilizer (X_2) received by a plant to the plant's height (Y).

(*) $Y = a_1X_1 + a_2X_2 + U$

Here a_1 and a_2 are fixed coefficients and U is an error term, representing other influences on Y besides X_1 and X_2. Suppose that, for some particular plant, the actual values of X_1 and X_2 are given by (1) and (2):

(1) $X_1 = x_1$

(2) $X_2 = x_2$

Jointly, (*), (1), and (2) constitute a structural equations model: that is, a set of exogenous values – given by (1) and (2) – and a set of functional dependencies – given, in this case, by (*) alone – that allow the calculation of the values of the endogenous variables in our model (in this case just Y). Specifically, provided that (*), (1), and (2) are all correct, then (modulo U) this allows us to calculate the actual value $Y = y$ of Y.

But our model does not simply allow for the calculation of the actual value of Y. It also encodes a set of counterfactuals. In particular, equation (*) is a counterfactual-supporting generalization, in the sense that "it gives us information about how the height of the plant depends upon the amount of water and fertilizer that it receives" (Hitchcock and Woodward 2003b, p. 183). How the value of Y would differ under counterfactual suppositions about the values of X_1 and X_2 can be evaluated (modulo U) with respect to our model by replacing equations (1) and (2) with equations (1') and (2') setting X_1 and X_2 equal to those alternative values, and then plugging those alternative values into the equation (*).

Structural equations accounts analyze causation in terms of the counterfactuals encoded in 'appropriate' models.[35] Consequently, they are a variety of counterfactual analysis. According to typical such analyses, the actual value x_1 of X_1 is a cause of the actual value y of Y just in case there is some possible value x_1' ($\neq x_1$) of X_1, some possible value y' ($\neq y$) of Y, and some 'permissible' value x_2^* (possibly $= x_2$) of X_2 such that, if X_1 had taken the value $X_1 = x_1'$, while X_2 had taken the value $X_2 = x_2^*$, then Y would have taken the value $Y = y'$. Many accounts take the *actual value* x_2 of X_2 to be a permissible value.[36] Consequently, they imply that the actual amount of water ($X_1 = x_1$) given to the plant was a cause of its actual height ($Y = y$) if there is some possible change to the amount of water given to the plant (to $X_1 = x_1'$) that would have made a difference to the height of the plant (i.e., would have changed the value of Y to some non-actual value $Y = y'$) if the actual amount of fertilizer given to the plant had been held constant ($X_2 = x_2$).

The key point for present purposes is that, as Hitchcock and Woodward argue, the counterfactual-supporting generalizations encoded by structural equations like (*) needn't be strict laws, but merely what they call 'invariant generalizations', which may admit of exceptions. They observe (Hitchcock and Woodward 2003a, p. 5) that (*) falls short of the standards of a strict, exceptionless law because, for example, it may fail under changes of background conditions not represented in (*) (e.g. if the plant were sprayed with weed killer). Nevertheless, the procedure described above for evaluating counterfactuals – simply substituting equations (1) and (2) with equations (1') and (2') specifying alternative values for X_1 and X_2, and then calculating the value of Y in accordance with (*) – can be thought of as a procedure for arriving at the 'closest possible world(s)' in which X_1 and X_2 take these alternative values (see Hitchcock 2001, p. 283). These are worlds in which background conditions are held constant or at least in which significant interfering factors (such as weed killer) are held absent: that is, worlds in which *ceteris* are held *paribus*.[37] Generalizations like that encoded by (*) are thus able to support the sort of counterfactual dependencies to which structural equations analyses of causation appeal, even though they appear to be mere *ceteris paribus* laws.

Several philosophers equate non-strict with *ceteris paribus* laws. Here are LePore and Loewer doing just that:

> "A nonstrict law is a generalization that contains a *ceteris paribus* qualifier that specifies that the law holds under 'normal or ideal conditions,' [...]. The generalizations one finds in the special sciences are mostly of this kind. In contrast, a strict law is one that contains no *ceteris paribus* qualifiers; it is exceptionless not just *de facto* but as a matter of

law." (1987, p. 632; see also p. 640)

Fodor (1989, p. 69) makes a similar equation.[38] The fact that Fodor and Lepore and Loewer make this equation may help to explain why they think that, having (respectively) shown *ceteris paribus* laws to be able to underwrite relations of nomic sufficiency and to support counterfactuals, the 'no strict laws' problem for high-level causation has been solved. But this equation seems to me to be *prima facie* wrong. There is a way in which generalizations can fail to be strict laws even though they do not contain *ceteris paribus* clauses. It is to the discussion of this that I now turn.[39]

6. *Minutiae Rectus* Laws

Not all exceptions to high-level generalizations arise due to the non-fulfillment of *ceteris paribus* clauses. That this is so is perhaps best illustrated with reference to a high-level generalization that appears not to include a *ceteris paribus* clause: namely, the Second Law of Thermodynamics.

The Second Law states that the total entropy of an isolated system increases over time (until equilibrium is reached, after which point it doesn't decrease).[40] It is well known that the Second Law admits of possible exceptions. Given an initial non-equilibrium state of an isolated system, it is *possible* though (for systems comprising a large number of molecules) incredibly unlikely that the microstate should be one that leads by the fundamental dynamic laws to a later state in which the system is still further from equilibrium: that is, in which the entropy of the system is *lower* than it was to start with.

Though the Second Law admits of such exceptions, these exceptions *do not* arise due to failures of a *ceteris paribus* clause to be satisfied. Plausibly, the Second Law doesn't contain a *ceteris paribus* clause. On the contrary, it includes a precise specification of its scope of application: it applies to thermodynamically isolated systems (including the universe as a whole). The reason that it does not contain a *ceteris paribus* hedge is that, unlike the laws of many high-level sciences, the properties that it relates – entropy levels of isolated systems at times – are not local, but rather global properties of the isolated systems in question. Consequently, there is no possibility of 'interference' from factors outside the space-time region over which these properties are instantiated. (The Second Law is therefore quite unlike the geological law *rivers become more meandering with age* – the property of being a young, relatively unmeandering river is not a global property of an isolated system.)

If, however, someone wishes to dispute the claim that the Second Law is not a *ceteris paribus* law – perhaps by insisting that its appeal to an ideal isolated system somehow amounts to a *ceteris paribus* clause[41] – then I shall not insist upon the point. The point that I *do* wish to insist upon is rather that there is a type of possible exception to it that has nothing to do with the violation of any *ceteris paribus* clause. Specifically, there is a class of exception of which it admits that is not due to the possible failure of its idealizations to obtain. Even assuming an ideal isolated system, the Second Law may be violated just as a consequence of certain unlikely microphysical realizations of the system's initial thermodynamic state.[42]

The type of exception under discussion is one of which many high-level generalizations admit, including many of those that clearly *do* include *ceteris paribus* clauses. Rather than having to do with the violation of the *ceteris paribus* clauses of high-level generalizations (due to influences from outside the systems they seek to characterize), this type of exception is a result of the *multiple realizability* of the properties that high-level generalizations relate. It is consequently a type of exception of which fundamental physical laws do not admit. We might call laws that admit of this sort of exception '*minutiae rectus* laws': laws that hold only when the high-level properties mentioned in the antecedent are realized in the right microphysical way.[43] That laws that *do* contain *ceteris paribus* clauses (due to their failure to characterize comprehensive closed systems) often admit of this type of exception also (due to the multiple realizability of the properties that they relate) means that many *ceteris paribus* laws are also *minutiae rectus* laws. Such laws thus admit of exceptions even when their *ceteris paribus* clauses are satisfied. The Second Law is an example of a *minutiae rectus* law that might *not* also count as a *ceteris paribus* law.

This second respect of non-strictness points to a limitation of the simple model (described in Section 4) that is different to, and independent of, that indicated by the typically *ceteris paribus* nature of high-level laws. The limitation is that (even *modulo* worries about their locality and the consequent possibility of interference) it is problematic to assume that the high-level properties of interest to the special sciences have microphysical realizers that straightforwardly map onto one another in the manner suggested by (*L*1)-(*L*4) and (*B*1)-(*B*4) (cp. Fodor 1974, pp. 111-12). A simple adjustment to our model will ensure that there is no straightforward mapping and that the exceptionless high-level laws (*H*1) and (*H*2) are no longer derivable. Specifically, suppose that the model is as before except that, rather than the microphysical laws being (*L*1)-(*L*4), they are instead (*L*1) and (*L*4)-(*L*6):

$$\forall x(\alpha_1 x \rightarrow \gamma_1 x) \qquad (L1)$$

$$\forall x(\beta_2 x \rightarrow \delta_2 x) \quad (L4)$$

$$\forall x(\alpha_2 x \rightarrow \delta_1 x) \quad (L5)$$

$$\forall x(\beta_1 x \rightarrow \gamma_2 x) \quad (L6)$$

It is now *not* the case that the αs simply map onto the γs and the βs onto the δs. We therefore cannot derive the strict high-level laws (*H*1) and (*H*2).

In our simple model, each of the high-level properties *A*, *B*, *C*, and *D* is realizable by merely two microstates. But the properties of concern to the high-level sciences (beliefs and desires, thermodynamic states, rates of inflation and economic growth, DNA structures, mineral properties, and so on) are typically microphysically realizable in a much larger (perhaps often infinite) number of ways. And in general, even assuming microphysical determinism, it needn't be the case that (all or only) the microphysical realizers of one high-level property of interest straightforwardly map onto (only or all) the microphysical realizers of another, as would be required for there to be an exceptionless, deterministic high-level law connecting the two. Indeed, the greater the number of possible micro-realizers of two high-level properties, the greater the number of possible ways there are for straightforward mapping to fail. It should therefore come as no surprise if high-level laws are rarely exceptionless, even when their *ceteris paribus* clauses (if any) are fulfilled.

The non-exceptionlessness of high-level laws due to their *minutiae rectus* nature poses a problem for the accommodation of high-level causal relations. The observations (due to Fodor, Lepore and Loewer, and Hitchcock and Woodward) that *ceteris paribus* laws can support relations of nomic sufficiency and counterfactual dependence hold true only on the assumption that the laws in question are exceptionless when their *ceteris paribus* clauses are fulfilled. But we have now seen that there is reason to think that high-level laws are not, in general, like that.

In the modified version of our model, in which the deterministic microphysical laws do not straightforwardly map the αs onto the γs and the βs onto the δs, it seems problematic to accommodate a causal relation between x's being *A* and its being *C*. Since some realizers of *A* (namely α_2s) now no longer lead to realizers of *C*, being *A* is no longer nomically sufficient for being *C*: something could be *A* (specifically α_2) without being *C* (and indeed it may be that there are nearby worlds in which x is *A* but not *C*). Likewise, since some realizers of *B* (namely β_1s) now lead to realizers of *C* (namely γ_2s), we can no longer infer (where x would be *B* if not *A*) that, if x were non-*A*, then it would be non-*C*.[44]

The present problem also afflicts structural equations versions of the counterfactual approach to causation. To see this, let V_A be a binary variable that takes value 1 if *x* possesses *A* and value 0 otherwise. We might suppose that, if $V_A = 0$, then *x* possesses *B* instead of *A*. Let V_C be a binary variable that takes value 1 if *x* possesses *C* and value 0 otherwise. We might suppose that, if $V_C = 0$, then *x* possesses *D* instead of *C*. Suppose that *x* is in fact *A* (that is, $V_A = 1$). Then the original version of our simple model – in which the realizers of *A* straightforwardly mapped onto the realizers of *C*, and the realizers of *B* onto *D* – could be represented by the pair of structural equations $\{V_A = 1, V_C = V_A\}$, the actual solution to which is $V_A = 1$, $V_C = 1$. Structural equations analyses of causation were able to say that *x*'s possession of *A* (represented by $V_A = 1$) is a cause of *x*'s possession of *C* ($V_C = 1$), since the value of V_C counterfactually depends upon that of V_A according to the structural equations. (To see this replace the equation '$V_A = 1$' with the equation '$V_A = 0$' and note that, in the new solution, the value of V_C is changed to 0.) But in the modified version of our example (in which some of the realizers of *A* now lead to realizers of *D* rather than *C*, and some of the realizers of *B* now lead to realizers of *C*), the structural equation $V_C = V_A$ can no longer be presumed to hold,[45] since it can no longer be presumed that (as the equation implies), if $V_A = 1$, then it would be that $V_C = 1$, nor that if $V_A = 0$, then it would be that $V_C = 0$. A causal relation between $V_A = 1$ and $V_C = 1$ therefore can no longer be inferred.

One might think that, in this case, it is no bad thing that nomic regularity and counterfactual theories (including their structural equations variants) cannot provide for the causal efficacy of some *x*'s being *A* for its being *C*. After all, given that (by stipulation) *x* would be *B* if it weren't *A*, and given that exactly one realizer of each of *A* and *B* leads to a realizer of *C* (with exactly one realizer of each leading instead to a realizer of *D*), it is far from intuitively clear that *x*'s being *A* (as opposed to its being α_1, say) *is* causally relevant to its being *C*.

Suppose, however, that the high-level properties in our model had been realizable in a greater number of ways. Specifically, suppose that *A* had been realizable by α_1, ..., and α_n, *B* by β_1, ..., and β_m, *C* by γ_1, ..., and γ_k, and *D* by δ_1, ..., and δ_h, where *n* and *m* are each very large. Suppose, moreover, that the *vast majority* of the αs map deterministically onto γs and that the vast majority of the βs map deterministically onto δs (where each of the δs is incompatible with any γ). In such a case it seems rather plausible to say (where *x* is both *A* and *C* and *x* would have been *B* if non-*A*) that *x*'s being *A* was a cause of its being *C*. Nevertheless, provided that *just one* realizer of *A* leads to a realizer of *D* (rather than *C*), then being *A* is not sufficient for being *C*. And provided that *just one* of the *closest* realizers of *B* leads to a realizer of *C*, then there is no counterfactual dependence of *C* upon *A* either. Specifically it is not true that if *x* had been non-*A*, then it would have been non-*C*, since it just might still have been *C*.[46]

Many apparently genuine instances of high-level causation involve properties that are related in precisely this way. Thermodynamically irreversible (that is, entropy-increasing) processes are a case in point. Suppose, for example, that I take an ice cube out of the freezer and place it in a large glass of hot water at 12:00 noon. By 12:30pm it has melted. It seems very plausible indeed to say that the ice cube's being in hot water at 12:00 caused it to melt by 12:30. Yet there are at least *some* possible microphysical realizations of the ice cube's being in the hot water at 12:00 that (together with appropriate realizations of the rest of some isolated system of which the ice-in-hot-water is part) lead deterministically to a 12:30 state in which the ice cube has *not* melted. There are even some that lead to its *increasing* in size (with the surrounding water becoming even hotter). So being an ice cube in hot water at 12:00 certainly isn't *sufficient* for melting by 12:30 (and this is so even if we suppose the ice-in-hot-water system to be isolated). Nor does the ice cube's melting counterfactually depend upon its being in hot water. If the ice hadn't been in the hot water, then perhaps it would have been back in the freezer. Still it's not true that it wouldn't have melted. It *might* have melted, with the rest of the contents of the freezer cooling slightly as it did so.[47] That is to say, there are possible microphysical realizers of an initial ice-in-freezer system that, by the deterministic microphysical laws, result in such a course of events.

It thus appears that, unlike non-strictness due to *ceteris paribus* clauses, the present respect of non-strictness – namely, that due to their *minutiae rectus* nature – poses a genuine difficulty for the ability of high-level laws to underwrite causal relations. In the next two sections, I will explore the question of whether this apparent difficulty can be shown to be illusory. In particular, I will examine whether a case can be made that high-level laws typically are not, after all, *minutiae rectus* laws and that they therefore do not suffer from a lack of strictness in any respect required for underwriting high-level causal relations.

7. Probabilistic High-Level Laws

It was seen in Section 6 that the apparent fact that high-level laws are typically *minutiae rectus* in nature poses a problem for the accommodation of high-level causal relations in a way that their typically *ceteris paribus* nature does not. It will now be argued that the most promising line of response to this problem is to distinguish between, on the one hand, *minutiae rectus* laws (which are non-strict, and non-exceptionless) and, on the other hand, strict and exceptionless probabilistic laws.[48] A case can be made that multiple realizability should not (after all) lead us to conclude that high-level laws are typically *minutiae rectus* laws, and as such liable to admit of exceptions, but rather that they are typically probabilistic.

The case is perhaps at its clearest when we consider the Second Law of Thermodynamics. It was observed in Section 6 that some possible microphysical realizers of an initial non-equilibrium macrostate of an isolated system (perhaps one containing an ice cube in a glass of hot water) may lead, by the micro-dynamic laws (which, for now, are being assumed to be deterministic), to later microstates which realize macrostates still further from equilibrium (where, for example, the ice cube has grown and the surrounding water has become hotter). Though the measure of such microstates – the volume they occupy in the phase space associated with the initial macrostate – is very small indeed, they nevertheless exist. It is for this reason that it appears that the Second Law is a *minutiae rectus* law, which admits of exceptions when one of these unlikely microphysical realizers of the initial non-equilibrium state is instantiated.

But the Second Law appears to be a *minutiae rectus* law, admitting of exceptions, only when construed *deterministically*. That it shouldn't be so construed, however, is made plausible by the fact that statistical mechanics – which relates the macro-properties of systems to the micro-properties of the particles that compose them – yields only a *probabilistic* version of the Second Law: according to statistical mechanics it is 'merely' overwhelmingly probable (and not certain) that an isolated system in initial disequilibrium will evolve in the direction of increasing entropy.

On the statistical mechanical (SM) picture, cases in which the initial macrostate of a system are realized in one of those rare ways that leads to sustained entropy decrease are not construed as exceptions to a deterministic Second Law. On the contrary, statistical mechanics construes the Second Law probabilistically and, like any probabilistic law, there are cases in which the properties that it relates (in this case earlier lower entropy and later higher entropy) fail to be co-instantiated, but where this failure is not construed as an exception to the law, because the law merely attaches a probability to such co-instantiation.

Some caution is needed here, however. The fact that a law is probabilistic isn't *in principle* incompatible with its being a *minutiae rectus* law. Suppose that a high-level law entails a certain probability p for an object's having the high-level property C conditional upon its having the high-level property A: that is, $P(C|A) = p$. It might nevertheless be the case that the microphysical plus bridge laws entail that, for some possible microphysical realizer α_i of A, $P(C|\alpha_i) \neq p$. This will be the case if, according to the microphysical laws, there is a probability p' ($\neq p$) of some or other microphysical realizer of C (that is, one or other of the γs) being instantiated *given* that α_i is instantiated. In such a case, one might reason that the high-level law holds only *minutiae rectus* and that the case where A is realized by α_i constitutes an exception to it. The idea would be that the high-level law holds only when A is realized by one of those microphysical properties, α_j, such that the microphysical laws plus bridge laws entail that $P(C|\alpha_j) = p$. (A deterministic *minutiae rectus* law of the form $A \rightarrow C$ can be construed as the special case where $P(C|A) = 1$ but, for some α_i, $P(C|\alpha_i) < 1$. Where the micro-dynamics are themselves deterministic, the latter implies that $P(C|\alpha_i) = 0$.)

But to construe the probabilistic version of the Second Law entailed by statistical mechanics as being, in this way, a *minutiae rectus* law would be to misconstrue it. Such a law entails that the probability of sustained entropy increase for an isolated system in initial disequilibrium is close to, but not equal to 1. In the classical case, however, the microphysical laws entail, for *each* possible micro-realizer of an initial non-equilibrium state of an isolated system (that is, for each point in the associated phase space), a probability 1 or 0 of sustained entropy increase (the probability is 1 for measure close, but not equal, to 1 of the possible micro-realizers). It is incorrect to construe a probabilistic Second Law as a *minutiae rectus* law because (at least in the classical case) there is no possible case in which the *minutiae* are *rectus*: every possible initial microstate of an isolated system yields (in conjunction with the deterministic microphysical and bridge laws) a trivial (1 or 0) probability of entropy increase. This diverges from the non-trivial probability entailed by a probabilistic Second Law. A probabilistic Second Law therefore does not have the status of a *minutiae rectus* law that simply states the probability of entropy increase that the microphysical laws associate with 'typical' realizers of an initial non-equilibrium state (since this would be 1).

On the contrary, the SM probabilities encoded in a probabilistic Second Law are derived from a probabilistic *averaging* over possible microstates (cp. Sober 1999, p. 555). Specifically (at least to a first approximation) the SM probabilities are a weighted average of the probabilities (which in the classical case are all 1s and 0s) with which the possible micro-realizers of the initial thermodynamic state of an isolated system lead to entropy increase, where the weighting is determined by a probability distribution over the set of micro-realizers compatible with that initial thermodynamic state (that is, over the associated region of phase space). The probability distribution in question is usually taken to be one that is uniform (on the standard Lebesgue measure). Consequently, the very high SM probability of entropy increase for a non-equilibrium isolated system is a result of the fact that a large volume of the associated phase space is occupied by microstates that (according to the microphysical laws) are on entropy-increasing trajectories toward the future.

If we drop the assumption that the fundamental dynamics are classical, and assume instead that they are the quantum mechanical dynamics, then probabilities enter into the fundamental dynamics themselves. Still, as Albert (2000, pp. 153-4; 2012, pp. 38-9) has argued, on most interpretations of quantum mechanics, probabilistic events do not occur in the right places to underwrite the statistical mechanical probabilities encoded in a probabilistic Second Law.[49] Consequently, the orthodox view of quantum statistical mechanics (described by Albert 2000 pp. 132-3, 154; 2012, pp. 38-9) is, to a first approximation, that the SM probabilities are a weighted average of the probabilities that the quantum dynamics assign to the various possible initial quantum states of a system leading to entropy increase, where the weighting is determined by a probability distribution over the set of quantum states compatible with the system's initial thermodynamic state.

Consequently, by the lights of standard statistical mechanics (whether classical or quantum), the correct interpretation of a case in which the initial macrostate of a system is realized in a way that leads to entropy decrease is not as a case in which there is an exception to the (probabilistic) Second Law, but as a case in which entropy decreases in accordance with the probabilistic Second Law, but as a matter of the low probability for entropy decrease entailed by the Law. The low probability entailed by the Law results from the fact that the distribution over initial microstates entails a low probability for the initial macrostate of the system being realized in a way that leads to entropy decrease.

Once it is understood that the Second Law should be construed as a probabilistic law that is not a *minutiae rectus* law, it becomes clear how it is able to underwrite causal relations. Take some thermodynamically irreversible process, such as our ice-melting-in-hot-water sequence, that is covered by the Law. While a probabilistic Second Law does not imply that the thermodynamic state of the ice-in-hot-water system at 12:00 noon was *sufficient* for a 12:30pm state of higher entropy (in which the ice cube has melted), it does imply that the former state was sufficient for a very high (SM) probability of the latter.[50]

Now consider possible macrostates of an ice-in-freezer system that might have been instantiated at 12:00 noon instead of the actual ice-in-hot-water system. For each of the *most likely* such macrostates, a probabilistic Second Law implies that there is a very low (SM) probability of its evolving into a 12:30pm macrostate in which the ice has melted and its surroundings have cooled (this would involve entropy decrease). It consequently implies that the SM probability of the ice's melting is lower conditional upon its being returned to the freezer than it is conditional upon its being placed in the hot water.

Similarly, while a probabilistic Second Law does not support the (false) counterfactual 'If the ice cube hadn't been in the hot water at 12:00 noon (but had rather been in the freezer), then it wouldn't have melted by 12:30pm' (since the ice cube still *might* have melted), it *does* support the (true) counterfactual 'If the ice cube hadn't been in the hot water at 12:00 noon (but had rather been in the freezer), then the (SM) probability of its melting would have been much lower than it actually was'. It supports this counterfactual because, in each of the *closest* worlds in which the ice is in the freezer at 12:00 noon, the macrostate of the ice-in-freezer system leads to a 12:30pm macrostate in which the ice has melted with a much lower SM probability than does the macrostate of the actual 12:00 noon ice-in-hot-water system.

The fact that the ice cube's being in the hot water raises the probability of its melting (in both the conditional probability and counterfactual senses described in the previous two paragraphs) is just the sort of fact to which probabilistic analyses of causation appeal.[51] The sequence, it might thus be claimed, is a paradigm case of probabilistic causation. No wonder straightforward regularity and counterfactual theories (including their structural equations variants) were unable to accommodate it: these are theories of deterministic causation being (mis)applied to a case of probabilistic causation!

A probabilistic Second Law appears to support a generalization along the lines of '*Ceteris paribus*, normal ice cubes placed in normal glasses of hot water quickly melt with very high (SM) probability'.[52] Why the need for the hedge (when the Second Law, which supports the generalization, itself appears to contain no hedge)? The reason is that the Second Law properly applies only to thermodynamically isolated systems. But ice-cube-in-hot-water systems are rarely isolated. Such a system may *approximate* an isolated one provided that mass-energy doesn't cross its boundaries in great enough quantities or in such a way as to significantly interfere with it. Cases in which this assumption is clearly violated include cases in which someone fishes the ice cube out of the water to put it back in the deep freeze, in which someone buries the glass in snow, and so forth. These are the sorts of scenarios in which the *ceteris paribus* clause is violated because the system no longer approximates an isolated one and so the Second Law no longer approximately applies.

The fact that typical ice-in-hot-water systems at best approximate isolated systems is one reason why, even where there is no egregious violation of the requirements for an isolated system (so that the *ceteris paribus* condition is fulfilled), we cannot associate the generalization with any precise SM probability. A second reason is that, even if such a system were genuinely isolated, a precise SM probability would be derivable only from a much more precise specification of its macrostate than is given by the predicate '... is a normal glass of hot water containing a normal ice cube'.

Nevertheless, a probabilistic Second Law implies things about the *range* of values within which the probability falls. The reason is that for *each* precisely-specified macrostate compossible with a system's satisfying the predicate '... is a normal glass of hot water containing a normal ice cube', the SM probability of its evolving (within the relevant time scale, and in the absence of egregious violation of the requirements for an ideal isolated system) into one of the macrostates compossible with its satisfying the predicate '... is a glass of tepid water containing no ice' is very high. In particular, it is (much) higher than the range of values associated with those precisely-specified macrostates compossible with the satisfaction of the predicate '... is a normal ice cube in a freezer'.[53]

Now it *may* be that the above reasoning can be generalized to a majority of other cases of apparent high-level causation. In particular, Albert (2000, 2012) and Loewer (2007a, 2008, 2012a, 2012b) have argued that – just as it underwrites a probabilistic generalization about ice cubes tending to melt in hot water – a probabilistic Second Law (or rather – as shall be seen in the next section – a set of axioms from which such a law derives) underwrites probabilistic versions of many of the generalizations of the high-level sciences and of folk theory. One reason for thinking that it does so is that at least very many high-level laws concern thermodynamically irreversible (i.e. entropy-increasing) processes. This is true, for example, of geological laws concerning the erosion of river banks and mountain ranges (see Elga 2001, p. 322); of meteorological laws concerning hurricane formation and the evolution of pressure systems (see Loewer 2008, p. 159); of biological laws concerning ageing (Albert 2000, p. 22), inheritance, and the workings of neurons; of chemical laws governing chemical reactions; of macrophysical laws concerning the slowing of objects due to friction (see Albert 2000, pp. 28-9); of gambling laws concerning coin flips (see Loewer 2007a, p. 306; 2008, p. 159; Albert 2012, p. 19); and so on. It appears that any probabilistic generalizations that the Second Law does underwrite ought to inherit its abilities to entail the existence of probability-raising relations between high-level properties (at least when the *ceteris paribus* clauses, if any, associated with those generalizations are satisified), and hence to support relations of high-level causation.

But even if a probabilistic Second Law (or a set of axioms that entail it) does not underwrite probabilistic versions of many of the generalizations of the high-level sciences,[54] it may nevertheless be the case that those generalizations can be understood by *analogy* to the Second Law: namely as probabilistic laws that are not *minutiae rectus* laws, and that are thus strict in the respect needed to underwrite causal relations. Indeed, a number of philosophers who are skeptical that statistical mechanics underwrites the generalizations of the high-level sciences nevertheless wish to make room for probabilistic high-level laws, in some cases taking them to derive – analogously to probabilistic versions of the principles of thermodynamics – from lawful probability distributions over underlying state spaces. Such views, which will be examined in greater detail in the next section, are to be found in Callender and Cohen (2009, 2010), Glynn (2010), and Ismael (2009, 2011). But first, there is a question to be answered about whether the Second Law and any probabilistic generalizations that it *does* support are really laws at all, strict or otherwise.

8. Are They Really Laws?

In the previous section it was argued that the Second Law and any other high-level generalizations that it underwrites are probabilistic laws that are not *minutiae rectus* laws. It was argued that, as such, they are able to support probability-raising relations (between high-level properties) of the sort appealed to by probabilistic analyses of causation.

Yet a probabilistic Second Law is not derivable from the fundamental dynamic (plus bridge) laws alone (as per the simple model of Section 4). As suggested in the previous section, a probability distribution (which is not itself given by the fundamental dynamic or bridge laws) over initial microstates is essential to its derivation. In the present section, it will be examined whether this undermines the status of the Second Law and the generalizations that it supports as genuine laws. One specific worry is that, since it is often supposed that objective chances derive only from genuine laws (see Schaffer 2003, pp. 36-7; 2007, p. 126; and Lewis 1994, p. 480), then unless the Second Law turns out to be a genuine law, the SM probabilities will not count as the sort of objective chances to which probabilistic analyses of causation appeal.

The non-derivability of a probabilistic Second Law from the microphysical laws is, of course, most vivid under the assumption of micro-determinism (an assumption that is made in classical statistical mechanics). Evidently, the non-trivial SM probabilities encoded in a probabilistic Second Law cannot be derived from deterministic microphysical laws (together with deterministic bridge laws) alone. As noted in the previous section the standard view is that, even if we relax the assumption of microphysical determinism and assume the micro-dynamics to be the quantum mechanical dynamics, a probability distribution over initial microstates is still essential to the derivation of the SM probabilities.

The non-derivability of a probabilistic Second Law from the fundamental dynamic laws alone is made clear when we consider a well-known attempt, described by Albert (2000, Chs. 3-4), to axiomatize statistical mechanics. Albert suggests that the SM probabilities can be derived from the following:[55]

(FD) the fundamental dynamical laws;

(PH) a proposition characterizing the initial conditions of the universe as constituting a special low-entropy state; and

(SP) a uniform probability distribution (on the standard Lebesgue measure) over the regions of microphysical phase space associated with that low-entropy state.[56]

I follow Albert (2000, p. 96; 2012, p. 20) in calling the second and third of these postulates the 'Past Hypothesis' and the 'Statistical Postulate' respectively. Albert and Loewer dub the conjunction FD & PH & SP 'the Mentaculus'.[57] PH and SP are essential to the derivation of the SM probabilities, but they are not derivable from the fundamental dynamic laws. Nor are they aptly construed as bridge principles.

On this formulation, the probability distribution over microstates needed to generate the SM probabilities is given by SP. Note, however, that on this formulation, a probability distribution isn't simply applied to the set of initial microstates compossible with the initial macrostate of each isolated system, but rather only to the set of those possible initial microstates of *the universe as a whole* that are compossible with the universe's low-entropy initial macrostate (as specified by PH).[58] Briefly, the reason for this is that applying a uniform probability distribution to the set of microstates compossible with some non-equilibrium macrostate, *S*, obtaining later than the beginning of the universe is liable to make it probable, not only that *S* will be followed by a higher entropy state, but also that *S* was *preceded* by a state of higher entropy. This ultimately leads to inconsistency with the Second Law.[59]

The argument that the SM probabilities are derivable from the Mentaculus goes roughly as follows. Consider the region of microphysical phase space associated with the low-entropy initial state of the universe implied by PH. Relative to the total volume of that region, the volume taken up by microstates that lead (by FD) to fairly sustained entropy increase until equilibrium is reached (and to the universe staying at or close to equilibrium thereafter) is extremely high. Consequently, the uniform probability distribution (given by SP) over the entire region yields an extremely high probability of the universe following such a path. When it comes to (approximately) isolated subsystems of the universe the idea is that, since a system's becoming approximately isolated is not itself correlated with its initial microstate being entropy-decreasing, it is extremely likely that any such subsystem that is in initial disequilibrium will increase in entropy over time (see Loewer 2007a, p. 302, 2012a, pp. 124-5; 2012b, p. 17; and Albert 2000, pp. 81-5).

Given that (a probabilistic version of) the Second Law isn't entailed by the fundamental dynamics alone, but (on this formulation) only by these in conjunction with PH and SP, one might wonder whether it, and any other probabilistic generalizations that are underwritten by it (or that are underwritten directly by the Mentaculus), are genuine, chance-entailing laws. Clearly, if the Second Law is a genuine law, the *source* of its lawfulness can't be its derivability from the fundamental dynamic laws. It must gain its lawfulness from some other source. A recently-popular view (developed by Loewer 2001, 2007a, 2012a; cp. also Callender and Cohen 2009, 2010; Dunn 2011; and Glynn 2010) is that its lawfulness derives from exactly the same source as that of the microphysical laws themselves: namely from the fact that it, like them, is entailed by a system that strikes the best balance between the theoretical virtues. Specifically, proponents of this view have appealed to Lewis's (1994) Best System Analysis (BSA) of laws, or variants thereof. It is argued that a probabilistic Second Law comes out as a genuine, chance-entailing, law of worlds like ours (including both nearby micro-deterministic and micro-probabilistic worlds) on such an analysis.[60]

According to the BSA, the theoretical virtuousness of an axiom system pertaining to what goes on in the course of the world's history counts directly towards its theorems earning their status as genuine laws, with the chance-entailing properties of genuine laws (Lewis 1994, pp. 478-9). Specifically, according to the BSA, the laws are the theorems of that deductive system which systematizes the entire history of the world in a way that achieves the best balance between the theoretical virtues of simplicity, strength, and fit.[61]

According to Lewis, a system is *strong* (or *informative*) to the extent that it says "either what will happen or what the chances will be when situations of a certain kind arise" (ibid.). There is reason to think that adding axioms to a system that already entails the fundamental dynamic laws, so that it additionally entails a probabilistic Second Law and other probabilistic high-level generalizations increases the informativeness of the system in question. The reason is that such high-level generalizations tell us what the chances will be when situations arise that are of a kind concerning which the fundamental dynamic laws are silent. The kinds of situation in question are, of course, situations of *high-level* kinds.

Take, for example, situations of the high-level kind *being an ice cube in a glass of hot water*. The fundamental dynamic laws (even together with the bridge laws) tell us nothing about what will happen or what the chances will be when situations of this kind arise. They tell us only about situations of the microphysical kinds α_1, ..., and α_n, where the latter are possible realizers of the ice-in-hot-water system. Moreover, because of its multiple realizability, the fact that a system is of the high-level kind *an ice cube in a glass of hot water* does not (together with the bridge laws) entail which of the microphysical kinds α_1, ..., or α_n it belongs to.

By contrast, a probabilistic Second Law *does* provide us with information about what the chances will be when situations of this high-level kind arise (it tells us, for instance, that the chance of the system's coming to be of the high-level kind *a glass of tepid water with no ice* within the half-hour is very high). Consequently a system that entails a probabilistic Second Law will be more informative than one that entails the fundamental dynamic laws alone.

In general, high-level laws and fundamental dynamic laws entail chance distributions conditional upon different sorts of proposition. High-level laws, such as a probabilistic Second Law, entail chance distributions conditional upon propositions about a system's macrostate $P(\cdot|M)$, while the fundamental dynamic laws entail distributions conditional upon propositions specifying a system's microstate $P(\cdot|m)$.[62] There is no conflict between divergent conditional chance distributions with different conditions. Indeed, an axiom system that entails both conditional distributions is more informative than one that entails only one (and leaves the other undefined).

It is a good question exactly how high a price in simplicity this greater strength or informativeness is bought at. As we have seen, in order to derive a probabilistic version of the Second Law, we must add PH and SP to the fundamental dynamic laws, FD. The questions of exactly how much simplicity the addition of PH and SP to our set of basic axioms costs, and of whether the strength gained is worth the price, are discussed (and disputed) by Loewer (2001, pp. 617-18), Schaffer (2007, pp. 130-1), Hoefer (2007, p. 560), and Glynn (2010, pp. 59-63).[63] I'm inclined to agree with Hoefer (*op cit.*) that the former question is not readily answered. The trouble is that there's no obviously most reasonable simplicity metric to apply (cp. Lewis 1994, p. 479). This, together with the fact that it's not obvious how to trade off simplicity against informativeness, makes it difficult to answer the latter question.

Still, I think it's clear that, if we are to make room for high-level causation, then we had better adopt a simplicity metric and an exchange rate between simplicity and informativeness such that a system (such as the Mentaculus) that entails probabilistic high-level laws comes out *best*. Moreover, in light of the overwhelming plausibility of examples of high-level causation like those described at the outset of this essay, I think that we *ought* to make room for high-level causation.

A potential objection to the view that the Mentaculus is the Best System for our world (and others like it) is that Lewis (1983, pp. 367-8) restricts candidates for best systemhood to those systems whose axioms refer only to perfectly natural properties. His reason for doing so is that the simplicity of a system is relative to the vocabulary in which it is expressed and that, by employing very unnatural predicates, we might make a strong system very (syntactically) simple indeed (ibid., p. 367). So without some restriction on the language in which a system is expressed, the desideratum of simplicity loses its bite. But, as Schaffer (2007, p. 130) points out, the Mentaculus contains predicates like 'low entropy' that correspond to properties that are not perfectly natural, and so doesn't even seem to be a candidate Best System.

I think a reasonable response involves a slight modification of the BSA. Observe that, as Lewis recognizes (1983, p. 368), naturalness admits of degrees. We may take naturalness of the predicates that it employs to be a theoretical virtue, to be weighed alongside the simplicity and strength of a system. If an axiom system is able to achieve great simplicity and strength by employing a not-too-unnatural predicate like 'low entropy' – as the Mentaculus does – then it is a plausible best system.

If the Mentaculus comes out as best, then the BSA will accommodate a probabilistic version of the Second Law as a genuine law.[64] On the BSA (which is proposed as an analysis of chance as well as of laws) it immediately follows that the SM probabilities are genuine, objective chances.[65] For, according to the BSA, "the chances are what the probabilistic laws of the best system say they are" (Lewis 1994, p. 480). But, if a probabilistic Second Law is a genuine law, and the SM probabilities are genuine objective chances, then a probabilistic Second Law is capable of underwriting exactly the sort of objective chance dependencies (that is, objective chance-raising relationships) that are appealed to by probabilistic analyses of causation. Since these objective chance dependencies obtain between high-level property exemplifications (such as *being an ice cube in a glass of hot water* and *melting*), a probabilistic Second Law underwrites probabilistic high-level causation.

As noted at the end of the previous section, it is controversial whether statistical mechanics, which Albert and Loewer claim to be captured by the Mentaculus, entails probabilistic versions of a great many of the generalizations of the special sciences. If Albert and Loewer are correct that it does, then, as theorems of the Best System, those probabilistic generalizations will count as probabilistic laws on the BSA, and will themselves be able to underwrite the objective chance dependencies in terms of which probabilistic causation is standardly analyzed.

But even if that's not the case, there are fallback positions available to the defender of high-level causation. For example, Callender and Cohen (2009, 2010) propose a modified BSA – a 'Better Best Systems Account (BBSA)' as they call it – which ensures that the simple, informative probabilistic generalizations of the special sciences come out as genuine laws and – applying their BBSA to chances as well as laws – that the probabilities that they entail come out as objective chances. Briefly, their proposal draws upon Lewis's observation that a system's simplicity depends upon the vocabulary in which it is expressed. But rather than following Lewis in restricting the systems under consideration to those whose axioms contain only perfectly natural kind predicates, their idea is that best systemhood should be taken to be *relative* to a set of basic kinds *K* (or predicates P_K).[66] Relative to different sets of kinds, different axiom systems strike the best balance between simplicity, strength, and fit. A generalization is a law relative to *K* just in case it is a theorem of the Best System relative to *K*.

Callender and Cohen's view is particularly conducive to counting special science generalizations as laws. In particular, on their view, the generalizations of a special science (such as biology, economics, or geology) count as laws of that science if they are theorems of the best system relative to the science's proprietary kinds or predicates (e.g. the biological, economic, or geological kinds).[67] A reasonable case can be made that the theorems of best systems relative to special science vocabularies will typically be *probabilistic* laws and not *minutiae rectus* laws.[68] Specifically, suppose that an axiom system entails a law relating the high-level kinds *A* and *C*. Then, other things being equal, that axiom system will fit history well if the law entails a probability $P(C|A)$ that closely matches the frequency with which instances of *A* are micro-realized in ways that lead deterministically to micro-realizers of *C*. Alternatively, if the microphysical laws are probabilistic, and the micro-realizers of *A* only lead probabilistically to micro-realizers of *C*, then the probability $P(C|A)$ will fit well if it is equal to a weighted average of the probabilities with which the various micro-realizers of *A* lead to micro-realizers of *C*, where the weighting depends upon the frequency with which *A* is realized in those various microphysical ways (cp. Sober 1999, p. 555). Indeed, Callender and Cohen (2010, pp. 437-8) and Callender (2011, p. 103, 112) suggest that the best axiomatizations for various special sciences will include probability distributions over underlying state-spaces (these need not be phase spaces however), where those distributions closely match the frequencies with which higher-level properties are realized in the spaces in question. On their view, the probabilistic theorems generated by the resulting axioms will be probabilistic laws of the sciences in question.

Similarly, although Ismael (2011, p. 433; cp. Ismael 2009, p. 106) expresses skepticism about the derivability of the probabilistic generalizations of the high-level sciences from a distribution over the initial conditions of the universe (a la Albert and Loewer), she nevertheless argues (Ismael 2009, 2011) that a probability measure over underlying state space is essential to their derivation (cp. Glynn 2010, pp. 60-2). Which measure is appropriate depends upon the relative frequencies with which macrostates are realized in various microphysical ways (Ismael 2009, p. 96; 2011, p. 433, 438; cp. Glynn 2010, p. 61): a distribution will be preferable if it closely matches those frequencies. Where the distribution over the micro-realizers of some high-level property *A* closely matches the frequency with which instances of *A* are realized in those various micro-physical ways, then, once again, a generalization will be entailed that yields a probability $P(C|A)$ that it is a weighted average of the probabilities with which the various micro-realizers of *A* lead to micro-realizers of *C*, where the weighting is close to the frequency with which *A* is realized in those various microphysical ways. While Ismael doesn't commit to (and has "reservations about" – 2011, p. 432) the Best Systems approach to laws and objective probability (cf. Glynn 2010, pp. 59-61), she nevertheless takes a probability distribution over underlying state space to be part of the "objective content" (Ismael 2009, p. 91) of any theoretical package from which probabilistic high-level generalizations can be derived, and takes such high-level generalizations to be laws and the probabilities that they entail to be objective.

In the previous section, I showed that it is vital that high-level generalizations be construed as probabilistic laws, rather than as *minutiae rectus* laws, if they are to be construed as strict in that respect (namely exceptionlessness once any *ceteris paribus* clause is satisfied) that is critical to their ability to underwrite causal relations. In this section, I have described a range of philosophically attractive accounts of lawhood that seem able to treat high-level generalizations as genuine probabilistic laws (and not as *minutiae rectus* laws, or as non-lawful generalizations). I will not attempt to adjudicate between these various accounts but, if one of them is along the right lines, then I think there is a strong case that there are indeed high-level laws that are strict in those respects needed to underwrite high-level causation.

The argument can also be turned around. As I argued in Sections 6 and 7, a law's holding only *minutiae rectus* seems to fatally undermine its ability to support the sort of relations (of nomic sufficiency, or of counterfactual or probabilistic dependence) in terms of which causation is analyzed on the most attractive philosophical theories of causation. Since it is hugely implausibility to deny the existence of high-level causal relations, this serves as an additional reason to endorse an account of lawhood which – like those described in this section – implies the existence of genuine probabilistic high-level laws that are not *minutiae rectus* laws.

9. Conclusion

It has been argued that the 'no strict laws' problem for high-level causation can be overcome. In that respect in which it is plausible that high-level laws are genuinely non-strict – namely, non-strictness in respect of containing *ceteris paribus* clauses – strictness isn't required in order to underwrite causal relations. On the other hand, I have argued that there is a possible respect of non-strictness – namely non-strictness in respect of being a *minutiae rectus* law – that has often not been properly distinguished and which appears to fatally undermine a law's ability to underwrite causal relations. *Prima facie*, it seems that many high-level laws are *minutiae rectus* laws. Yet I have argued that a strong case can be made that high-level laws are in fact better interpreted as *probabilistic* laws rather than as *minutiae rectus* laws. As such, high-level laws are able to support the sort of objective chance dependencies (between high-level properties) to which probabilistic analyses of causation appeal. I have argued, moreover, that there is a range of independently promising philosophical accounts of lawhood that can support this probabilistic interpretation of high-level laws.

Acknowledgements

For detailed comments on earlier drafts, I would like to thank Christopher Hitchcock, Thomas Kroedel, Wolfgang Spohn, and Joel Velasco. I would also like to thank audiences to presentations of various forerunners of this paper: namely those at the 2010 *Causation Across Levels Workshop* at the IHPST in Paris, a 2010 meeting of the *Causality and Probability Research Colloquium* at the University of Konstanz, the 2012 workshop *The Objective Reality of Causality* also at Konstanz, and a 2012 presentation at the University of St. Andrews. Funding for this research was provided by the Deutsche Forschungsgemeinschaft (SP279/15-1) and by the J. S. McDonnell Causal Learning Collaborative.

[1] By contrast, Russell (1913) famously argued that *fundamental physics* neither seeks nor discovers causes. Russell's claim is highly controversial, and has been challenged by Suppes (1970, pp. 5-6) and Hitchcock (2007a, pp. 55-6). For further discussion, see Field (2003), Ney (2009), and the papers collected in Price and Corry (2007).

[2] Davidson focused specifically on the problem posed for *mental* causation by the lack of strict *psychological* and *psycho-physical* laws.

[3] Though see Davidson (1967) for some relevant remarks.

[4] Salmon (1994, 1997) and Dowe (2000) have both attempted to characterize causal processes without the use of counterfactuals (and without appeal to nomic sufficiency or probabilistic dependence). See Hitchcock (1995, 2009) and Choi (2002) for compelling objections.

[5] By contrast to the generalizations of the high-level sciences, the laws of fundamental physics are generally supposed by philosophers to be strict. (Though, for a dissenting view, see Cartwright 1983, 1999.) Consequently – and somewhat ironically in light of the Russellian claim mentioned in note 1 – we find ourselves in the position of having greater trouble (in this regard at least) in providing a theoretical accommodation of high-level than of microphysical causation. In this paper, I shall not discuss the question of whether there is microphysical causation. My treatment of high-level causation is independent of this.

[6] See, for example, Malcolm (1968) and Kim (1989a).

[7] See, for example, Kim (1998, pp. 46-7, 77-87, 112-18) and Block (2003, esp. pp. 138-40).

[8] In fact, there is a range of promising responses to the exclusion problem already on the market, none of which constitutes a solution to the no strict laws problem. See, for example, Yablo (1992), Crane (1995), Horgan (1997), Bennett (2003), Loewer (2007b), and Kroedel (2008).

[9] See, for example, Block (1990, pp. 139-40) and Crane (1995, p. 223), as well as the critics of Davidson cited in the next footnote.

[10] Arguments to this effect are found, for example, in Stoutland (1980), Honderich (1982), Sosa (1984), Kim (1989b, 1993), and McLaughlin (1989, 1993).

[11] This formulation presupposes the view that property exemplifications (rather than, or as well as, coarsely individuated events) can enter into causal relations. Those not comfortable with this supposition could substitute, here and in what follows, occurrences of "x's being X was a cause of y's being Y" with "x was a cause of y in virtue of x's being X and y's being Y" (with x and y interpreted as coarse-grained events). Various other locutions would also have to be reformulated. But none of this would affect the thrust of the argument, which would then concern whether coarsely individuated events enter into causal relations *in virtue* of their high-level properties, or whether those properties are mere epiphenomena.

[12] The notion of a strict probabilistic law will be discussed in detail in Section 7.

[13] Here and throughout I speak interchangeably of 'causal efficacy' and 'causal relevance'. Jackson and Pettit (1988, pp. 391-7; 1990) have challenged this conflation by arguing that a high-level property, even if it lacks efficacy, may be causally relevant by *programming for* an efficacious lower-level property (cp. Block 1990, pp. 162-3). This relevance-in-virtue-of-programming relation neither requires strict high-level laws nor is the target of standard philosophical analyses of causation. And nor is it my concern here. My concern is rather to investigate the prospects for a philosophical vindication of high-level causation in the most full-blooded sense (call it 'efficacy' if you wish) rather than in some secondary or inferior sense, such as mere 'programming'. In my view this is what intuition and scientific practice demand.

[14] Until Section 7, I shall focus specifically upon the ability of *regularity* and *counterfactual* theories to accommodate high-level causation. I will disregard *probabilistic* theories (until Section 7) because I will (until then) make the simplifying assumption of determinism. And I won't explicitly discuss *process* theories because, to the extent that the notion of a causal process is ultimately to be cashed out in terms of nomic regularities or counterfactual (or probabilistic) dependencies, establishing the existence of such relations among high-level properties ought to reassure us that there exist high-level processes too.

[15] Standard examples of special sciences include biology, chemistry, geology, meteorology, economics, and psychology. Thermodynamics may or may not count as a *special* science, depending upon exactly how one characterizes the latter (cp. Dunn 2011, pp. 86-7). The same may be true for certain other macrophysical sciences. I certainly don't suppose that there is always and everywhere a sharp distinction between macrophysics and the special sciences.

[16] This formulation is similar to those given by Kim (1984, p. 262; 1987, p. 316; 2002, p. 9) and Yablo (1992, p. 254). It is adequate only on the assumption that the set P of basic physical properties is closed under conjunction, so that the possession of any combination of basic physical properties is itself a basic physical property (see Kim 2002, pp. 8-9). Otherwise, the correct formulation is: necessarily, for each property h_i in H, if any object x has h_i, then there is some $P' \subseteq P$ ($P' \neq \emptyset$) such that x has all the properties in P' and, necessarily, anything that has all the properties in P' has h_i.

[17] The thesis that the properties in H are multiply realizable by those in P, as captured by (**M**) + (**S**), is tenable only on the assumption that the set P is *not* closed under *disjunction*. After all, the disjunction of each of the *p*s that necessitate some given h_i is (by (**S**)) necessitated by h_i (as well as

necessitating h_j). So if this disjunction were itself a member of $\boldsymbol{P}$, (**M**) would be violated. But the assumption of $\boldsymbol{P}$'s non-closure under disjunction is plausible: as Yablo (1992, p. 255n) observes, if one were to reject it, one would be forced to give up the very reasonable assumption that the sharing of physical properties makes for physical similarity (see also Kim 1992, p. 13; and Fodor 1974, pp. 103-4, 108-10).

[18] I think the qualification is appropriate since (for one thing) it is at present very unclear how the General Theory of Relativity – which serves as the basis for contemporary cosmology and astrophysics – is related to quantum mechanics, which is our best microphysical theory.

[19] In classical statistical mechanics, the microstate of a system is characterized by its point in phase space, as described in the main text. In quantum statistical mechanics, by contrast, the microstate of a system is its quantum state. In that case, the set of possible thermodynamic states of a system bear the relation (**M**) + (**S**) to the set of its possible quantum states (cp. Albert 2000, pp. 132-3, 154).

[20] This is so however one resolves the 'chicken-and-egg' dilemma of whether the notion of a fundamental law is to be analyzed in terms of that of a perfectly natural property or vice versa.

[21] The model resembles one described by Fodor (1974), though Fodor (ibid. esp. pp. 111-12) ultimately rejects the notion that strict high-level laws can be derived from fundamental physical laws (see also Fodor 1997).

[22] That these generalizations must be at least nomically necessary follows directly from the assumption that A, B, C, and D stand (respectively) in the relation defined by (**M**) + (**S**) to the sets $\{\alpha_1, \alpha_2\}$, $\{\beta_1, \beta_2\}$, $\{\gamma_1, \gamma_2\}$, and $\{\delta_1, \delta_2\}$. Note that, for economy of expression, I don't distinguish sharply between laws and statements of law here and in what follows. I trust that this won't result in any confusion.

[23] Microphysical determinism is assumed for simplicity.

[24] This is an assumption to which plausible empiricist analyses of lawhood, such as Lewis's Best System Analysis (BSA), are committed (see Lewis 1983, pp. 367-8; 1994, p. 478). (The BSA will be outlined in Section 8.) Armstrong (1983, esp. pp. 44-5, 145), a realist about laws, also allows that all logical consequences of laws count as *derived* laws. Fodor (1974, p. 109) denies that the property of lawhood is preserved under entailment (see also Kim 1992; and Fodor 1997), though Sober (1999, pp. 552-4) criticizes Fodor for this. (Sober's critique – esp. ibid. p. 554 – also applies to Kim.)

[25] To keep the model simple, I have supposed each of the relevant laws to relate the possession of two or more properties by a single object or system, x.

[26] Schurz (2002) distinguishes between 'normic' *ceteris paribus* clauses, which state that the associated generalization holds under *normal* conditions, and 'theoretically definite (exclusive)' *ceteris paribus* clauses, which state that the associated generalization holds under certain theoretically well-defined *ideal* conditions (cp. Cartwright 1983, p. 45). Schurz (ibid.) makes a related distinction between *comparative ceteris paribus* clauses ('other relevant factors held constant') and *exclusive ceteris paribus* clauses ('other relevant factors absent'). The details of these distinctions needn't concern us here.

I believe that Schurz, in a talk entitled 'Comparative Versus Exclusive Ceteris Paribus Laws: Content and Testability', given at the 2011 British Society for the Philosophy of Science Annual Conference, referred to laws with *exclusive ceteris paribus* clauses as '*ceteris rectus* laws'. This notion of a *ceteris rectus* law differs significantly from that of a *minutiae rectus* law discussed in this paper. As shall be explained in the next section, the latter is not a type of (or a variant on the notion of a) *ceteris paribus* law.

[27] Compare Schurz (2002, pp. 366-70). What about macrophysical theories like thermodynamics? Do these furnish examples of 'comprehensive closed theories' that yield non-*ceteris paribus* high-level laws? This is a possibility that will be explored in the next section. In the present section, I will focus upon examples of generalizations drawn from high-level scientific theories that are clearly not comprehensive and closed.

[28] Compare Earman and Roberts (1999, pp. 462-3).

[29] This was among the reasons that Russell (*op. cit.*) had for doubting that the relations that physics discovers are causal (see note 1 above).

[30] I have been supposing that, where *A* and *C* are ordinary, locally-instantiated high-level properties, their micro-realizers (α_1, α_2, etc., and γ_1, γ_2, etc.) are themselves locally instantiated. Microphysical laws (like (*L*1) and (*L*2)) relating those realizers are therefore themselves liable to hold only *ceteris paribus*, in the absence of interference from microphysical factors not mentioned in their antecedents. But suppose we instead take the realizers of *A* and *C* to be more comprehensive states of affairs (for example, the microstates of the entire Cauchy surfaces upon which *A* and *C* are instantiated). Then some αs may indeed nomically suffice for γs (and not do so only *ceteris paribus*). But other αs (by including interfering factors) may nomically suffice for the absence of any γ. So if we suppose ordinary, locally-instantiated high-level properties to be realized by comprehensive microphysical states, then the problem with the simple model (highlighted by the *ceteris paribus* nature of high-level laws) is not that their realizers are themselves related only by *ceteris paribus* laws, but

rather that not all realizers of one such high-level property will map onto those of another in the neat manner suggested by (*L*1)-(*L*2) and (*B*1) & (*B*3) of our simple model. It turns out that there is independent reason to think that the simple model fails in this latter manner, as shall be seen in the next section.

[31] Cartwright (1983, 1999) has challenged this orthodox view that physical theories are (at least potentially) comprehensive, and that physical laws can therefore (in principle) be characterized without *ceteris paribus* clauses. Nothing of substance in what follows turns upon whether or not we accept the orthodoxy on these matters.

[32] That is to say the *ceteris paribus* clause may be, in the terminology of Schurz (2002), 'indefinite' or, in the terminology of Earman and Roberts (1999, pp. 461-2), 'non-lazy' (see also Earman, Roberts, and Smith 2002, pp. 283-4).

[33] Lange (2002, esp. pp. 407-11) defends the view that there may be such a fact of the matter, even when we are unable to specify the conditions in question so as to render the law 'fully explicit'.

[34] See, for example, Glymour and Wimberly (2007), Halpern (2008), Halpern and Hitchcock (ms.), Halpern and Pearl (2005), Hitchcock (2001, 2007b), Menzies (2004), Pearl (2009, Ch. 10), and Woodward (2003, pp. 74-86).

[35] For discussion of what constitutes an 'appropriate' structural equations model, see Halpern and Hitchcock (2010).

[36] See Hitchcock (2001, pp. 286-7, 289-90), Woodward (2003, pp. 77, 83-4), Glymour and Wimberly (2007, esp. p. 58), and Halpern and Pearl (2005, p. 853-5).

[37] Though Hitchcock and Woodward themselves prefer to call generalizations like (*) 'invariant generalizations' rather than '*ceteris paribus* laws', this *does* appear merely to be a matter of terminology (see Reutlinger *et al.* 2011; and Hitchcock and Woodward 2003a, p. 3). Specifically, the generalizations represented by structural equations like (*) are closely related to what Schurz (2002) calls 'comparative *ceteris paribus* laws' (see Reutlinger *et al.*, *op. cit.*).

[38] As do, among others, Callender and Cohen (2009, p. 25), Reutlinger *et al.* (2011), Schrenk (2007, p. 221), Schurz (2002, p. 351), and Woodward (2002, pp. 303-4).

[39] Ironically, this further respect of non-strictness was first observed by Fodor himself in a different paper (1974, pp. 111-12). In light of this, Fodor's (1989) equation of non-strict with *ceteris paribus* laws may seem surprising. In a third paper, Fodor (1991) speaks of high-level laws that admit of 'absolute exceptions'. Such laws are (at least closely related to) what I'm calling *minutiae rectus* laws.

Fodor (ibid.) attempts to treat such laws as a variety of *ceteris paribus* law. This seems to me problematic for reasons that will become clear in the next section. Schiffer (1991) observes that Fodor's (1989) argument that *ceteris paribus* laws can underwrite nomic sufficiency relations between high-level properties doesn't apply to laws that admit of absolute exceptions. This is a point that will be developed in the next section.

[40] In thermodynamics, the *entropy* of a system can be understood as the proportion of the system's total energy that is *un*available for external mechanical work. Thermodynamic *equilibrium* is reached where the entropy of a system is maximal. This occurs where the values of macro-variables such as temperature, pressure, and mass density are uniform throughout the system.

[41] For example, it seems that Schurz (2002, pp. 369-70) would wish to count the Second Law as a 'theoretically definite (exclusive)' *ceteris paribus* law in virtue of the fact that it appeals to an ideal isolated system. The idea is that, in the theoretically well-defined circumstances in which the isolated system requirement is not met (namely, when mass-energy crosses the boundaries of the system), the *ceteris paribus* clause is violated due to the influence of factors other than the initial thermodynamic state of the system itself.

[42] Someone *might* maintain that the Second Law should be construed as including an implicit *ceteris paribus* condition that supposes away such microphysical realizations. That would be to construe the Second Law's implicit form as something like "the total entropy of a non-equilibrium isolated system increases over time, except when the initial microstate of the system is such that it doesn't". I don't think this is a happy construal (cp. Earman and Roberts 1999, p. 465), for it comes close to rendering the Second Law empty when clearly it isn't.

[43] I'm not claiming that such generalizations should be construed as having a *minutiae rectus* clause built in to them as part of their content (as a high-level generalization might have a *ceteris paribus* clause built into it). That would trivialize them for the reason described in the previous footnote. Rather I'm simply claiming that, as a matter of fact, high-level generalizations typically hold only *minutiae rectus*, and have exceptions where the *minutiae* are not *rectus*: that is, where the properties mentioned in the antecedent of the law are realized in certain unusual microphysical ways.

[44] At least we can't do so unless there is a fact of the matter about *which* realizer of *B* would be instantiated by x if x were non-A *and* the fact of the matter is that β_2 would be instantiated. There seems no good reason to suppose there typically will be such a fact and that it will be the fact that we need.

[45] Alternatively, we could say that $V_C = V_A$ does hold, as a *minutiae rectus* law, but that it fails to support the counterfactuals appealed to by structural equations analyses of causation.

[46] Again, this also means that our equation $V_C = V_A$ doesn't hold (or that it holds only as a non-counterfactual-supporting *minutiae rectus* law).

[47] I take it to be very plausible that the might-counterfactual ($\phi \Diamond\!\!\rightarrow \psi$) and the would-not counterfactual ($\phi \Box\!\!\rightarrow \sim\psi$) are *contraries*: they can't both be true together. This is weaker than the (also plausible) view (see Lewis 1973a) that the two counterfactuals are *duals* (i.e. that ($\phi \Diamond\!\!\rightarrow \psi$) $\leftrightarrow$ $\sim(\phi \Box\!\!\rightarrow \sim\psi)$). The thesis that the two counterfactuals are at least contraries is defended by Hájek (ms.), who also draws upon the fact that there are possible micro-realizers of thermodynamic states of systems that lead to sustained future entropy decrease to argue that most ordinary would-counterfactuals about the outcomes of thermodynamically irreversible processes are false. Hájek replies to several possible responses, including 'contextualist' and 'pragmatist' responses, according to which (respectively) the would-counterfactuals in question express true propositions in some contexts, or are true but infelicitous. Moss (2012) argues that a pragmatist response is supported by the embedding behavior of would-counterfactuals. I find this strategy – of appealing to linguistic intuitions about embedded counterfactuals – dialectically slightly strange given that it is Hájek's view that most ordinary *judgments* about counterfactuals (as well as the counterfactuals themselves!) are false. But, more importantly, Moss's specific examples seem to me to elicit the intuitions that they do only by eliciting the dubious intuition of the truth of conditional excluded middle, a principle which looks doubtful for precisely the same reasons as Hájek offers for thinking most ordinary would-counterfactuals are false. This is not, however, the place to argue the point in detail.

[48] Like a deterministic law, a probabilistic law is exceptionless if its consequent is always true when its antecedent is fulfilled. Unlike a deterministic law, this does not mean that the properties mentioned in the antecedent must always be accompanied by the properties mentioned in the consequent. Rather, it means that the properties mentioned in the antecedent must always be accompanied by the *probability distribution* (over the instantiation of various further properties) mentioned in the consequent.

[49] Albert (ibid. pp. 150-62; 2012, pp. 39-40) argues that the GRW interpretation is an exception to this and that, if it is correct, the SM probabilities might be derivable from the fundamental dynamics themselves, without the need for a probability distribution over the possible initial quantum states of a system.

[50] By contrast, if the Second Law were construed as a probabilistic *minutiae rectus* law, then it would not be capable of underwriting the nomic sufficiency of the initial thermodynamic state of the ice-in-hot-water system for a high probability of the ice's melting, given the existence of possible realizers of such an initial state that confer low probability on melting. Likewise, the Second Law would not be capable of underwriting the nomic sufficiency of various alternative macrostates for a high probability of entropy increase (as described in the following two paragraphs of the main text).

[51] Probability-raising understood in the conditional probability sense figures in the probabilistic analyses of causation developed by Reichenbach (1971), Good (1961a, 1961b), Suppes (1970), Kvart (2004), and Glynn (2011). Understood in the counterfactual sense, it figures in the analyses given by Lewis (1986), Menzies (1989), and Ramachandran (2004). Glynn (2011) develops an analysis of causation that is a probabilistic analogue of the deterministic structural equations approaches.

[52] The 'normality' of the glass of hot water is intended to cover the fact that the glass is of a usual sort of size, that it is not unusually cold, that it has a decent amount of hot water in it, and so on. The 'normality' of the ice cube covers the fact that it is an ice cube of the same sort of size, temperature, purity, etc. as those that people typically put in their drinks. In other words, to require the 'normality' of the ice cube and the glass of hot water is to indicate, in an extremely imprecise way, the ranges within which the values of various macro-variables pertaining to the ice-cube-in-hot-water system must lie to get a high SM probability of melting. I don't think the two occurrences of 'normal' are redundant given the *ceteris paribus* clause: even in the absence of non-negligible interference from outside the ice-in-hot-water system, an abnormally large and cold ice cube in an abnormally cool glass containing an abnormally small amount of hot water is unlikely to melt quickly. This is not, however, a point on which I need to insist.

[53] In actual fact, Albert (2000, e.g. pp. 94-6; 2012, pp. 24-8) and Loewer (2012b, p. 18) argue that we can (in principle) get an SM probability distribution for the future evolution of a system *given only* relatively imprecise information about its macrostate. Perhaps the information that the system comprises a normal glass of hot water containing a normal ice cube is enough. We are, of course, liable to get a *different* SM probability distribution given a *precise* specification of the system's initial macrostate. But conditional upon either sort of specification, the SM probability of the system's evolving into a later macrostate in which the ice has melted is very high, and higher than that associated with (more or less precise) specifications of initial macrostates of reasonable ice-in-freezer systems.

[54] For expressions of skepticism, see Callender and Cohen (2010, pp. 437-8), Callender (2011, p. 103), and Dunn (2011, p. 84). See Weslake (forthcoming) and Frisch (forthcoming) for sustained defenses of such skepticism.

[55] See also Loewer (2001, p. 610; 2007a, p. 300; 2012a, p. 124; 2012b, p. 16) and Albert (2012, p. 20).

[56] In the quantum case, the uniform probability distribution is not over classical phase space, but over the set of quantum states compossible with PH.

[57] The term is from the Coen brothers movie *A Serious Man*, in which a character is working on 'the probability map of the universe', which he calls 'the Mentaculus' (see Loewer 2012a, p. 124; 2012b, p. 16).

[58] See Earman (2006, esp. pp. 418-20) for skepticism about whether such a distribution can explain the temporally asymmetric behavior of the sub-systems of the universe that we observe.

[59] For details, see Albert (2000, Ch. 4).

[60] Glynn (2010, pp. 63-5) argues that *any* satisfactory account of laws must be able to accommodate a probabilistic Second Law as a genuine law for worlds like ours (cp. also Loewer 2007a, pp. 305-6). For reasons of space I shall focus on variants of the BSA, which I believe to be the most attractive philosophical accounts of lawhood available.

[61] A system's *fit* is the probability that it assigns to the actual course of history (ibid. p. 480). See Elga (2004) for a critique of, and suggested amendment to, Lewis's notion of fit. The question of which notion is the appropriate one to use needn't detain us here.

[62] The latter distributions are liable to be trivial – that is the chances will all be 1s and 0s – if the fundamental dynamics are deterministic.

[63] See also Callender and Cohen (2009, p. 10) and Dunn (2011, p. 91).

[64] Albert (2000, p. 96; 2012, pp. 20, 28) takes PH and SP to be genuine laws, but Loewer (2001; 2007a, p. 305) was the first argue that the BSA entails their lawfulness (together with that of the probabilistic version of the Second Law that they entail), as well as the objective chancehood of the SM probabilities. Hoefer (2007, pp. 559-60) and Glynn (2010, pp. 62-3) answer some objections to Loewer's argument due to Schaffer (2007, pp. 130-1).

[65] For detailed defenses of the claim that the SM probabilities are objective chances, see Loewer (2001), Hoefer (2007), Glynn (2010), and Frigg and Hoefer (2010).

[66] For similar proposals, see Dunn (2011, pp. 88-90) and Schrenk (2008).

[67] Frisch (forthcoming) suggests an alternative variant of the Best Systems approach which makes it

plausible that the laws of the special sciences, together with those of fundamental physics, are part of a single 'big' (non-vocabulary relative) best system. Frisch achieves this by, firstly, taking a pragmatic approach to the vocabulary that the axioms must be expressed in: they must be framed in terms of predicates referring to properties *that we are interested in* (whether those properties be perfectly natural or not). Secondly, he adds a desideratum (to weigh alongside the simplicity, strength, and fit of a system) according to which a system is better if its theorems can be derived from its axioms in fewer steps. (Specifically, the informativeness of a theorem is discounted by the length of the required proof.) The idea is that a system to which informative special science generalizations are added as axioms will achieve high informativeness-relative-to-distance-from-the-axioms (and so will plausibly come out *best*), since these generalizations no longer have to be derived as theorems, as they are from the Mentaculus.

Callender and Cohen (2009, pp. 10, 28) and Callender (2011, pp. 106-12) treat thermodynamics and statistical mechanics as at least on a par with the special sciences in the sense that they are likely upshots of best system competitions conducted in their own proprietary vocabularies, which include predicates – such as *entropy* – that don't refer to fundamental natural kinds. Similarly, Dunn (2011, pp. 80-1, 91) explicitly treats thermodynamics as just one special science among many (cp. Loewer 2012b, pp. 15, 18). Winsberg (2008, p. 884) objects that there is no distinctive proprietary vocabulary for statistical mechanics: there is only the thermodynamic language and the microphysical language. Weslake (forthcoming) suggests that the correct response is to see statistical mechanics as a best system for the *conjunction* of the fundamental kinds and the thermodynamic kinds. Frisch's (forthcoming) 'big' best systems account seems suited to accommodating axiom systems that, in this way, draw upon multiple scientific vocabularies.

[68] Indeed, given that Lewis (1994, p. 478) restricts candidates for best systemhood to those systems whose theorems are true, it is difficult to see how those that have theorems that hold only *minutiae rectus* (as opposed to those that have probabilistic theorems), and which therefore seem to be literally false universal generalizations, could even be in the running. (Though one option would be to require merely the approximate truth of theorems of candidate best systems; cp. Dunn 2011, p. 91.) However, we are in danger of proving too much here, for there is also a question about whether *ceteris paribus* generalizations can be understood as literally true (see Lange 2002, pp. 411-14). If not, then – since most special science generalizations are *ceteris paribus* generalizations – it is difficult to see how any best system account could treat special science generalizations as genuine laws. Yet the

usual view of *ceteris paribus* laws (at least among their friends) is that they *are* literally true, and that the challenge is to show how they can be non-trivially so (and thus able to add strength to a system that entails them). Fodor (1989, 1991) and Schrenk (2008) are examples of attempts to give non-trivial truth-conditions for *ceteris paribus* laws. Callender and Cohen (2009, p. 25) assume that *ceteris paribus* laws will have some non-vacuous truth conditions, and are therefore candidates for inclusion in a Best System. Schrenk (2008) provides an argument that this is so.

References

Albert, D. (2000): *Time and Chance* (Cambridge, MA: Harvard University Press).

Albert, D. (2012): 'Physics and Chance', in Y. Ben-Menahem and M. Hemmo (eds.): *Probability in Physics* (Berlin: Springer), pp. 17-40.

Armstrong, D. (1983): *What is a Law of Nature?* (Cambridge: Cambridge University Press).

Bennett, K. (2003): 'Why the Exclusion Problem Seems Intractable, and How, Just Maybe, to Tract It', *Noûs*, 37, pp. 471-497.

Block, N. (1990): 'Can the Mind Change the World?', in G. Boolos (ed.): *Meaning and Method: Essays in Honor of Hilary Putnam*, (Cambridge: Cambridge University Press), pp. 137-170.

Block, N. (2003): 'Do Causal Powers Drain Away?', *Philosophy and Phenomenological Research*, 67, pp. 133-150.

Callender, C. (2011): 'The Past Histories of Molecules', in C. Beisbart and S. Hartmann (eds.) *Probabilities in Physics* (New York: OUP), pp. 83-113.

Callender, C. and Cohen, J. (2009): 'A Better Best System Account of Lawhood', *Philosophical Studies*, 145, pp. 1-34.

Callender, C. and Cohen, J. (2010): 'Special Sciences, Conspiracy and the Better Best System Account of Lawhood', *Erkenntnis*, 73, pp. 427-447.

Cartwright, N. (1983): *How the Laws of Physics Lie* (Oxford: OUP).

Cartwright, N. (1999): *The Dappled World* (Cambridge: CUP).

Choi, S. (2002). 'Causation and gerrymandered world lines: a critique of Salmon', *Philosophy of Science*, 69, pp. 105-117.

Crane, T. (1995): 'The Mental Causation Debate', *Proceedings of the Aristotelian Society*, *Supplementary Volume*, 69, pp. 211-236.

Davidson, D. (1967): 'Causal Relations', *Journal of Philosophy*, 64, pp. 691-703.

Davidson, D. (1970): 'Mental Events', in L. Foster & J. Swanson (eds.): *Experience & Theory* (Amherst: University of Massachusetts Press), pp. 79-101.

Davidson, D. (1974): 'Psychology as Philosophy', in S. Brown (ed.): *Philosophy of Psychology* (New York: Barnes & Noble), pp. 41-52.

Davidson, D. (1993): 'Thinking Causes', in J. Heil and A. Mele (*eds.*) (1993): *Mental Causation* (Oxford: Clarendon Press), pp. 3-17.

Dowe, P. (2000): *Physical Causation* (Cambridge: CUP).

Dunn, J. (2011): 'Fried Eggs, Thermodynamics, and the Special Sciences', *British Journal for the Philosophy of Science*, 62, pp. 71–98.

Earman, J. (2006): 'The "Past Hypothesis": Not Even False', *Studies in History and Philosophy of Modern Physics*, 37, pp. 399-430.

Earman, J, and J. Roberts (1999): '*Ceteris Paribus*, There is no Problem of Provisos', *Synthese*, 118, pp. 439-478.

Earman, J., J. Roberts, and S. Smith (2002): 'Ceteris Paribus Lost', *Erkenntnis*, 57, pp. 281-301.

Elga, A. (2001): 'Statistical Mechanics and the Asymmetry of Counterfactual Dependence', *Philosophy of Science*, 68 (Proceedings), pp. 313-324.

Elga, A (2004): 'Infinitesimal Chances and the Laws of Nature', in F. Jackson and G. Priest (eds.): *Lewisian Themes: The Philosophy of David K. Lewis* (Oxford: OUP), pp. 68-77.

Field, H. (2003): 'Causation in a Physical World', in M. Loux and D. Zimmerman (eds.): *Oxford Handbook of Metaphysics* (Oxford: OUP), pp. 435-460.

Fodor, J. (1974): 'Special Sciences (or: The Disunity of Science as a Working Hypothesis) ', *Synthese*, 28, pp. 77-115.

Fodor, J. (1989): 'Making Mind Matter More', *Philosophical Topics*, 17, pp. 59-79.

Fodor, J. (1991): 'You Can Fool Some of The People All of The Time, Everything Else Being Equal; Hedged Laws and Psychological Explanations', *Mind*, 100, pp. 19-34.

Fodor, J. (1997): 'Special Sciences: Still Autonomous After All These Years', *Noûs*, 31, pp. 149-163.

Frigg, R. and Hoefer, C. (2010): 'Determinism and Chance from a Humean Perspective', in D. Dieks *et al.* (eds.): *The Present Situation in the Philosophy of Science*, 1, pp. 351-371.

Frisch, M. (forthcoming): 'Why Physics Can't Explain Everything', in A. Wilson (ed.): *Asymmetries of Chance and Time* (Oxford: OUP).

Glymour, C. & Wimberly, F. (2007): 'Actual Causes and Thought Experiments', in J. K. Campbell, M. O'Rourke, & H. Silverstein (eds.), *Causation and Explanation* (Cambridge, MA: MIT Press), pp.

43-67.

Glynn, L. (2010): 'Deterministic Chance', *British Journal for the Philosophy of Science*, 61, pp. 51-80.

Glynn, L. (2011): 'A Probabilistic Analysis of Causation', *British Journal for the Philosophy of Science*, 62, pp. 343-392.

Good, I. (1961a): 'A Causal Calculus (I) ', *British Journal for the Philosophy of Science*, 11, pp. 305-318.

Good, I. (1961b): 'A Causal Calculus (II)', *British Journal for the Philosophy of Science*, 12, pp. 43-51.

Hájek, A. (ms.): 'Most Counterfactuals are False'.

Halpern, J. Y. (2008): 'Defaults and Normality in Causal Structures', in G. Brewka & J. Lang (eds.), *Proceedings of the Eleventh International Conference on Principles of Knowledge Representation and Reasoning* (Menlo Park, CA: AAAI Press.), pp. 198-208.

Halpern, J. Y. & Hitchcock, C. (2010): 'Actual Causation and the Art of Modeling', in R. Dechter, H. Geffner, & J. Y. Halpern (eds.), *Heuristics, Probability and Causality: A Tribute to Judea Pearl* (London: College Publications), pp. 383-406.

Halpern, J. Y. & Hitchcock, C. (ms.): 'Graded Causation and Defaults'.

Halpern, J. Y. & Pearl, J. (2005): 'Causes and Explanations: A Structural-Model Approach, Part I: Causes', *British Journal for the Philosophy of Science* 56, pp. 843-887.

Hitchcock, C. (1995): 'Salmon on Explanatory Relevance', *Philosophy of Science*, 62, pp. 304-320.

Hitchcock, C. (2001). 'The Intransitivity of Causation Revealed in Equations and Graphs', *Journal of Philosophy*, 98, 273-299.

Hitchcock, C. (2007a): 'What Russell Got Right', in H. Price and R. Corry (eds.) (2007): *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited* (Oxford: Clarendon Press), pp. 45-65.

Hitchcock, C. (2007b). 'Prevention, Preemption, and the Principle of Sufficient Reason', *Philosophical Review*, 116, pp. 495-532.

Hitchcock, C. (2009): 'Problems for the Conserved Quantity Theory: Counterexamples, Circularity, and Redundancy', *Monist*, 92, pp. 72-93.

Hitchcock, C. and Woodward, J. (2003a): 'Explanatory Generalizations, Part I: A Counterfactual Approach', *Noûs*, 37, pp. 1-24.

Hitchcock, C. and Woodward, J. (2003b): 'Explanatory Generalizations, Part II: Plumbing Explanatory Depth', *Noûs*, 37, pp. 181-199.

Hoefer, C. (2007): 'The Third Way on Objective Chance: A Sceptic's Guide to Objective Chance', *Mind*, 116, pp. 549-596.

Honderich, T. (1982): 'The Argument for Anomalous Monism', *Analysis*, 42, pp. 59-64.

Horgan, T. (1997): 'Kim on Mental Causation and Causal Exclusion', *Noûs*, 31, Supplement: *Philosophical Perspectives*, 11, pp. 165-184.

Ismael, J. (2009): 'Probability in Deterministic Physics', *Journal of Philosophy*, 106, pp. 89-108.

Ismael, J. (2011): 'A Modest Proposal About Chance', *Journal of Philosophy*, 108, pp. 416-442.

Jackson, F. and Pettit, P. (1988): 'Functionalism and Broad Content', *Mind*, 97, pp. 381-400.

Jackson, F. and Pettit, P. (1990): 'Program Explanation: A General Perspective', *Analysis*, 50, pp. 107-117.

Kim, J. (1979): 'Causality, Identity, and Supervenience in the Mind-Body Problem', *Midwest Studies in Philosophy*, 4, pp. 31-49.

Kim, J. (1984): 'Epiphenomenal and Supervenient Causation', *Midwest Studies in Philosophy*, 9, pp. 257-270.

Kim, J. (1987): '"Strong" and "Global" Supervenience Revisited', *Philosophy and Phenomenological Research*, 48, pp. 315-326.

Kim, J. (1989a): 'Mechanism, Purpose, and Explanatory Exclusion', *Philosophical Perspectives*, 3, pp. 77-108.

Kim, J. (1989b): 'The Myth of Nonreductive Materialism', *Proceedings of the American Philosophical Association*, 63, pp. 31-47.

Kim, J. (1992): 'Multiple Realization and the Metaphysics of Reduction', *Philosophy and Phenomenological Research*, 52, pp. 1-26.

Kim, J. (1993): 'Can Supervenience and "Non-Strict Laws" Save Anomalous Monism', in J. Heil and A. Mele (*eds.*) (1993): *Mental Causation* (Oxford: Clarendon Press), pp. 19-26.

Kim, J. (1998): *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*, (Cambridge, MA: MIT Press).

Kim, J. (2002): 'The Layered Model: Metaphysical Considerations', *Philosophical Explorations*, 5, pp. 2-20.

Kroedel, T. (2008): 'Mental Causation as Multiple Causation', *Philosophical Studies*, 139, pp. 125-143.

Kvart, I. (2004): 'Causation: Probabilistic and Counterfactual Analyses', in J. Collins, N. Hall, and L. A. Paul (eds.) (2004): *Causation and Counterfactuals* (Cambridge, MA: MIT Press), pp. 359-386.

Lange, M. (2002): 'Who's Afraid of *Ceteris-Paribus* Laws? Or: How I Learned to Stop Worrying and Love Them', *Erkenntnis*, 57, pp. 407-423.

LePore, E. and Loewer, B. (1987): 'Mind Matters', *Journal of Philosophy*, 84, pp. 630-642.

Lewis, D. (1973a): *Counterfactuals* (Cambridge, MA: Harvard University Press).

Lewis, D. (1973b): 'Causation', *Journal of Philosophy*, 70, pp. 556-567.

Lewis, D. (1979): 'Counterfactual Dependence and Time's Arrow', *Noûs*, pp. 455-476.

Lewis, D. (1983): 'New Work for a Theory of Universals', *Australasian Journal of Philosophy*, 61, pp. 343-377.

Lewis, D. (1986): 'Postscripts to "Causation"', in his *Philosophical Papers*, Vol. 2 (Oxford: OUP), pp. 172-213.

Lewis, D. (1994): 'Humean Supervenience Debugged', *Mind*, 103, pp. 473-490.

Loewer, B. (2001): 'Determinism and Chance', *Studies in History and Philosophy of Science*, 32, pp. 609-620.

Loewer, B. (2007a): 'Counterfactuals and the Second Law', in H. Price and R. Corry (eds.) (2007): *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited* (Oxford: Clarendon Press), pp. 293-326.

Loewer, B. (2007b): 'Mental Causation, or Something Near Enough', in B. McLaughlin and J. Cohen (eds.): *Contemporary Debates in Philosophy of Mind* (Oxford: Blackwell), pp. 243-264.

Loewer, B. (2008): 'Why There *Is* Anything Except Physics', in J. Hohwy and J. Kallestrup (eds.): *Being Reduced: New Essays on Reduction, Explanation and Causation* (Oxford: OUP), pp. 149-163.

Loewer, B. (2012a): 'Two Accounts of Laws and Time', *Philosophical Studies*, 160, pp. 115-137.

Loewer, B. (2012b): 'The Emergence of Time's Arrows and Special Science Laws from Physics', *Interface Focus*, 2, pp. 13-19.

Malcolm, N. (1968): 'The Conceivability of Mechanism', *Philosophical Review*, 77, pp. 45-72.

McLaughlin, B. (1989): 'Type Epiphenomenalism, Type Dualism, and the Causal Priority of the Physical', *Philosophical Perspectives*, 3, pp. 109-135.

McLaughlin, B. (1993): 'On Davidson's Response to the Charge of Epiphenomenalism', in J. Heil and A. Mele (*eds.*) (1993): *Mental Causation* (Oxford: Clarendon Press), pp. 27-40.

Menzies, P. (1989): 'Probabilistic Causation and Causal Processes: A Critique of Lewis', *Philosophy of Science*, 56, pp. 642-663.

Menzies, P. (2004): 'Causal Models, Token Causation, and Processes', *Philosophy of Science*, 71, pp. 820-832.

Moss, S. (2012): 'On the Pragmatics of Counterfactuals', *Noûs*, 46, pp. 561-586.

Ney, A. (2009): 'Physical Causation and Difference-Making', *British Journal for the Philosophy of Science*, 60, pp. 737-764.

Pearl, J. (2009): *Causality: Models, Reasoning, and Inference.* 2nd Ed. (Cambridge: Cambridge University Press).

Price, H. and Corry, R. (eds.) (2007): *Causation, Physics, and the Constitution of Reality: Russell's Republic*

Revisited (Oxford: Clarendon Press).

Putnam, H. (1967): 'Psychological Predicates', in W. Capitan and D. Merrill (eds.): *Art, Mind, and Religion* (Pittsburg: University of Pittsburg Press), pp. 37-48.

Ramachandran, M. (2004): 'A Counterfactual Analysis of Indeterministic Causation' in J. Collins, N. Hall, and L. A. Paul (eds.) (2004): *Causation and Counterfactuals* (Cambridge, MA: MIT Press), pp. 387-402.

Reichenbach, H. (1971): *The Direction of Time*, M. Reichenbach (ed.) (Berkeley and Los Angeles, CA: University of California Press).

Reutlinger, A., Schurz, G. and Hüttemann, A. (2011): '*Ceteris Paribus* Laws', *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2011/entries/ceteris-paribus/>.

Russell, B. (1913): 'On the Notion of Cause', *Proceedings of the Aristotelian Society*, 13, pp. 1-26.

Sober, E. (1999): 'The Multiple Realizability Argument against Reductionism', *Philosophy of Science*, 66, pp. 542-564.

Sosa, E. (1984): 'Mind-Body Interaction and Supervenient Causation', *Midwest Studies in Philosophy*, 9, pp. 271-281.

Salmon, W. (1984): *Scientific Explanation and the Causal Structure of the World* (Princeton, NJ: Princeton University Press).

Salmon, W. (1994): 'Causality without Counterfactuals', *Philosophy of Science*, 61, pp. 297-312.

Salmon, W. (1997): 'Causality and Explanation: A Reply to Two Critiques', *Philosophy of Science*, 64, pp. 461-477.

Schaffer, J. (2003): 'Principled Chances' *British Journal for the Philosophy of Science*, 54, pp. 27-41.

Schaffer, J. (2007): 'Deterministic Chance?', *British Journal for the Philosophy of Science*, 58, pp. 113-140.

Schiffer, S. (1991): 'Ceteris Paribus Laws', *Mind*, 100, pp. 1-17.

Schrenk, M. (2007): 'Can Capacities Rescue Us From Ceteris Paribus Laws?' in M. Kistler, B. Gnassounou (eds.): *Dispositions in Philosophy and Science* (Aldershot, UK: Ashgate), pp. 221–247.

Schrenk, M. (2008): 'A Lewisian Theory for Special Science Laws', in S. Walter and H. Bohse (eds): *Ausgewählte Beiträge zu den Sektionen der GAP.6* (Paderborn: Mentis), pp. 121-131.

Schurz, G. (2002): '*Ceteris Paribus* Laws: Classification and Deconstruction', *Erkenntnis*, 57, pp. 351-372.

Stoutland, F. (1980): 'Oblique Causation and Reasons for Action', *Synthese*, 43, pp. 351-367.

Suppes, P. (1970): *A Probabilistic Theory of Causality*, *Acta Philosophical Fennica* (Amsterdam: North-Holland).

Weslake, B. (forthcoming): 'Statistical Mechanical Imperialism', in A. Wilson (ed.): *Asymmetries of Chance and Time* (Oxford: OUP).

Winsberg, E. (2008): 'Laws and Chances in Statistical Mechanics', *Studies in History and Philosophy of Modern Physics*, 39, pp. 872-888.

Woodward, J. (2002): 'There is No Such Thing as a *Ceteris Paribus* Law', *Erkenntnis*, 57, pp. 303-328.

Woodward, J. (2003): *Making Things Happen* (Oxford: OUP, 2003).

Yablo, S. (1992): 'Mental Causation', *Philosophical Review*, 101, pp. 245-280.