

Laws, causation and dynamics at different levels

Jeremy Butterfield*

Trinity College, University of Cambridge, Cambridge CB2 1TQ, UK

I have two main aims. The first is general, and more philosophical (§2). The second is specific, and more closely related to physics (§§3 and 4). The first aim is to state my general views about laws and causation at different ‘levels’. The main task is to understand how the higher levels sustain notions of law and causation that ‘ride free’ of reductions to the lower level or levels. I endeavour to relate my views to those of other symposiasts. The second aim is to give a framework for describing dynamics at different levels, emphasizing how the various levels’ dynamics can mesh or fail to mesh. This framework is essentially that of elementary dynamical systems theory. The main idea will be, for simplicity, to work with just two levels, dubbed ‘micro’ and ‘macro’, which are related by coarse-graining. I use this framework to describe, in part, the first four of Ellis’ five types of top-down causation.

Keywords: emergence; causation; reduction; dynamics; supervenience; coarse-graining

1. INTRODUCTION

I have two main aims. The first is general, and more philosophical (§2). It concerns not just this Theme Issue’s topic, top-down causation, but the general relations between ‘levels’. The second aim is specific, and more closely related both to top-down causation and to physics, in particular dynamical systems theory (§§3 and 4).

In discussing relations between levels, I will take it that the overall task is to understand how the higher levels sustain notions of law, causation and explanation that are ‘autonomous’, or ‘ride free’, from whatever reductions there might be to the lower level or levels (or at least: notions that *seem* to be autonomous or to ride free!). This is a large task, with a large literature of controversy, both nowadays and in the past. The reasons for the controversy are obvious. People disagree about how to understand the notions of level and reduction, and also those of law, causation and explanation. They disagree about the extent to which, and the sense in which, the higher levels are autonomous or ride free. And these disagreements are fuelled by having different sets of scientific examples in mind.

These disagreements become more vivid (and more comprehensible) when one considers historical changes in the disputants’ scientific examples. One broad example is the demise of vitalism. Before the century-long rise of microbiology, biochemistry and molecular biology (from say 1860 to 1960), it was perfectly sensible to believe that biological processes depended on certain ‘vital forces’; and therefore that, as the slogan puts it, ‘biology is not reducible to chemistry and

physics’—in a much stronger sense of ‘not reducible’ than one could believe today. Other broad examples extending over many decades are: (i) the rise of atomism and statistical mechanics in explaining irreversible macroscopic processes and (ii) the rise of atomism, the periodic table and then quantum chemistry, in explaining chemical bonding. Again: before these developments, one could believe in irreducibility, e.g. of chemical bonding to physics, in a much stronger sense than one can today. In short: what historians now call ‘the second scientific revolution’ from 1850 onwards has given us countless successful reductions of behaviour (both specific processes and general laws) at a higher (often macroscopic) level to facts at a lower (often microscopic) level.

Thus, the overall philosophical task, both nowadays and in yesteryear, is: first, to state and defend notions of level and reduction, and of law, causation and explanation; and second, to use them to assess, in a wide range of contemporary scientific examples, the extent to which, and the sense in which, the higher levels are autonomous, or ride free, from the lower levels. But nowadays, after the triumphs of the second scientific revolution, we must expect the extent of, and/or senses for, such autonomy of the higher levels to be more restricted and/or more subtle. One aspect of this overall task is the topic of this special issue: assessing the prospects for top-down causation.

My own contribution will proceed in two stages. First (§2), I will summarize some of my own views about the overall task, relating them to top-down causation and the views of some other authors. For example, I will briefly endorse some views of Sober’s about reduction and causation, and List & Menzies’ recent defence of top-down causation (§2.2 and 2.4). My overall views are defended in detail elsewhere [1,2].

*jb56@cam.ac.uk

One contribution of 15 to a Theme Issue ‘Top-down causation’.

They mostly concern emergence, reduction and supervenience (§2 will report my construals of these contested terms). I should admit at the outset that I will have nothing distinctive to say about the notions of law, causation and explanation. But, in fact, I take a broadly Humean view of all three; and it will be clear that this will fit well with my views on emergence, reduction and supervenience.

Second, in §§3 and 4, I will give a framework for describing dynamics at different levels, emphasizing how the various levels' dynamics can mesh, or fail to mesh. This framework is essentially that of elementary dynamical systems theory. The main idea will be, for simplicity, to work with just two levels, dubbed 'micro' and 'macro', which are related by coarse-graining. I then consider two topics, in §§3 and 4, respectively.

First, there is the question whether a micro-dynamics, together with a coarse-graining prescription, induces a well-defined macro-dynamics (§3). I describe how physics provides some precise and important examples of such 'meshing' (e.g. in statistical mechanics), as well as examples of where it fails. I will stress that failure of meshing need not be a problem, let alone a mystery: the pilot-wave theory, in the foundations of quantum mechanics, will provide a non-problematic example. I also discuss how to secure meshing by redefining the coarse-graining; and relate the topic to the philosophical views of Fodor & Papineau on multiple realizability, and of List on free will.

Second, I use the framework to describe, in part, the first four of Ellis' [3,4] five types of top-down causation (§4). There are various choices to be made in giving such a dynamical systems description of Ellis' typology; but I maintain that the fit is pretty good. In particular, I note that Ellis calls my first topic above, i.e. the meshing of micro- and macro-dynamics, 'coherent higher level dynamics' or 'the principle of equivalence of classes'; and he takes it as a presupposition of his typology of top-down causation.

Finally, a clarification. This paper has some 'reductionist' features, which might be misleading. Thus, in §2 I will join Sober & Papineau in rejecting the multiple realizability argument against 'reductionism'. And in §4, I will not try to articulate the differences between my formal descriptions, in the jargon of dynamical systems, of Ellis' types of top-down causation, and Ellis' own informal and richer descriptions. These features might suggest that I deny any or all of the following three claims:

- (i) There are, or can be, laws and/or explanation and/or causation at 'higher levels', or in the special sciences.
- (ii) There is a good notion of causation beyond that of functional dependence of one quantity on another.
- (iii) Top-down causation, at least of Ellis' types, needs more than my dynamical systems framework.

But, in fact, I endorse (i)–(iii). It is just that they are not centre-stage in my discussion.

2. REDUCTION, SUPERVENIENCE AND CAUSATION

My first aim is to summarize some of my views about the relations between levels. Sections 2.1 and 2.2 discuss reduction and 'multiple realizability', and §2 discusses supervenience. Broadly speaking, I will deny the widespread (perhaps even orthodox?) views that multiple realizability prevents reduction, and that levels are typically related by supervenience without reduction. Section 2.4 concerns causation: here I will endorse Shapiro & Sober's, and List & Menzies', recent arguments for top-down causation.

2.1. Reduction

I have analysed the relations between reduction, emergence and supervenience, elsewhere [1,2]. In short, I construe these notions as follows. I take emergence as a system's having behaviour, i.e. properties and/or laws, that is novel and robust relative to some natural comparison class. Typically, the behaviour concerned is collective or macroscopic; and it is novel compared with the properties and laws that are manifest in (theory of) the microscopic details of the system. I take reduction as a relation between theories: namely, deduction using appropriate auxiliary definitions. (As we will see, this is in effect a strengthening of the traditional Nagelian conception of reduction.) And I take supervenience as a weakening of this concept of reduction; namely, to allow infinitely long definitions (more details in §2.3).

Then my main claim was that, with these meanings, emergence is logically independent both of reduction and of supervenience. In particular, one can have emergence with reduction, as well as without it. Physics provides many such examples, especially where one theory is obtained from another by taking a limit of some parameter. That is, there are many examples in which we deduce a novel and robust behaviour, by taking the limit of a parameter.¹ And emergence is also independent of supervenience: one can have emergence without supervenience, as well as with it.

Broadly speaking, this main claim gives some support to the 'autonomy' of higher levels (cf. claim (i) at the end of §1), namely by reconciling such autonomy with the existence of reductions to lower levels. Some of my other claims had a similar reconciling intent: e.g. my joining Sober & Papineau in holding that multiple realizability is no problem for reductionism.² I shall develop this position a little by discussing Nagelian reduction (this subsection), multiple realizability (§2.2) and supervenience (§2.3).

Nagel's idea is that reduction should be modelled on the logical idea of one theory being a *definitional extension* of another.³ Writing t for 'top' and b for 'bottom', we say: T_t is a *definitional extension* of T_b , iff one can add to T_b a set D of definitions, one for each of T_t 's

¹I analysed four examples [2]. Footnote 6 and §3.2 will mention yet other examples.

²There were other claims I will not need here, e.g. that emergence does not require limits, in particular not 'singular' limits.

³The main source is Nagel [5, pp. 351–363]; see also Hempel [6, ch. 8]. Schaffner [7] is a masterly review not only of Nagel's position, but also of others' critiques, defences and modifications of Nagel.

non-logical symbols, in such a way that T_t becomes a sub-theory of the augmented theory $T_b \cup D$. That is, in the augmented theory, we can prove every theorem of T_t . Here, a *definition* is a statement, for a predicate, of co-extension; and for a singular term, of co-reference. To be precise: for a predicate P of T_t it would be a universally quantified biconditional with P on the left-hand side stating that P is co-extensive with a right-hand side that is an open sentence ϕ of T_b built using such operations as Boolean connectives and quantifiers. Thus, if P is n -place: $(\forall x_1) \dots (\forall x_n)(P(x_1, \dots, x_n) \equiv \phi(x_1, \dots, x_n))$. (The definitions are often called ‘bridge laws’ or ‘bridge principles’.)

A *caveat*. I said that Nagel held that reduction ‘should be modelled on’ the idea of definitional extension, because definitional extension is (a) sometimes too weak as a notion of reduction and (b) sometimes too strong.

As to (a): Nagel [5, pp. 358–363] holds that the reducing theory T_b should explain the reduced theory T_t ; and following Hempel, he conceives explanation in deductive–nomological terms. Thus he says, in effect, that T_b reduces T_t iff:

- (i) T_t is a definitional extension of T_b ; and
- (ii) in each of the definitions of T_t ’s terms, the *definiens* in the language of T_b must play a role in T_b ; so it cannot be, for example, a heterogeneous disjunction.

As to (b): definitional extension is sometimes too strong as a notion of reduction; as when T_b *corrects*, rather than implies, T_t . Thus, Nagel says that a case in which T_t ’s laws are a close approximation to what strictly follows from T_b should count as reduction, and be called ‘approximative reduction’.

More important for us is the fact that definitional extensions, and thereby Nagelian reduction, can perfectly well accommodate what philosophers call *functional definitions*. These are definitions of a predicate or other non-logical symbol (or in ontic, rather than linguistic, jargon: of a property, relation, etc.) that are second-order, i.e. that quantify over a given ‘bottom set’ of properties and relations. The idea is that the *definiens* states a pattern among such properties, typically a pattern of causal and law-like relations between properties. So an n -tuple of bottom properties that instantiates the pattern in one case is called a *realizer* or *realization* of the *definiendum*. And the fact that, in different cases, different such n -tuples instantiate the pattern is called *multiple realizability*. Examples of functional or second-order properties, and so of multiple realizability, are legion. For example, the property of being locked is instantiated very differently in padlocks using keys, combination locks, etc.

2.2. The multiple realizability argument refuted

Multiple realizability is undoubtedly a key idea, philosophically and scientifically, for our overall task: understanding relations between levels, and especially how higher levels can be autonomous, or ride free, from lower levels. Admittedly, there is not much to be said by way of a theory about being locked and similarly

for countless other multiply realizable properties, like being striped, or being mobile, or being at least 50 per cent metallic or . . . For being locked, being striped, etc. do not define, or contribute to defining, significant levels. But some multiple realizable properties do so: cf. claim (i) at the end of §1.

Here is a schematic biological example (my thanks to a referee). Fitness is multiply realized by the different morphological, physiological and behavioural properties of organisms. Indeed, *very* multiply realized, what makes a cockroach fit is very different from what makes a daffodil fit. Thus, fitness is a higher order or ‘more abstract’ similarity of organisms (as are its various degrees). And unlike being locked, etc., it contributes to defining a significant level: there are general truths about it and related notions to be expressed and explained. For this it is not enough to have a theory about cockroaches, and another one for daffodils, etc. Rather, we need the theory of natural selection.

In short, multiple realizability is undoubtedly important for understanding relations between levels. But many philosophers go further than this. Some think that multiple realizability provides an argument against reduction. The leading idea is that the *definiens* of a multiply realizable property shows it to be too ‘disjunctive’ to be suitable for scientific explanation, or to enter into laws. And some philosophers think that multiple realizability prompts a non-Nagelian account of reduction; even suggesting that definitional extensions cannot incorporate functional definitions.

I reject both these lines of thought. Multiple realizability gives no argument against definitional extension; nor even against stronger notions of reduction like Nagel’s, which add further constraints additional to deducibility, e.g. about explanation. That is, I believe that such constraints are entirely compatible with multiple realizability. This was shown very persuasively by Sober [8]. But as these errors are unfortunately widespread, it is worth first rehearsing, then refuting, the multiple realizability argument.

We again envisage two theories, T_b and T_t , or two sets of properties, \mathcal{B} and \mathcal{T} , defined on a set \mathcal{O} of objects. The choice between theories and sets of properties makes almost no difference to the discussion; and I shall here mostly refer to \mathcal{B} and \mathcal{T} , rather than T_b and T_t . So multiple realizability means that the instances in \mathcal{O} of some ‘top’ property $P \in \mathcal{T}$ are very varied (heterogeneous) as regards (how they are classified by) their properties in \mathcal{B} .

The *multiple realizability argument* holds that, in some cases, the instances of P are so varied that even if there is an extensionally correct definition of P in terms of \mathcal{B} , it will be so long and/or heterogeneous that:

- (a) explanations of singular propositions about an instance of P cannot be given in terms of \mathcal{B} , whatever the details about the laws and singular propositions involving \mathcal{B} ; and/or
- (b) P cannot be a natural kind, and/or cannot be a law-like or projectible property, and/or cannot enter into a law, from the perspective of \mathcal{B} .

Usually an advocate of (a) or (b) is not ‘eliminativist’, but rather ‘anti-reductionist’. P and the other

properties in \mathcal{T} satisfying (a) and/or (b) are not to be eliminated as cognitively useless. Rather, we should accept the taxonomy they represent, and thereby the legitimacy of explanations and laws invoking them. Probably the most influential advocates have been: Putnam [9] for version (a), with the vivid example of a square peg fitting a square hole, but not a circular one; and Fodor [10] for version (b), with the vivid example of $P =$ being money.

I believe that Sober [8] has definitively refuted this argument, in its various versions, whether based on (a) or on (b), and without needing to make contentious assumptions about topics like explanation, natural kind and law of nature. As he shows, it is instead the various versions of the argument that make contentious assumptions! I will not go into details. Suffice it to make three points, by way of summarizing Sober's refutation.⁴ The first two correspond to rebutting (a) and (b); the third point is broader and arises from the second.

As to (a): the anti-reductionist's favoured explanations in terms of \mathcal{T} do *not* preclude the truth and importance of explanations in terms of \mathcal{B} . As to (b): a disjunctive definition of P , and other such disjunctive definitions of properties in \mathcal{T} , is no bar to a deduction of a law, governing P and other such properties in \mathcal{T} , from a theory T_b about the properties in \mathcal{B} . Nor is it a bar to this deduction being an explanation of the law.

The last sentence of this refutation of (b) returns us to the question whether to require reduction to obey further constraints apart from deduction. The tradition, in particular Nagel himself, answers Yes (as I reported in *caveat* (a), §2.1). Nagel in effect required that the *definiens* play a role in the reducing theory T_b . In particular, it cannot be a very heterogeneous disjunction. (Recall that the *definiens* is the right-hand side of a bridge principle.) The final sentence of the last paragraph conflicts with this view. At least, it conflicts if this view motivates the non-disjunctiveness requirement by saying that non-disjunctiveness is needed if the reducing theory T_b is to explain the laws of reduced theory T_t . But I reply: so much the worse for the view. Sober puts this reply as a rhetorical question [8, p. 552]: 'Are we really prepared to say that the truth and lawfulness of the higher level generalization is *inexplicable*, just because the... derivation is peppered with the word "or"? I agree with him: of course not!

2.3. Supervenience? The need for precision

So far my main points have been that reduction in a strong Nagelian sense is compatible with both emergence (§2.1) and multiple realizability (§2.2). But these points leave open the questions how widespread is reduction, and what are the other, perhaps typical or even widespread, relations between theories at different levels. My view is that, within physics and even between physics and other sciences, reduction—at least in the approximative sense mentioned in *caveat* (b) of §2.1—is indeed widespread. I develop this view

⁴Agreed, in philosophy, there is always more to say. I do not pretend that Sober's paper is the last word on the subject: in a large literature, I recommend Shapiro [11, especially §IVf. p. 643f] and Papineau [12, especially §4, p. 183f].

in Butterfield [1, §3.1.2] and [2, §4f], partly in terms of the unity of nature (cf. opening remarks about the second scientific revolution in §1).

As to what relation or relations hold when reduction fails, philosophers' main suggestion has been: *supervenience*. Roughly speaking, this notion is a strengthening of the idea of multiple realizability. I will first explain the notion, and then state four misgivings about it.

Again we envisage two sets of properties, \mathcal{B} and \mathcal{T} , defined on a set \mathcal{O} of objects. We say that \mathcal{T} *supervenes on* \mathcal{B} (also: *is determined by* or *is implicitly defined by* \mathcal{B}) iff any two objects in \mathcal{O} that match in all properties in \mathcal{B} also match in all properties in \mathcal{T} . Or equivalently, any two objects that differ in a property in \mathcal{T} must also differ in some property or other in \mathcal{B} . (One also says that \mathcal{B} *subvenes* \mathcal{T} .) A standard (i.e. largely uncontroversial!) example takes \mathcal{O} to be the set of pictures, \mathcal{T} their aesthetic properties (e.g. 'is well-composed') and \mathcal{B} their physical properties (e.g. 'has magenta in top-left corner').

It turns out that supervenience is a weakening of the notion of definitional extension given in §2.1; namely to allow that a definition in the set D might have an *infinitely* long *definiens* using \mathcal{B} . The idea is that, for a property $P \in \mathcal{T}$, there might be infinitely many different ways, as described using \mathcal{B} , that an object can instantiate P : but provided that, for any instance of P , all objects that match it in their \mathcal{B} -properties are themselves instances of P , then supervenience will hold.

Thus, many philosophers have held that in cases where one level or theory seems irreducible to another, yet to be in some sense 'grounded' or 'underpinned' by it, the relation is in fact one of supervenience. They say that the irreducible yet grounded level or theory (specified by its taxonomy of properties \mathcal{T}) supervenes on the other one. That is, there is supervenience without definitional extension: at least one definition in D is infinitely long.

At first sight, this looks plausible: recall from the start of §2.2 that examples of multiple realizability are legion. But we should note four misgivings about it. The first two are widespread in the literature; the third and fourth are more my own. The first and fourth are philosophical limitations of supervenience; the second and third, scientific limitations.

First: philosophers of a metaphysical bent who discuss reduction, emergence and related topics find it natural to require that, in reduction, the 'top' properties \mathcal{T} are shown to be identical to properties in (or perhaps composed from) \mathcal{B} ; and that this is so whether the reduction is finite, as in definitional extension, or infinite, as in supervenience. But the identities of properties (and the principles for composing properties) are controversial issues in metaphysics; and the holding of a supervenience relation is not generally agreed to imply identity. So, for such philosophers, supervenience leaves a major question unanswered.

Second: although the distinction between finite and infinitely long definitions is attractively precise, it seems less relevant to the issue whether there is a reduction than another, albeit vague, distinction: namely, the distinction between definitions and deductions that are short enough to be comprehensible, and those that are not. Recall that, according to the notion of a definitional extension given in §2.1, a definition in D

can be so long as to be incomprehensible, e.g. a million pages—to say nothing of the length of the deductions!

Thus, the remarks usually urged to show a supervenience relation in some example, e.g. that no one knows how to construct a finite definition of ‘is well-composed’ out of ‘has magenta in top-left corner’ and its ilk, are not compelling. Our inability to complete, or even begin, such a definition is no more evidence that a satisfactory definition would be infinite than that it would be incomprehensibly long. In other words, we have no reason to deny that the example supports a definitional extension, albeit an incomprehensibly long one. And so far as science is concerned, definitional extension with incomprehensibly long definitions and deductions is useless: that is, it may as well count as a failure of reduction. Philosophers, including Nagel himself, have long recognized this point: recall the *caveat* (a) in §2.1.⁵

The third misgiving is similar to the second, in that both accuse supervenience of having—for all its popularity in philosophy—limited scientific value. But where the second sees supervenience’s allowance of infinite disjunctions as a distraction from the more important issue of comprehensibility, the third sees supervenience’s allowance of infinite disjunctions as a distraction from the more important issue of the limiting processes that occur in the mathematical sciences, and in particular in examples of emergence in physics. That is, because supervenience’s infinity of ‘ways (in terms of \mathcal{B}) to be $P \in \mathcal{T}$ ’ bears no relation to the taking of a limit (e.g. through a sequence of states, or of quantities, or of values of a parameter), it sheds little or no light on such limits, in particular on the emergent behaviour that they can produce. Agreed, this sort of accusation can only be made to stick by analysing examples: suffice it to say here that Butterfield [2] analyses four such.

The fourth misgiving concerns philosophers’ appeal to supervenience, not as a relation between two independently specified levels or theories, but as a tool for precisely formulating physicalism: the doctrine that, roughly speaking, all facts supervene on the physical facts. Here, my complaint is for physicalism to be precise, you need to state precisely what are ‘the physical facts’ (or what is ‘the physical supervenience basis’). Sad to say, in the philosophical literature, both proponents and opponents of physicalism tend to be vague about this. Here is one example which has been discussed widely (more details in [1, §5.2.2]). (I also recommend Sober’s very original discussion of how the

definition of, and our reasons for, physicalism are usefully cast in terms of probability, especially the Akaike framework for statistical inference [16].)

If there had been fundamental many-body forces (called by C. D. Broad, the British emergentist of the early twentieth century, ‘configurational forces’), then the theory of a many-body system would not be supervenient on (let alone a definitional extension of) a theory of its components that used only two-body interactions.⁶ Of course, if there had been such forces, physicists would have made it their business to investigate them, so that a phrase like ‘the physical facts’ would have come to include facts about such configurational forces, as well as facts about the familiar two-body forces. At least it would have come to include such facts if the configurational forces turned out to fit into the familiar general frameworks of (classical or quantum) mechanics, e.g. having a precise quantitative expression as a term in a Hamiltonian. But this says more about the elasticity of the word ‘physics’, or about universities’ departmental structure, than about the truth of a substantive doctrine of ‘physicalism’!

2.4. Causation

So far, I have ignored issues about time-evolution, and in particular causation: I have stressed what one might call ‘synchronic issues’, rather than ‘diachronic issues’. But, from now on, diachronic issues will be centre-stage. As I mentioned in §1, I take a broadly Humean view of causation, but do not advocate a specific account. Nor will I need such an account for the rest of this paper’s aims. There are three such aims. In this subsection, I will report and recommend two recent arguments broadly in favour of top-down causation. My final aim, in later sections, will take longer: it is to describe Ellis’ types of top-down causation, in terms of functional dependence (cf. claims (ii) and (iii) at the end of §1).

Of course, there is much to say about top-down causation apart from what follows in the rest of this paper; and even apart from my fellow symposiasts—in a large literature, I recommend Bedau [14, pp. 157–160, 175–178]. And I cannot trace the consequences of what follows, for other authors’ views. But I commend what follows to advocates of top-down causation, such as that of Ellis and Atmanspacher, Auletta, Bishop, Jaeger and O’Connor. For I think it makes precise some of their claims: such as that higher level facts or events constrain, modify or form a context for the lower level, which is therefore not independent of the higher level [22]; or that lower level facts or events are necessary but not sufficient for higher level ones [21].

So I turn to reporting and recommending the two arguments. The first is Shapiro & Sober’s argument, not so much *for* top-down causation, as *against* a contrary position, namely epiphenomenalism. This is

⁶Here I applaud Scerri’s [17, pp. 3–4] critique of Hendry’s surely maverick view [18, §3] that configurational forces *are* needed in modern quantum chemistry to explain molecular chirality and shape. On the contrary, I take it to be well established that the explanation lies in superselection (classical quantities) being rigorously emergent in appropriate limits (e.g. Primas [19], Amann [20] and Bishop & Atmanspacher [21]). So this is another example of the reconciling claim as in §2.1 that emergence and reduction are compatible.

⁵The idea of incomprehensible definitions and deductions, and thereby the need for higher level concepts and laws, is often illustrated with cellular automata such as Conway’s Game of Life, with, for example, ‘glider’ as a higher level concept (cf. Dennett [13, pp. 196–200], Bedau [14, pp. 164–178] and O’Connor [15, pp. 1–3]). The idea is: these concepts and laws are definable and deducible from Life’s basic rules, but only by processes so grotesquely long that you would be ill-advised—mad!—to try and follow them, rather than investigating the higher level behaviour directly. Besides: a theorem in logic (Beth’s theorem) shows that, under certain conditions (namely first-order finitary languages), the finite–infinite distinction collapses in the sense that if every term of T_t is implicitly definable in T_b , then T_t is a definitional extension (i.e. with finite definitions) of T_b . This point was first emphasized by Hellman and Thompson; more details are in Butterfield [1, §5.1].

the doctrine that higher level states cannot be causes, i.e. they are causally ineffective. Thus, the idea of epiphenomenalism is that such states are pre-empted, as causes, by lower-level states. (Here, we could say ‘property’, ‘fact’ or ‘event’, instead of ‘state’: it would make no difference to what follows.) Shapiro & Sober rebut this, by adopting an account of causation in terms of intervention (or, in another jargon: in terms of manipulation). On the other hand, the second argument is List & Menzies’ positive argument for top-down causation; it is based on an account of causation in terms of counterfactuals. Fortunately, both accounts of causation are plausible, and I will not need to choose between them. Nor will I need to develop the accounts’ details. For the argument by each pair of authors needs only the basic ideas of the account.

I can present both arguments in terms of the same example of a pair of levels: the mental and the physical. Both arguments are very general, and apply equally to other examples of pairs of levels. But this example has various advantages. It is vivid and widely discussed in philosophy. All these four authors use it. Although Shapiro & Sober also discuss higher level causation in biology, especially evolutionary biology, for List & Menzies, this example is the main focus. So their central case of top-down causation is mental causation, e.g. my deciding to raise my arm causing it to go up. Besides, all these authors rebut various formulations and defences of epiphenomenalism about the mental with respect to the physical by Kim, who is probably the most prolific recent writer on the mental–physical relationship.

Thus here, in terms of the mental and the physical, is what Shapiro & Sober [23] call ‘the master argument for epiphenomenalism’:

How could believing or wanting or feeling cause behavior? Given that any instance of a mental property M has a physical micro-supervenience base P , it would appear that M has no causal powers in addition to those that P already possesses. The absence of these additional causal powers is then taken to show that the mental property M is causally inert.

([23, p. 241]; notation changed)

In addition to Kim’s formulations of this argument that Shapiro & Sober go on to document, compare Kim [24, p. 19]. The argument is a cousin of what is often called ‘the exclusion argument’, also often advocated by Kim: for a discussion, see, for example, Humphreys [25].

Assessing this argument depends of course on one’s account of causation. (And, one might guess, on one’s account of realization or supervenience—but, in fact, the varieties of these notions turn out not to matter.) But the argument fails utterly, on each of two plausible accounts of causation: the interventionist account adopted by Shapiro & Sober, and the counterfactual account adopted by List & Menzies. And the failure follows from just the basic features of each pair of authors’ account.

Thus the leading idea of Shapiro & Sober’s rebuttal is:

To find out whether M causally contributes to N , you manipulate the state of M while holding fixed

the state of any common cause C that affects both M and N ; you then see whether a change in the state of N occurs. . . . *It is not relevant, or even coherent, to ask what will happen if one wiggles M while holding fixed the micro-supervenience base P of M* Because a supervenience base for M provides a sufficient condition for M , where the entailment has at least the force of nomological necessity, asking this question leads one to attempt to ponder the imponderable—would N occur if a sufficient condition for M occurred but M did not?

([23, pp. 238–240]; notation changed)

Besides, Shapiro & Sober strengthen this rebuttal by augmenting their account of causation with probabilities (see also Sober [16, pp. 145–149]). There is a happy concordance here between interventionist and probabilistic accounts of causation. But, again, I will not need to give details.

I turn to List & Menzies [26] and Menzies & List [27]. In short, they make two points: (i) a general point, in common with other authors, which supports the idea of causation occurring at higher levels, and between levels (cf. claim (i) at the end of §1), and (ii) the specific argument in favour of top-down causation, and against epiphenomenalism. Broadly speaking, their first point supports the idea that we should understand causation in terms of counterfactuals (especially ‘if C had not occurred, then E would not have occurred’). On the other hand, their specific argument assumes some such counterfactual account.⁷

The general point is that a cause needs to be specific enough to produce its effect—but not more specific. The point is clear from how we think about causes in countless examples. What caused the bull to be in a rage? Answer: the bull’s seeing the matador’s red cape nearby. If the cape happens to be crimson and the matador to be standing 3 m from the bull, nevertheless, the cause is as stated. It is not the more specific fact of the bull’s seeing the matador’s crimson cape 3 m away.

Nor is this point just a matter of our intuitive verdicts in countless examples. It is upheld by plausible accounts of causation, in terms of counterfactuals (stemming from Lewis [28]). The key idea is that C ’s causing E requires that if C were not to hold, then E would not hold. And, indeed, if the cape were not red, but say green, then the bull would not be enraged. But if C is too specific, this requirement tends to fail. We cannot conclude that, if the cape were not crimson, the bull would not be enraged—for if the cape were not crimson, it might well have been some other shade of red, and then the bull would still have been enraged.

This point applies in countless examples where the contrast between appropriate and too-specific causes corresponds to a contrast between levels, such as the mental and the physical. Witness the bull example (which is adapted from Yablo [29]); or in the other ‘direction’, a mental state causing a physical one, as

⁷But, again, that is a small price to pay, as such accounts are plausible, especially when compared with the traditional rival idea (called ‘nomological sufficiency’) that a cause, taken together with the laws, should imply the effect.

in the time-honoured example of arm-raising. What caused my arm to go up? My deciding to raise it.

To sum up: this point teaches us that more specific information about the facts and events in some example is *not* always illuminating about causal relations; and (*a fortiori*), that it is wrong to think that the ‘best’ or ‘real’ causal explanation of the facts and events always lies in the most specific and detailed information about them. (Of course, it is ‘reductionists’ rather than others who are most likely to suffer these bad temptations.) And so this point supports the idea that there are causal relations between facts or events at higher levels, or between different levels.⁸

List & Menzies build on this general point so as to refute the ‘master argument’—that a mental state M cannot cause a later physical, say neural, state N , as its realizing or subvening neural state P pre-empts it as a cause. They show that, on plausible accounts of causation invoking counterfactuals, the argument fails. Indeed, they state a causal assumption about how M causes N , which in many actual cases we have good reason to believe true, and which implies that P is *not* a cause of N . (This assumption is that, even if M were realized by a physical state other than its actual realizer P , N could still obtain. List & Menzies’ jargon is that M ’s causing N is ‘realization-insensitive’; see [26, §V].)

To sum up, here is excellent news for advocates of top-down causation, at least if they like interventionist or counterfactual accounts of causation. More power to these authors’ elbows. Or, expressed in the spirit of their results: more power to their wills, so as to cause their elbows to move!

3. DYNAMICS AT DIFFERENT LEVELS

I now turn to this paper’s second main aim. In this section, I give a framework for describing dynamics at different levels, emphasizing how two levels’ dynamics can mesh or fail to mesh. Section 4 then applies the framework to describe some of Ellis’ [3] types of top-down causation.

There will of course be some connections with topics addressed in §2. Here are two examples. (i) Section 3.2 describes how Papineau’s critique of Fodor’s version of the multiple realizability argument (see §2.2) is a matter of two levels’ dynamics failing to mesh. (ii) I admit at the start of §4 that my descriptions of Ellis’ types are partial, because they avoid controversies about what is required for causation, beyond functional dependence of quantities (cf. claims (ii) and (iii) at the end of §1).

3.1. The framework introduced

For simplicity, I will work with just two levels, dubbed ‘micro’ and ‘macro’, which are related by coarse-graining. There will be several other simplifying assumptions, as follows.

- (i) We think of the micro-level as a state space \mathbb{S} , with the micro-dynamics as a map $T: \mathbb{S} \rightarrow \mathbb{S}$

⁸Agreed, it does not by itself imply such relations. But nor should we expect a general principle about causation to dictate on such specific matters as between which levels there are causal relations.

(so time is discrete). Since T is a function, we assume a past-to-future micro-determinism.

- (ii) But T need not be invertible, so future-to-past determinism can fail. Besides, most of what follows applies if T is just a binary relation, i.e. one-many as well as many-one, so that there is past-to-future micro-*indeterminism*. (This indeterminism need not reflect quantum mechanics—under the orthodox interpretation! It could reflect the system being *open*.)
- (iii) We will not need to discuss in detail the set \mathbb{Q} of quantities. But we envisage that each $Q \in \mathbb{Q}$ is a real-valued function on \mathbb{S} . For a classical physical state s , $Q(s)$ would be thought of as the system’s intrinsic or possessed value for Q when in s . But, in quantum theory, $Q(s)$ would naturally be taken to be a Born-rule expectation value, e.g. s is a Hilbert space vector, and $Q(s) := \langle s | \hat{Q} | s \rangle$. In either classical or quantum physics, we naturally think of \mathbb{Q} as separating \mathbb{S} in the sense that, for any distinct $s_1 \neq s_2 \in \mathbb{S}$, there is a $Q \in \mathbb{Q}$ such that $Q(s_1) \neq Q(s_2)$.
- (iv) We think of the macro-level as given by a partition P of \mathbb{S} , i.e. a decomposition of \mathbb{S} into mutually exclusive and jointly exhaustive subsets $C_i \subset \mathbb{S}$. The C_i , i.e. the cells of the partition $P = \{C_i\}$, are the *macro-states*. So C stands for ‘cell’ or ‘coarse-graining’.
- (v) Filling out (ii)–(iv), we naturally think of each C_i as specified by a set of values for some subset of \mathbb{Q} , the macro-quantities $\mathbb{Q}_{\text{mac}} \subset \mathbb{Q}$. That is, each C_i is the intersection of level-surfaces, one for each quantity in \mathbb{Q}_{mac} . But we can downplay quantities, and just focus on the macro-states C_i , i.e. the cells of the partition P .

3.2. Meshing of dynamics: examples and counterexamples, in physics and philosophy

I now consider the way in which the dynamics at the two levels can mesh with each other, or fail to do so. Physics provides precise and important examples of such ‘meshing’ (e.g. in statistical mechanics), as well as examples where it fails. I also relate meshing to the views of Fodor & Papineau on multiple realizability, and to the views of List on free will.

The framework, especially assumptions (i) and (iv) of §3.1, yields natural definitions of:

- (a) how the micro-dynamics T defines a macro-dynamics, and
- (b) whether these two dynamics mesh.

As to (a), recall that any function $f: X \rightarrow Y$ between any two sets X and Y defines a function $\bar{f}: \mathcal{P}(X) \rightarrow \mathcal{P}(Y)$ between their power sets, by the obvious rule $A \subset X \mapsto \bar{f}(A) := \{y \in Y \mid y = f(x) \text{ for some } x \in A\} \subset Y$. We apply this idea to $T: \mathbb{S} \rightarrow \mathbb{S}$, and then restrict \bar{T} to the partition $P = \{C_i\} \subset \mathcal{P}(\mathbb{S})$. In general, the image $\bar{T}(C_i)$ of a cell C_i is *not* a subset of a single macro-state. That is, two distinct $s_1 \neq s_2 \in C_i$ are sent by T to distinct macro-states. In other words, we have macro-*indeterminism*, despite T giving micro-determinism.




Figure 1. Micro-determinism induces indeterministic macro-dynamics.

The micro- and macro-dynamics do *not* mesh; see figure 1, where the union symbol, \cup , beside some of the upward lines indicates coarse-graining.

On the other hand, as to (b): if for every cell $C_i \in P$, its image $\bar{T}(C_i)$ is a cell in P , $\bar{T}(C_i) = C_j$ for some j , then the micro-dynamics T *does* induce a deterministic macro-dynamics, and we will say that the micro- and macro-dynamics *mesh*. So in figure 1, the upward \cup arrow from $T(s_1)$ would not be ‘stray’, but would point to C_j . In mathematical jargon, coarse-graining and time-evolution *commute*. Another jargon: coarse-graining is *equivariant* with respect to the group actions representing time-evolution (i.e. for discrete time: actions of the group of integers \mathbb{Z} , on \mathbb{S} and on P).

Physics provides many important examples of meshing dynamics. The most obvious examples concern conserved quantities, especially in integrable systems. Recall from (iii) and (v) of §3.1 the idea of macro-quantities $\mathbb{Q}_{\text{mac}} \subset \mathbb{Q}$, i.e. the idea that each cell C_i is the intersection of level-surfaces, one for each quantity in \mathbb{Q}_{mac} . If we instead define $\mathbb{Q}_{\text{cons}} \subset \mathbb{Q}$ as those quantities whose values are constant under time-evolution T (which could now even be indeterministic), then, obviously, the intersection of the level-surfaces of these conserved quantities will be invariant under T . So there is meshing dynamics, though the macro-dynamics is trivial, i.e. every quantity we are concerned with is constant in value. And one could go on to investigate the ‘integrable’ systems for which \mathbb{Q}_{cons} is ‘rich enough’, in the sense that these intersections are the minimal sets invariant under time-evolution; see also §3.3.

But there are important examples of meshing, when the system is not integrable, and even has only a few conserved quantities. Indeed, that is putting it mildly! Several of the most famous and fundamental equations of macroscopic physics (such as the Boltzmann, Navier–Stokes and diffusion equations) are the meshing macro-dynamics induced by a micro-dynamics. Or rather: they are the meshing macro-dynamics once we make the framework of §3.1 more realistic by allowing that:

- the meshing may not last for all times;
- the meshing may apply, not for all micro-states s , but only for all except a ‘small’ class;
- the coarse-graining may not be so simple as partitioning \mathbb{S} ; and indeed
- the definition of the micro-state space \mathbb{S} may require approximation and/or idealization, especially by taking a limit of a parameter: in particular, by letting the number of microscopic constituents tend

to infinity, while demanding of course that other quantities, such as mass and density, remain constant or scale appropriately.

This point, especially (d), returns us to the claim in §2.1 that emergent (i.e. novel, robust) behaviour may be deduced from a theory of the microscopic details, often by taking a limit of some parameter. As I said there, I gave four such examples [2]. The equations just listed—Boltzmann, etc.—provide others.⁹

Having emphasized dynamical meshing, not least for its philosophical importance in reconciling emergence and reduction, I should add that, on the other hand, *failure* of meshing—in mathematical jargon, the non-commutation of the diagram in figure 1; in physical jargon, macro-indeterminism—need not pose any difficulties or ‘worries’ for our understanding the situation. For there can be other factors that make the non-commutation, the macro-indeterminism, well understood, and even well controlled and natural.

For a physicist, the obvious and striking example of this is the pilot-wave theory in the foundations of quantum mechanics as follows. The deterministic evolution T on the pilot-wave micro-states s (comprising, for example, positions $\mathbf{q}_i \in \mathbb{R}^3$ of corpuscles $i = 1, \dots, N$, as well as an orthodox quantum state ψ) induces an *indeterministic* evolution on ψ by a precise version of the textbooks’ projection postulate rule that, when ψ divides into disjoint wave-packets, it is to be replaced by whichever of the (renormalized) packets has support containing the actual configuration $\langle \mathbf{q}_i \rangle \in \mathbb{R}^{3N}$. Besides, a mathematically natural probability measure on the micro-states s (dependent on the quantum state ψ : namely, $|\psi|^2$) makes the macro-indeterministic dynamics probabilistic, with the induced probabilities being the orthodox Born-rule probabilities (which are empirically correct for myriadly many kinds of experiments). So, in this example, the macro-indeterminism is entirely understood, well-controlled and natural (cf. Bohm & Hiley [33, especially ch. 3] and Holland [34, especially ch. 3]).

For a philosopher also, there is an obvious and striking example of macro-indeterminism, i.e. non-commutation of figure 1. Namely, the venerable idea of free will (ah, the joys of interdisciplinarity!). More precisely, a philosopher attracted by compatibilism—i.e. the view that determinism and free will are compatible—can take this sort of macro-indeterminism induced by a deterministic micro-dynamics to be exactly what free will involves. Or, still more precisely, what free will could be taken to involve in a world governed by such a micro-dynamics. For a full defence of this version of compatibilism (including the discussion of such background assumptions as non-reductive physicalism), I recommend List [35].

On the other hand, I agree that, in these less well-defined philosophical contexts about the relations between levels, failure of meshing can be ‘worrying’. One example of such a worry is Papineau’s critique

⁹Uffink & Valente [30] expound and assess the principal deduction of the Boltzmann equation, namely Lanford’s theorem. This example is especially topical in view of Villani’s 2010 Fields medal (cf. Ambrosio [31] and Yau [32]).




Figure 2. Fodor's hope: meshing dynamics.

[12, pp. 180–185] of Fodor's [10] vision of special sciences as autonomous. Thus, I agree that Papineau is right to press failure of meshing as a problem for Fodor. He argues that Fodor just assumes without justification that there will be meshing: Fodor's discussion appeals to a diagram like figure 2, which (with a trivial change from my notation) pictures the micro-laws as indeed preserving the macro-categories, i.e. as inducing a well-defined dynamics by coarse-graining.¹⁰

Papineau goes on (p. 186f) to argue—surely rightly—that in many cases, in everyday life and the special sciences, the meshing shown in figure 2 is secured by the macro-categories being defined by *selection processes*.

I take up the topic of selection in the discussion of Ellis in §4. For the moment, I just note that Ellis also addresses the idea of meshing dynamics. He calls the commutation we see in figure 2 'coherent higher level dynamics', 'effective same level action' and 'the principle of equivalence of classes' [3, pp. 7–8 and fig. 3b, p. 45]. And, at least as I read him, his typology of five kinds of top-down causation presupposes such meshing. So, before turning to that typology, it is appropriate to discuss, albeit briefly, how one might secure such meshing (see §3.3).

3.3. Meshing secured by redefining the macro-states

One response to the failure of meshing is to redefine the macro-states, so as to secure it. I shall briefly consider this response from a general, and so mathematical, perspective. This will amount to considering how to get macro-determinism 'by construction': or as one might gloss it less charitably, 'by fiat!' So I should emphasize that, in a real scientific context, one would usually invoke considerations about how to respond, which are much more specific than the general ideas (like taking unions of the cells of the given partition) which I now mention. For example, one might invoke considerations like those in §3.2: about conserved quantities, or about allowances (a)–(d), or about selection processes, as discussed by Papineau.

I will make three comments. (1) The first is a 'false start'. (2) The second is a response to the false start;

¹⁰Agreed, Fodor's and Papineau's jargon is different from mine. Both authors talk about laws, and/or causal relations, within each of the two levels, not about dynamics; and about realization or supervenience as the relation between the laws and their instances, not about coarse-graining. But I submit these are the only, or almost only, differences of jargon.

Another example of concern over failure of meshing in the context of the mental–physical relationship is Atmanspacher's discussion of the emergence of mental states from neurodynamics ([36, §§4.1, 6]; cf. Atmanspacher & beim Graben [37, §2]).

and (3) the third is a look at what one might call the 'converse' to failure of meshing. All three will be recognizable as, in effect, some of the first steps in anyone's study of discrete-time dynamical systems.

3.3.1. A false start

Given a cell C_i that 'gets broken up' by the time-evolution T , there seem at first sight to be two possible tactics for changing the partition so as to secure meshing. But each runs into difficulties.

- (i) We might define a coarser notion of macro-state (a smaller partition with larger cells) by taking the union of the cells that contain images $T(s_1)$, $T(s_2)$ of states s_1, s_2 in C_i that get sent to different cells. That is, writing as usual TC_i for the image-set, i.e. $TC_i = \{s' | s' = Ts, \text{ some } s \in C_i\}$, we might define

$$[TC_i] = \bigcup_{C_j \cap TC_i \neq \emptyset} C_j. \quad (3.1)$$

However, the sets TC_i , as C_i runs through P , are not a partition of \mathbb{S} , as different TC_i can overlap. So to get a partition with a meshing dynamics, we have to define a *chain* of overlapping sets TC_i , i.e. a sequence $\langle TC_0, TC_1, \dots, TC_N \rangle$, with $TC_0 \cap TC_1 \neq \emptyset$, $TC_1 \cap TC_2 \neq \emptyset$, etc; and then define the partition whose cells are maximal chains of image-sets TC_i . This partition has, by construction, a meshing dynamics. But it is liable to be 'uninformative': that is, the unions of maximal chains of image-sets TC_i are liable to be large.

- (ii) We might define a finer notion of macro-state (a larger partition with smaller cells) by decomposing the cell C_i that gets broken up by T into subsets, according to which cells its elements $s \in C_i$ get sent to. That is, we might define, for each j in the index set of the given partition P ,

$$(TC_i)_j = C_i \cap T^{-1}(C_j); \quad (3.2)$$

and then consider the decomposition of C_i into its subsets $(TC_i)_j$. Then the sets $(TC_i)_j$, as i, j run through the index set of P , obviously form a partition of \mathbb{S} (perhaps with, harmlessly, various 'copies' of the empty set, i.e. cases where $(TC_i)_j = \emptyset$). And since this partition has smaller cells than P did, this tactic is not in danger of being uninformative in the way that the first tactic, in (i), was. But now the trouble is that this partition need not have a meshing dynamics. For though indeed T maps all of $(TC_i)_j$ into C_j , T need not map all of $(TC_i)_j$ into some cell of the partition we have just defined, i.e. into some $(TC_k)_l$.

3.3.2. A response

Rather than starting from a given partition P with a non-meshing dynamics, and asking how to modify it so as to get meshing, it is easier to start with just the micro-dynamics and consider defining *ab initio* a partition with a meshing dynamics. Indeed, it is easy to address the stronger question of defining cells each of which is invariant under the dynamics, i.e. mapped into themselves by T .

Thus, for any state $s \in \mathbb{S}$, the set

$$[Ts] := \{s' \mid \text{there is a chain } s_0 = s, \\ s_1 = Ts, \dots, s_n = Ts_{n-1}, \dots, s_N = s'\} \quad (3.3)$$

is obviously the smallest set invariant under T that contains s . This statement also holds true if T is not a function, but one-many, i.e. the time-evolution is indeterministic (so that many chains could start at s). If T is a function, then $[Ts]$ either is a cycle, i.e. $s \mapsto s_1 \mapsto s_2 \mapsto \dots \mapsto s$, or is an ω -sequence.

By containing only ‘descendants’ of s , definition of $[Ts]$ of equation (3.3) obviously favours the ‘initial time’. If T is a function, we can avoid this favouritism by instead defining

$$[[Ts]] := \{s' \mid \text{either there is a chain } s_0 = s, \\ s_1 = Ts, \dots, s_n = Ts_{n-1}, \dots, s_N = s', \\ \text{or there is a chain } s_0 = s', s_1 = Ts', \dots, \\ s_n = Ts_{n-1}, \dots, s_N = s\}, \quad (3.4)$$

which is obviously the smallest set invariant under T that contains both s and any of its ‘ancestors’. If T is not a function, then we need to allow for chains from ancestors of s that do not pass through s . Then

$$[[[Ts]]] := [[Ts]] \cup \bigcup_{s' \in [[Ts]]} [Ts'] \quad (3.5)$$

is by construction the smallest set invariant under T that contains both s and any of its ‘ancestors’.

3.3.3. The converse scenario

So much by way of discussing the failure of meshing, i.e. the scenario in figure 1, and how one might respond to it by redefining the macro-states. That scenario also prompts one to consider the converse scenario: that is, micro-indeterminism inducing, by coarse-graining, macro-determinism, as in figure 3.

Indeed, this scenario can happen in a tightly defined theoretical context, in a scientifically important way. This means in particular: the macro-determinism is not ‘won on the cheap’, by having very large cells (in other words, by using a logically weak taxonomy of micro-states). Brownian motion provides an example. Think of the Langevin equation’s probabilistic description of Brownian motion. Different realizations of the noise (i.e. different, unknown, trajectories of the atoms bombarding the large Brownian particle) give the particle different spatial trajectories: micro-indeterminism. But averaging over the realizations gives a deterministic evolution of the probability density for the particle’s position (this evolution is given by a Fokker–Planck equation).¹¹

¹¹Probability theory also provides similar examples where time-evolution is downplayed, but which retain the key idea of combining variety of micro-states with macroscopic uniformity. One main example is the method of arbitrary functions; where the probability of a macro-state is approximately the same for any of a wide class of probability density functions f on the micro-states $s \in \mathbb{S}$, essentially because the partition P defining the macro-states is very intricate. Butterfield [2, §4] gives more details, including a discussion of emergence.




Figure 3. Micro-indeterminism induces deterministic macro-dynamics.

4. ELLIS’ TYPES OF TOP-DOWN CAUSATION

I turn to using the framework of §3 to describe, in part, the first four of Ellis’ five types of top-down causation [3,4]. I think these descriptions are worth formulating, because they are precise. But I say ‘in part’, because this precision comes with a price-tag; indeed two price-tags.

First, I will face several choices about how best to formalize Ellis’ ideas. For clarity and brevity, I will make simple choices: even to the extent of *not* representing micro-states and micro-dynamics, and so also not representing top-down causation. But I shall suggest that my choices are innocuous: one can see how one could elaborate the descriptions so as to represent micro-states, top-down causation, etc.

But second, and perhaps more important, I will also simplify by setting aside non-formal, indeed philosophical, questions about what, beyond ideas like functional dependence and coarse-graining, is needed for causation, in particular top-down causation. (Recall my concessions (ii) and (iii) at the end of §1.) For example, I set aside, the distinction, made by Ellis [3, p. 6, p. 8 et seq.] and Auletta *et al.* [22], between causal effectiveness and causal power; and their taking top-down causation to require the macro-level to have causal power over the micro-level. And apart from the views of Ellis and kindred spirits like Jaeger & Calkins, I will not try to connect my descriptions of Ellis’ third and fourth types, which concern adaptation and evolution and so biology, to biological details, or to this Theme Issue’s other relevant authors [38–40].

These simplifications are evident already in my proposed description of Ellis’ first type, which he calls *algorithmic top-down causation*. Ellis writes [3, p. 8]: ‘algorithmic top-down causation occurs when high-level variables have causal power over lower level dynamics through system structuring, so that *the outcome depends uniquely on the higher level structural, boundary and initial conditions*’. In line with my simplifications, I will take this quotation to mean *just* meshing dynamics, or commutation, as in §3.2: the ‘Fodor’s hope’ of figure 2. For meshing dynamics matches the quotation’s idea of a higher level outcome being uniquely determined by higher level facts. But that this is indeed a simplification is clear from Ellis also calling meshing dynamics ‘coherent higher level action emerging from lower level dynamics’ and ‘the principle of equivalence of classes’ (*ibid.*), and his taking it as a presupposition of his typology of top-down causation.

I will devote §4.1–4.3 to describe each of Ellis’ second, third and fourth types. For clarity, I will begin each subsection with Ellis’ own description [3] of the type. For brevity, I omit his fifth type (called ‘intelligent top-down causation’), which involves the use of symbolic representation to investigate the outcome of goal choices. But, by the end of my description of his fourth type (§4.3), it will be clear how the fifth type might be partially described with my framework.

4.1. Top-down causation via non-adaptive information control

Ellis writes [3, p. 12]: ‘in non-adaptive information control, higher level entities influence lower level entities so as to attain specific fixed goals through the existence of feedback control loops, whereby information on the difference between the system’s actual state and desired state is used to lessen this discrepancy . . . unlike (algorithmic top-down causation), *the outcome is not determined by the boundary or initial conditions; rather it is determined by the goals*’.

I will represent this notion of control in my framework using two simple ideas (which will be useful later): which I will call *culling* and *conforming*. In both culling and conforming, some macro-state C^* (cell C^* of the partition P) is designated as ‘desired’. The difference will be that, in culling, one simply sends to some ‘rubbish’ or ‘dead’ state, 0 say, all states other than the desired one C^* . So this will be modelled with a characteristic function χ_{C^*} (though with the ‘yes’ or ‘winning’ value being C^* itself, rather than 1). In conforming, on the other hand, states other than the desired one are sent to the desired one: so this will be modelled by a constant map with value C^* . The details are as follows, with due precision about the difference between maps at the macro- and micro-levels.

Culling. Writing 0 for the rubbish state at the macro- or micro-level, and adjoining the rubbish macro-state to the partition P so as to define a partition $P^0 := P \cup \{0\}$, we have:

- (a) at the macro-level: a characteristic map $\chi_{C^*} : P \rightarrow P^0$; with $\chi_{C^*}(C_i) := 0$, unless $C_i = C^*$ in which case $\chi_{C^*}(C_i) := C^* \equiv C_i$;
- (b) at the micro-level: a characteristic map $\chi_{C^*} : \mathbb{S} \rightarrow \{\mathbb{S}, 0\}$, or if one prefers $\chi_{C^*} : \mathbb{S} \rightarrow \{C^*, 0\}$, with a similar ‘culling’ definition.

Conforming. We have:

- (a) at the macro-level: a constant map $\kappa_{C^*} : P \rightarrow P$ sending all cells in P to C^* : $\kappa_{C^*}(C_i) := C^*$, for all C_i ;
- (b) at the micro-level: *any* micro-dynamics given by a function (or indeed relation, i.e. multi-valued function) T on \mathbb{S} that sends all of \mathbb{S} into C^* . That is, any T such that $T(\mathbb{S}) \subset C^*$ will induce $\kappa_{C^*} : P \rightarrow P$ as its meshing macro-dynamics.

Various combinations and liberalizations of these two ideas are possible. The simplest combination is to conform and then cull: $\chi_{C^*} \circ \kappa_{C^*}$. Then no $s \in \mathbb{S}$ ends up as ‘rubbish’. Instead, all $s \in \mathbb{S}$ get sent to a

micro-state in the desired macro-state C^* . This is a perfect control, with the system ending in the desired macro-state, whatever its initial $s \in \mathbb{S}$.

4.2. Top-down causation via adaptive selection

Ellis writes [3, pp. 14–15]: ‘Adaptive processes take place when many entities interact . . . for example individuals in a population, and *variation takes place in the properties of these entities, followed by selection of preferred entities that are better suited to their environment or context*. Higher level environments provide niches that are either favourable or unfavourable to particular kinds of lower level entities; those variations that are better suited to the niche are preserved and the others decay away. Criteria of suitability in terms of fitting the niche can be thought of as fitness criteria guiding adaptive selection. On this basis, a *selection agent* or *selector* accepts one of the states and rejects the rest; this selected state is then the current system state that forms the starting basis for the next round of selection Thus, this is top-down causation from the context to the system. An equivalence class of lower level variables will be favoured by a particular niche structure in association with a specific fitness criteria. Unlike feedback control, this process does not attain preselected internal goals by a specific set of mechanisms or systems; rather it creates systems that favour the meta-goals embodied in the fitness criteria’.

In representing these ideas, one of course faces several choices about how much detail to represent explicitly. For example, should one represent explicitly:

- (i) the environment/context (even niches?), or just the adapting entities;
- (ii) as regards these entities: individuals (and their variation?), or just the population;
- (iii) as regards the environment and the entities: micro-states, or just macro-states;
- (iv) selection as a process with several possible outcomes, or just two (survival or death!), as in the notion of culling in §4.1;
- (v) several rounds of selection, or just one;
- (vi) adaptation using the notion of conforming as in §4.1; and
- (vii) adaptation *within* a round of selection, e.g. in the lifetime of an individual, or just over many rounds?

In answering these questions, I propose to keep things pretty simple. In terms of this list, I will represent explicitly:

- (i) the environment;
- (ii) individuals and their variation; and
- (v) several rounds of selection.

But I will *not* represent:

- (iii) micro-states;
- (iv) outcomes of selection other than survival or death, as in culling; and
- (vi) and (vii) adaptation, in terms of conforming, in a single round or over many rounds.

But I admit that these are just *choices* of what to represent: no doubt, several other possible choices are equally (or more) useful.

As to (i), the environment: I will be simplistic. I take the environment to be unchanging, so that there is no co-evolution. There is an environment state-space \mathbb{S}^e on which there is a partition $P^e = \{C_k^e\}$. Niches will be represented only implicitly; namely, by the way in which the macro-states of the individuals (and so of the population) that are selected for depend on the macro-state C_k^e of the environment.

As to (ii), the individuals and their variation: again, I will be simplistic. I assume that, in each round of selection, there are N individuals, each with an individual state-space \mathbb{S} on which there is a partition $P = \{C_i\}$. So the population of N individuals has a Cartesian product state-space \mathbb{S}^N (neglecting the tensor products of quantum theory!), with the product partition P^N whose cells are given by N -tuples $\langle C_{i_1}, C_{i_2}, \dots, C_{i_N} \rangle$. So, variation consists in not all individuals being in one cell: i.e. the population macro-state is not an N -tuple with all components equal to some single C_i . (Here one could adopt the occupation number formalism from statistical physics.)

Finally, as to (v), several rounds of selection: again, I will be simplistic. Not only will I take selection to be essentially culling as in §4.1, albeit with a designated/desired macro-state C^* that is a function of the environment's macro-state, i.e. $C^* = C^*(C_k^e)$. Also, I will assume that after culling:

- (a) the M ($M \leq N$) surviving individuals simply persist in the desired macro-state C^* that they (are so lucky to!) have been in (or, if you prefer, each is replaced by an offspring in the macro-state C^*); and
- (b) $N - M$ new individuals spring up (as from dragons' teeth!), in some randomly selected combination of macro-states, to proceed to the next culling, along with the M individuals in C^* who have survived from the last round.

Now it is obvious how this toy-model achieves adaptation. As I have assumed that the environment is unchanging, i.e. always in a certain macro-state C_k^e , the culling always favours the same macro-state, $C^* = C^*(C_k^e)$, of individuals. Therefore, over sufficiently many rounds of selection, the random production, at the start of each round, of unfit variations, i.e. macro-states $C_i \neq C^*(C_k^e)$, decays away. That is, over time, the population (and the individuals) achieves adaptation in the sense of notion of conforming as in §4.1: i.e. a constant map taking the desired macro-state C^* as its value.

This scenario can be summed up in terms of functions on macro-states. First, we adapt to the partition P^N the notation we used in §4.1 for culling. That is,

- (a) we adjoin the rubbish macro-state 0 to the partition P^N so as to define a partition $P^{(N,0)} := P^N \cup \{0\}$; and
- (b) we adopt the obvious component-wise definition of the characteristic function $\chi_{C^*} \equiv \chi_{C^*(C_k^e)}: P^N \rightarrow P^{(N,0)}$, defined on the partition P^N and with co-domain $P^{(N,0)}$.

We also need to represent the birth, after each round of culling, of the new individuals. We do this by postulating that, after each round, we apply a function $\beta: P^{(N,0)} \rightarrow P^N$, which is defined (i) to keep constant any component equal to C^* and (ii) to replace any component equal to 0 by some 'living', non-rubbish macro-state $C_i \in P$ (as hinted by β 's co-domain being just P^N). But I shall not formalize the idea that the new individuals' macro-states, given by (ii) of β , are randomly chosen: I shall simply imagine that, for each round of culling, the function β is in general different; so that we postulate a sequence of functions, $\beta_1, \beta_2, \beta_3, \dots$, each subject to the requirements (i) and (ii).

Thus, each round of selection and birth is an application of the characteristic culling function $\chi_{C^*(C_k^e)}$, followed by an application of one of the β functions, representing the birth of new individuals. So the (macro)-histories of all the individuals, and of the population, that are possible, for a given environment macro-state C_k^e and given random functions β_n , are encoded in the sequence of functions

$$P^N \xrightarrow{\chi_{C^*(C_k^e)}} P^{(N,0)} \xrightarrow{\beta_1} P^N \xrightarrow{\chi_{C^*(C_k^e)}} P^{(N,0)} \xrightarrow{\beta_2} P^N \xrightarrow{\chi_{C^*(C_k^e)}} P^{(N,0)} \xrightarrow{\beta_3} \dots \xrightarrow{\beta_p} P^N. \quad (4.1)$$

These functions are defined so that a generic initial population macro-state is mapped eventually to a state where all individuals are in C^* : adaptation! Thus, suppose that, in the first generation, the first, third and N th individuals happen to be in the desired macro-state C^* , while the second, fourth (and no doubt other!) individuals are not and so get culled. And suppose that, after many, say p , generations, the functions β_n have been 'random enough' to have thrown up at some time (i.e. for some n) the desired macro-state C^* for every component. This would give a (macro)-history as follows:

$$\begin{aligned} \langle C_{i1}, C_{i2}, C_{i3}, C_{i4}, \dots, C_{iN} \rangle &\equiv \langle C^*, C_{i2}, C^*, C_{i4}, \dots, C^* \rangle \\ &\xrightarrow{\chi_{C^*(C_k^e)}} \langle C^*, 0, C^*, 0, \dots, C^* \rangle \xrightarrow{\beta_1} \langle C^*, C'_{i2}, C^*, C'_{i4}, \dots, C^* \rangle \\ &\xrightarrow{\chi_{C^*(C_k^e)}} \dots \xrightarrow{\beta_p} \langle C^*, C^*, C^*, \dots, C^* \rangle. \end{aligned} \quad (4.2)$$

To sum up this section, here is a toy-model of adaptive selection. Agreed, it is very simple. Indeed, it does not represent micro-states; so that, in terms of my framework, it cannot represent top-down causation. But I submit that, if we elaborated the model so as to include micro-states, we would be in a situation like that in §4.1. We would face issues about whether the model's macro-dynamics meshes with its micro-dynamics. For example, non-meshing would be threatened by indeterminism, owing to each individual, and so the population as a whole, being an open system. But we could write down a meshing dynamics (perhaps, for realism, availing ourselves of allowances like (a)–(d) in §3.2); and thus get a formal description of Ellis' third type of top-down causation—at least in the sense of causation as functional dependence.

4.3. Top-down causation via adaptive information control

Ellis writes [3, p. 18]: ‘Adaptive information control takes place when there is adaptive selection of goals in a feedback control system, thus combining both feedback control (§4.1) and adaptive selection (§4.2). The goals of the feedback control system are irreducible higher level variables determining the outcome, but are not fixed as in the case of non-adaptive feedback control; they can be adaptively changed in response to experience and information received. The overall process is guided by fitness criteria for selection of goals, and is a form of adaptive selection in which goal selection relates to future rather than the present use of the feedback system. This allows great flexibility of response to different environments, indeed in conjunction with memory it enables learning and anticipation... and underlies effective purposeful action as it enables the organism to adapt its behaviour in response to the environment in the light of past experience, and hence to build up complex levels of behaviour’.

In representing these ideas, one again needs to make choices, and exercise judgement, about which details to be explicit about. I will again be simplistic. I will also build on choices of §4.2, and its ensuing notations. This means that, although I will represent the idea that the goal is not fixed but depends on history, I will *not* capture one prominent feature of control or feedback: namely, the system’s time-evolution being guided towards the goal, by, for example, the system calculating the difference between its present state and the goal-state, and then engineering its change in the next time-step so as to reduce this difference.

Instead, I will postulate, like I did in §4.2, that a system evolves randomly, until it happens to hit its goal—after which it stays in that state. Agreed, that does not merit the names ‘control’, ‘feedback’ or ‘learning’. But it will be clear that, at the cost of more definitions and notation, my framework could equally well describe these notions, namely in terms of time-evolutions using such difference-reducing rules. Thus, I submit that what follows, though simple, is enough for my purpose; namely, to show that one could modify a model of evolution such as that of §4.2 so as to represent Ellis’ ‘adaptive information control’.

Recall that, in §4.2, the goal C^* was a function of just the (time-constant) environment macro-state C_k^e . We had $C^* = C^*(C_k^e)$; and throughout the process of evolution (and adaptation), the same C^* acted as the goal (the attractor macro-state for each individual). I now modify this so as to make each individual’s goal a function of its history, and also the environment’s history (recall Ellis’ mention of memory and past experience).

So let us imagine N persisting individuals (despite the talk on offspring and births in §4.2), labelled by $j = 1, 2, \dots, N$. Time is discrete, with the generic time-point labelled n . So the environment passes through a sequence of macro-states

$$C_{k_1}^e \mapsto C_{k_2}^e \mapsto C_{k_3}^e \dots \mapsto C_{k_n}^e \mapsto \dots \quad (4.3)$$

and individual j passes through a sequence of macro-states

$$C_{i_j}^I \mapsto C_{i_j}^{II} \mapsto C_{i_j}^{III} \dots \mapsto C_{i_j}^{(n)} \mapsto \dots \quad (4.4)$$

We postulate that the goal of each individual j at each time-point n is given by a goal-function $C_{n,j}^*$, which takes as its argument a prior (macro)-history both of j and of the environment (but for simplicity: not other individuals!), and as its value an individual’s macro-state, i.e. an element C_i of the partition P of the individual state-space.

So let us write $H_{n,j}$ for a (macro)-history both of j and of the environment up to the time-point n , i.e. $H_{n,j}$ is a $2n$ -tuple $\langle C_{i_j}^I, C_{i_j}^{II}, \dots, C_{i_j}^{(n)}; C_{k_1}^e, C_{k_2}^e, \dots, C_{k_n}^e \rangle$, and let us write $\mathcal{H}_{n,j}$ for the class of these $2n$ -tuples. Summing up, we have the goal-function

$$C_{n,j}^* : \mathcal{H}_{n,j} \ni H_{n,j} = \langle C_{i_j}^I, C_{i_j}^{II}, \dots, C_{i_j}^{(n)}; C_{k_1}^e, C_{k_2}^e, \dots, C_{k_n}^e \rangle \mapsto C_{n,j}^*(H_{n,j}) \in P. \quad (4.5)$$

We now adjoin these definitions to the scenario of §4.2. There, each time-step (‘round’) involved a culling and rebirth for those individuals who had not attained their goal-state, while those in the goal-state C^* simply persisted in it. Now, using our present metaphor of N persisting individuals with an ever-lengthening history, each round must (i) keep those individuals that have attained their goal-state in that state and (ii) assign to any other individual some random macro-state as its next state.

The combination of (i) and (ii) means that, provided assignments of macro-states of (ii) are sufficiently random that for each individual they sometimes assign its present goal-state, then in the long run, all the individuals attain—and stay in—their goals. In short, again, we have adaptation.

As regards spelling out the formal details of (i) and (ii): I will skip the details about (ii), which are a straightforward adaptation of the β (‘birth’) functions of §4.2. As to (i), there are two features we need to require, the first being more important.

- (a) If up to time-point $n + 1$, the individual j and environment has experienced the joint history $H_{n,j}$, so that j ’s goal is then $C_{n,j}^*(H_{n,j})$, and if j happens to be in the state $C_{n,j}^*(H_{n,j})$, then we require that j will forever remain in $C_{n,j}^*(H_{n,j})$. That is, we require that, once any individual enters its goal, it stays forever.
- (b) It is natural to have goal-functions ‘stay consistent’, i.e. go on endorsing any goal that is attained. That is, it is natural to require that for any j and n , and joint history $H_{n,j}$, with initial segments $H_{m,j}$ ($m \leq n$): if the m th component of (the individual-history first-half of) $H_{n,j}$ is $C_{m,j}^*(H_{m,j})$, then all the later initial segments of $H_{n,j}$ i.e. $H_{p,j}$ with $p > m$, yield endorsements of this goal, i.e. $C_{p,j}^*(H_{p,j}) = C_{m,j}^*(H_{m,j})$. (Owing to (a), the system will in fact stay in $C_{m,j}^*(H_{m,j})$ at all times $p > m$.)

To sum up, these two points encode the idea that a goal-state is an attractor.

I thank the John Templeton Foundation and George Ellis for the invitation; Harald Atmanspacher, Robert Bishop, Adam Caulton and especially Christian List and Elliott Sober for comments; the symposium participants for discussion; and, above all, George Ellis for the inspiration of his work, and for very helpful detailed comments on an earlier version.

REFERENCES

- 1 Butterfield, J. 2011 Emergence, reduction and supervenience: a varied landscape. *Found. Phys.* **41**, 920–960. (doi:10.1007/s10701-011-9549-0)
- 2 Butterfield, J. 2011 Less is different: emergence and reduction reconciled. *Found. Phys.* **41**, 1065–1135. (doi:10.1007/s10701-010-9516-1)
- 3 Ellis, G. 2008 On the nature of causation in complex systems. *Trans. R. Soc. S. Afr.* **63**, 69–84. (doi:10.1080/00359190809519211)
- 4 Ellis, G. F. R. 2012 Top-down causation and emergence: some comments on mechanisms. *Interface Focus* **2**, 126–140. (doi:10.1098/rsfs.2011.0062)
- 5 Nagel, E. 1961 *The structure of science: problems in the logic of scientific explanation*. New York, NY: Harcourt.
- 6 Hempel, C. 1966 *Philosophy of natural science*. Englewood Cliffs, NJ: Prentice-Hall.
- 7 Schaffner, K. In press. Ernest Nagel and reduction. *J. Phil.*
- 8 Sober, E. 1999 The multiple realizability argument against reductionism. *Phil. Sci.* **66**, 542–564. (doi:10.1086/392754)
- 9 Putnam, H. 1975 *Mind, language and reality*, vol. 2, ch. 14, pp. 291–303. Cambridge, UK: Cambridge University Press.
- 10 Fodor, J. 1974 Special sciences (or: the disunity of science as a working hypothesis). *Synthese* **28**, 97–115. (Reprinted by Bedau, M. & Humphreys, P. (eds) 2008 *Emergence: contemporary readings in philosophy and science*. Cambridge, MA: MIT Press/Bradford Books.) (doi:10.1007/BF00485230)
- 11 Shapiro, L. 2000 Multiple realizations. *J. Phil.* **97**, 635–654. (doi:10.2307/2678460)
- 12 Papineau, D. 2010 Can any sciences be special? In *Emergence in mind* (eds C. Macdonald & G. Macdonald), pp. 179–197. Oxford, UK: Oxford University Press. (<http://www.kcl.ac.uk/content/1/c6/04/17/78/SpecialSciencesinMacdonalds.pdf>)
- 13 Dennett, D. 1991 Real patterns. *J. Phil.* **87**. (Reprinted by Bedau, M. & Humphreys, P. (eds) 2008 *Emergence: contemporary readings in philosophy and science*, pp. 27–51. Cambridge, MA: MIT Press/Bradford Books.)
- 14 Bedau, M. 2003 Downward causation and autonomy in weak emergence. *Principia Revista Internacional de Epistemologica* **6**. (Reprinted by Bedau, M. & Humphreys, P. (eds) 2008 *Emergence: contemporary readings in philosophy and science*. Cambridge, MA: MIT Press/Bradford Books.)
- 15 O'Connor, T. 2011 A causally unified world in which fundamental physics is not causally complete. *Interface Focus* **1**.
- 16 Sober, E. 1999 Physicalism from a probabilistic point of view. *Phil. Stud.* **95**, 135–174. (doi:10.1023/A:1004519608950)
- 17 Scerri, E. R. 2012 Top-down causation regarding the chemistry–physics interface: a sceptical view. *Interface Focus* **2**, 20–25. (doi:10.1098/rsfs.2011.0061)
- 18 Hendry, R. 2010 Ontological reduction and molecular structure. *Stud. History Phil. Mod. Phys.* **41**, 183–191. (doi:10.1016/j.shpsb.2010.03.005)
- 19 Primas, H. 1981 *Chemistry, quantum mechanics and reductionism*. Berlin, Germany: Springer.
- 20 Amann, A. 1993 The *Gestalt* problem in quantum theory: generation of molecular shape by the environment. *Synthese* **97**, 125–256. (doi:10.1007/BF01255834)
- 21 Bishop, R. & Atmanspacher, H. 2006 Contextual emergence in the description of properties. *Found. Phys.* **36**, 1753–1777. (doi:10.1007/s10701-006-9082-8)
- 22 Auletta, G., Ellis, G. & Jaeger, L. 2008 Top-down causation by information control: from a philosophical problem to a scientific research program. *J. R. Soc. Interface* **5**, 1159–1172. (doi:10.1098/rsif.2008.0018)
- 23 Shapiro, L. & Sober, E. 2007 Epiphenomenalism—the do’s and the don’t’s. In *Thinking about causes* (eds P. Machamer & G. Wolters), pp. 235–264. Pittsburgh, PA: University of Pittsburgh Press.
- 24 Kim, J. 1999 Making sense of emergence. *Phil. Stud.* **95**. (Reprinted by Bedau, M. & Humphreys, P. (eds) 2008 *Emergence: contemporary readings in philosophy and science*, pp. 3–36. Cambridge, MA: MIT Press/Bradford Books.) (doi:10.1023/A:1004563122154)
- 25 Humphreys, P. 1997 How properties emerge. *Phil. Sci.* **64**, 1–17. (Reprinted by Bedau, M. & Humphreys, P. (eds) 2008 *Emergence: contemporary readings in philosophy and science*. Cambridge, MA: MIT Press/Bradford Books.) (doi:10.1086/392533)
- 26 List, C. & Menzies, P. 2009 Non-reductive physicalism and the limits of the exclusion principle. *J. Phil.* **106**, 475–502.
- 27 Menzies, P. & List, C. 2010 The causal autonomy of the special sciences. In *Emergence in mind* (eds C. Macdonald & G. Macdonald), pp. 139–168. Oxford, UK: Oxford University Press. (doi:10.1093/acprof:oso/9780199583621.003.0008)
- 28 Lewis, D. 1973 Causation. *J. Phil.* **70**, 556–567. (doi:10.2307/2025310)
- 29 Yablo, S. 1992 Mental causation. *Phil. Rev.* **101**, 245–280. (doi:10.2307/2185535)
- 30 Uffink, J. & Valente, G. 2010 Time’s arrow and Lanford’s theorem. *Seminaire Poincare Le temps* **15**, 141–173.
- 31 Ambrosio, L. 2011 The work of Cedric Villani. *Not. Am. Math. Soc.* **58**, 464–468.
- 32 Yau, H.-T. In press. The work of Cedric Villani. In *Proc. 2010 Int. Congress of Mathematicians*.
- 33 Bohm, D. & Hiley, B. 1992 *The undivided Universe*. London, UK: Routledge.
- 34 Holland, P. 1993 *The quantum theory of motion*. Cambridge, UK: Cambridge University Press. (doi:10.1017/CBO9780511622687)
- 35 List, C. 2011 Free will, determinism, and the possibility to do otherwise. See <http://personal.lse.ac.uk/list/PDF-files/FreeWill.pdf>
- 36 Atmanspacher, H. 2012 Identifying mental states from neural states under mental constraints. *Interface Focus* **2**, 74–81. (doi:10.1098/rsfs.2011.0058)
- 37 Atmanspacher, H. & beim Graben, P. 2007 Contextual emergence of mental states from neurodynamics. *Chaos Complex. Lett.* **2** 151–168.
- 38 Love, A. C. 2012 Hierarchy, causation and explanation: ubiquity, locality and pluralism. *Interface Focus* **2**, 115–125. (doi:10.1098/rsfs.2011.0064)
- 39 Noble, D. 2012 A theory of biological relativity: no privileged level of causation. *Interface Focus* **2**, 55–64. (doi:10.1098/rsfs.2011.0067)
- 40 Okasha, S. 2012 Emergence, hierarchy and top-down causation in evolutionary biology. *Interface Focus* **2**, 49–54. (doi:10.1098/rsfs.2011.0046)