The Metaphysics Within Physics

TIM MAUDLIN

OXFORD
UNIVERSITY PRESS

2

Why Be Humean?

Everything only connected by 'and' and 'and'.
(Elizabeth Bishop, *Over 2,000 Illustrations and a Complete Concordance*)

The title of this chapter is not a rhetorical question. Nor is it a question that the essay aspires to answer. It is, rather, a sincere plea for enlightenment. There is, in some contemporary metaphysics, an explicit preference, or desire, or in some cases demand, for 'Humean' theories. Humean, or 'empiricist', theories of law and of chance are sought; theories that posit irreducible nomic or modal or dispositional or causal facts are dismissed as unHumean. David Lewis has characterized a central motivation for some of his theories as a desire to 'uphold not so much the truth of Humean supervenience but the *tenability* of it' (Lewis 1986a, p. xi), a somewhat modest but still mysterious ambition. Why, to put it bluntly, should one *want* to be Humean? What is the appeal of 'Humean Supervenience' such that metaphysical accounts should aspire to it? Although Lewis and others issue calls to rally to Hume's banner, no strategic justification for this campaign is offered. I suspect that the reason for this reticence is that the motivations will not withstand close scrutiny in the light of day. The aim of this essay is to unshutter the windows.

Any examination of Humean Supervenience must begin with a statement of the doctrine. Lewis provides the *locus classicus*:

> Humean supervenience is named in honor of the greater [*sic*] denier of necessary connections. It is the doctrine that all there is to the world is a vast mosaic of local matters of fact, just one little thing and then another. (But it is no part of the thesis that these local matters of fact are mental.) We have geometry: a system of external relations of spatio-temporal distance between points. Maybe points of spacetime itself, maybe point-sized bits of matter or aether fields, maybe both. And at those points we have local qualities: perfectly natural intrinsic properties which need nothing bigger than a point at which to be instantiated. For short: we have an arrangement of qualities. And that is all. All else supervenes on that. (ibid., p x)

and again:

> The question turns on an underlying metaphysical issue. A broadly Humean doctrine (something I would very much like to believe if at all possible) holds that all the facts there are about the world are particular facts, or combinations thereof. This need not be taken as a doctrine of analyzability, since some combinations of particular facts cannot be captured in any finite way. It might better be taken as a doctrine of supervenience: if two worlds match perfectly in all matters of particular fact, they match in all other ways too–in modal properties, laws, causal connections, chances .... (ibid. 111)

Although he does not remark it, Lewis's Humeanism comprises two logically independent doctrines. The first, which we may call *Separability*, claims that all fundamental properties are local properties and that spatio-temporal relations are the only fundamental external physical relations. To be precise:

> Doctrine 1 (Separability): The complete physical state of the world is determined by (supervenes on) the intrinsic physical state of each spacetime point (or each pointlike object) and the spatio-temporal relations between those points.

Separability posits, in essence, that we can chop up space-time into arbitrarily small bits, each of which has its own physical state, much as we can chop up a newspaper photograph into individual pixels, each of which has a particular hue and intensity. As the whole picture is determined by nothing more than the values of the individual pixels plus their spatial disposition relative to one another, so the world as a whole is supposed to be decomposable into small bits laid out in space and time.

The doctrine of Separability concerns only how the total physical state of the universe depends on the physical state of localized bits of the universe. The second component of Lewis's Humeanism takes care of everything else:

> Doctrine 2 (Physical Statism): All facts about a world, including modal and nomological facts, are determined by its total physical state.

I have employed the unlovely neologism 'Physical Statism' to distinguish Doctrine 2 from Physicalism. Physicalism holds that two worlds that agree in all physical respects (i.e. with respect to all items that would be mentioned in a perfected physics) agree in all respects. Physicalism is a much less contentious thesis than Doctrine 2. Doctrine 2 essentially adds to Physicalism the further requirement that all physical facts about the world are determined by its total physical state, by the disposition of physical properties. If one holds,[1] for example, that the laws of nature do not supervene on the total physical state of the world (at least so far as that state can be specified independently of the laws), then one can be a Physicalist while denying Physical Statism. One can hold that worlds which agree on both their physical state and their physical laws agree on all else, while denying that the laws are determined by the state. Lewis's Humeanism importantly maintains the stronger claim.

In order to clearly distinguish Doctrine 2 from Physicalism, we must remark a condition on acceptable analyses accepted by the Physical Statist but not by the Physicalist:

> Non-circularity condition: The intrinsic physical state of the world can be specified without mentioning the laws (or chances, or possibilities) that obtain in the world.

When allied with the doctrine of Separability, the non-circularity condition implies that the physical state of every space-time point is metaphysically independent of the laws that govern the world. This in turn implies that the fundamental physical quantities, such as electric charge, mass, etc., are metaphysically independent of the laws of electromagnetism, gravitation, and so on. This is a controversial thesis, but one that Lewis accepts. It will not come in for further notice here.

The conjunction of Separability with Physical Statism is not peculiar to Lewis. Consider John Earman's account of what it is to be an empiricist about laws of nature:

> The filling I prefer for the blank in (F1) produces the following constraint:
>
> (E1) For any W_1, W_2, if W_1 and W_2 agree on all occurrent facts, then W_1 and W_2 agree on laws.
>
> I will refer to (E1) as the empiricist loyalty test on laws, for I believe it captures the central empiricist intuition that laws are parasitic on occurrent facts. Ask me what an occurrent fact is and I will pass your query on to empiricists. But in lieu of a reply, I will volunteer that the paradigm form of an occurrent fact is: the fact expressed by the sentence $P(o, t)$, where '*P*' is again a suitable kosher predicate, '*o*' denotes a physical object or spatial location, and '*t*' denotes a time. There may also be general occurrent facts (I think there are), but these are also parasitic on the singular occurrent facts. Conservative empiricists may want to restrict the antecedent of (E1) so as to range only over observable facts while more liberal empiricists may be happy with unobservable facts such as the fact that quark *q* is charming and flavorful at *t*. In this way we arrive at many different versions of the loyalty test, one for each persuasion of empiricist. (Earman 1984, p. 195)

[1] Cf. 'A Modest Proposal Concerning Laws, Counterfactuals, and Explanations', Chapter 1, this volume; Carroll 1994.

Note that Earman's 'paradigm form' of an occurrent fact is a fact about a particular (presumably small) space-time region, thus endorsing Separability. Note also that among the particular occurrent facts about a small space-time region there had best not be facts about the laws which govern that region, else an empiricist analysis of laws could consist in stating, independently of all other facts, which laws obtain at every region of space-time.

The interest in dissecting Humean Supervenience into Separability and Physical Statism arises, in the first place, from the remarkable fact that contemporary physics strongly suggests that the world is not Separable. If quantum theory is even remotely on the right track, then the best physical theories will continue, as they do now, to posit fundamental non-Separable physical states of affairs. This discovery casts the question of *motivating* a desire to defend Doctrine 1 into a peculiar light, for one knows beforehand that the motivations, whatever they may be, turn out to lead away from the truth. So before asking why one might want to be Humean, we shall review the evidence that the world is not Humean. Only then will we seek the motivations for defending Separability, and then lastly turn to the possible motivations for Physical Statism.

1. NON-SEPARABILITY IN QUANTUM THEORY

The notion that the physical state of the world is separable is not a philosopher's fancy. In a now famous letter to Max Born, Albert Einstein stated the doctrine succinctly and lucidly:

> If one asks what, irrespective of quantum mechanics, is characteristic of the world of ideas of physics, one is first of all struck by the following: the concepts of physics relate to a real outside world, that is, ideas are established relating to things such as bodies, fields, etc., which claim 'real existence' that is independent of the perceiving subject-ideas which, on the other hand, have been brought into as secure a relationship as possible with the sense-data. It is further characteristic of these physical objects that they are thought of as arranged in a space-time continuum. An essential aspect of this arrangement of things in physics is that they lay claim, at a certain time, to an existence independent of one another, provided these objects 'are situated in different

parts of space'. Unless one makes this kind of assumption about the independence of the existence (the 'being-thus') of objects which are far apart from one another in space—which stems in the first place from everyday thinking—physical thinking in the familiar sense would not be possible. It is also hard to see any way of formulating and testing the laws of physics unless one makes a clear distinction of this kind. This principle has been carried to extremes in the field theory by localizing the elementary objects on which it is based and which exist independently of each other, as well as the elementary laws which have been postulated for it, in the infinitely small (four-dimensional) elements of space.

The following idea characterizes the relative independence of objects far apart in space (A and B): external influence on A has no direct influence on B; this is known as the 'principle of contiguity', which is used consistently in the field theory. If this axiom were to be completely abolished, the idea of the existence of (quasi-) enclosed systems, and thereby the postulation of laws which can be checked empirically in the accepted sense, would become impossible. (Born 1971, pp. 170–1)

It is no accident that Einstein discusses this principle in connection with quantum mechanics, for Einstein saw, perhaps earlier than anyone else, that the formalism of the quantum theory seems to reject Separability. Let's review quickly why that is so.

The quantum theory, no matter how interpreted, employs as a fundamental device the so-called *quantum state* or *wavefunction* of a system. These quantum states obey a *principle of superposition* to the effect that if **A** represents one quantum state of a system and **B** another, then $\alpha\mathbf{A} + \beta\mathbf{B}$ represents a third possible state of the system, where α and β are complex numbers such that $|\alpha|^2 + |\beta|^2 = 1$. For our example, we will use the *spin states* of electrons; the particular mathematical details of how these states are represented will be of no concern.

The spin of an electron can be measured in any direction in space, and when measured one is certain to get one of two possible results which correspond, using the usual measurement apparatus, to the particle being deflected either in one direction or in the opposite direction by a magnetic field. So if we want to measure the spin of an electron in the z-direction, we shoot the electron through a particularly oriented device and the electron will either be deflected up or down. In the first case we say the electron has z-spin up, in the second z-spin down. It turns out that there is a unique spin state for an electron in which it is *guaranteed* to be deflected upwards in an z-spin measurement, and such an electron is said to have z-spin up; similarly for z-spin down. Using an obvious notation, we represent the z-spin up state as $|z\uparrow>$ and the z-spin down state as $|z\downarrow>$. Similarly there are up and down spin states for spin measured in any other direction. Quite significantly, any spin state in any direction can be written down as a superposition of the up and down spin states for any other direction. So, for example, we can write the z-spin states as superpositions of the up and down spin states in the x-direction (i.e. a direction orthogonal to the z-direction) as follows:

$$|z\uparrow> = \frac{1}{\sqrt{2}}|x\uparrow> + \frac{1}{\sqrt{2}}|x\downarrow>$$
$$|z\downarrow> = \frac{1}{\sqrt{2}}|x\uparrow> - \frac{1}{\sqrt{2}}|x\downarrow>.$$

z-spin up is a superposition of x-spin up and x-spin down, as is z-spin down.

We can now introduce the only rule of quantum mechanical calculation we will need. Suppose an electron is in some arbitrary spin state **A**, and we decide to measure its x-spin. The quantum formalism tells us to calculate the probabilities of up and down results in the following way. First, write down the state **A** in terms of the states $|x\uparrow>$ and $|x\downarrow>$, that is, write **A** as a superposition of x-spin states $\alpha|x\uparrow> + \beta|x\downarrow>$. Then the probability of getting an up result when measuring x-spin is just $|\alpha|^2$ and the probability of a down result is $|\beta|^2$. To take a concrete example suppose one measures the x-spin of an electron in the z-spin up state $|z\uparrow>$. Since $|z\uparrow> = 1/\sqrt{2}|x\uparrow> + 1/\sqrt{2}|x\downarrow>$, the chance of getting x-spin up is 1/2, as is the chance for x-spin down. This simple rule is all of the quantum theory we will require.

(Incidentally, these rules illustrate why the so-called Heisenberg Uncertainty Principle is built into the fundamental structure of the quantum theory rather than being imposed as an *additional* constraint on the formalism. There is, for example, no spin state in which one could predict with certainty the result of both a z-spin and an x-spin measurement. For the only states which guarantee a particular z-spin result are $|z\uparrow>$ and $|z\downarrow>$, and each of these is, as a purely mathematical matter of fact, a superposition of x-spin states.)

None of these facts about quantum spin as yet implies anything about Separability, for we have only been discussing the spin state of a single particle. But once we apply these same principles to systems containing more than one particle, we get some startling results.

Let us now consider a system that consists in two electrons. What sorts of spin states are available to it?

There are, to begin, the sorts of uninteresting states that the principle of Separability would lead one to expect. It is possible, for example, that particle 1 and particle 2 should both be in the state $|z\uparrow>$, or particle 1 can be in $|z\uparrow>$ while particle two is in $|z\downarrow>$, or particle 1 can be in $|z\uparrow>$ while particle 2 is in $|x\uparrow>$. Using an obvious notation, these states for the two-particle system are written $|z\uparrow>_1|z\uparrow>_2$, $|z\uparrow>_1|z\downarrow>_2$, and $|z\uparrow>_1|x\uparrow>_2$ respectively. In general, for any spin state that the first particle alone can be in and any state that the second alone can be in, there is a state of the joint system which assigns exactly the first state to the first particle and the second state to the second. These states are called *product states* of the composite system and are, metaphysically at least, as boring as could be. A product state assigns a perfectly determinate spin state to each of the two particles, and the state of the composite is nothing but the logical sum of the states of the components.

If all quantum states of composite systems were mere product states, then quantum theory would pose no threat to Separability. But from the principle of superposition it follows that there are more possibilities available for the composite system than product states. Since any pair of states of a system can be superposed to yield a new state, any pair of product states can be superposed. For our purposes, we will consider only two such states: the Singlet State and the m = 0 Triplet State:

$$\text{Singlet State}: \quad \frac{1}{\sqrt{2}}|z\uparrow>_1|z\downarrow>_2 - \frac{1}{\sqrt{2}}|z\downarrow>_1|z\uparrow>_2$$

$$m = 0 \text{ Triplet State}: \quad \frac{1}{\sqrt{2}}|z\uparrow>_1|z\downarrow>_2 + \frac{1}{\sqrt{2}}|z\downarrow>_1|z\uparrow>_2.$$

Let's play with these states a bit to see what sorts of statistical properties they display.

Suppose we decide to measure the z-spins of both particles. If the system were in the state $|z\uparrow>_1|z\downarrow>_2$ we would get an up result for particle 1 and down for particle 2. If it were in $|z\downarrow>_1|z\uparrow>_2$ we would get down for 1 and up for 2. So in the Singlet State there is a 50 per cent chance of up-down and a 50 per cent chance of down-up, and similarly for the m = 0 Triplet State.

Suppose we decide simply to measure the z-spin of particle 1, and either measure spin in some other direction or nothing at all on particle 2. There is a fundamental quantum mechanical principle that applies here: the statistical predictions for the results of measurement on one part of a composite system are unchanged by conditionalizing on the fact that any sort of measurement was made on another part. This principle is sometimes called parameter independence in the literature, and it holds in all interpretations of the

quantum theory when applied to the statistical predictions one can derive from the wavefunction alone.[2] So we know that the statistics displayed by particle 1 under a z-spin measurement are unchanged whether z-spin or anything else is measured on particle 2. And we know that if z-spin is measured on particle 2, particle 1 will come out z-spin up 50 per cent of the time. Therefore the prediction for particle 1 (or particle 2) when z-spin is measured is simply 50 per cent chance of up and 50 per cent down. In fact, the statistical predictions for *any* direction of spin for either particle are 50 per cent up and 50 per cent down: in any direction one wishes to choose, each particle in the Singlet State or the m = 0 Triplet State has an even chance of coming out up or down.

This last fact is of considerable interest, for *no pure quantum state for a single particle displays these statistics.* For example, in the state $|z\uparrow>_1$, particle 1 is certain to yield an up result if z-spin is measured. And in every pure single particle quantum spin state, there is some direction such that the result of a spin measurement in that direction is certain. It follows that the Singlet and m = 0 Triplet States *cannot be written as simple product states but only as superpositions of product states.* Such states are called *entangled states*, and they engender the most shocking and radical metaphysical innovations to be found in the quantum theory.

The central challenge that quantum theory poses for Separability can now be stated. Suppose there are two electrons, well separated in space (perhaps at opposite ends of a laboratory), that are in the Singlet State. If the principle of Separability held, then each electron, occupying a region disjoint from the other, would have its own intrinsic spin state, and the spin state of the composite system would be determined by the states of the particles taken individually, together with the spatio-temporal relations between them. But, as we have seen, no pure state for a single particle yields the same predictions as the Singlet State, and if one were to ascribe a pure state to each of the electrons, their joint state would be a product state rather than an entangled state. The joint state of the pair simply cannot be analyzed into pure states for each of the components.

[2] The reason for the careful qualifications in this sentence is that not all interpretations of the quantum theory display parameter independence at the level of all fundamental ontological posits. Bohm's theory, for example, violates it for individual systems, but still obeys it at the level of statistical predictions. Any theory that disobeyed parameter independence at the level of the statistical predictions from the wavefunction would also violate the so-called quantum no-signaling theorems, and would permit superluminal signals.

The attentive reader will have noticed that the qualifier 'pure' has snuck into the preceding discussion. No *pure* state for an individual particle yields 50–50 chances for spin measurements in all directions; if each particle is in a *pure* state then the pair is in a product state, and so on. And indeed, the qualifier 'pure' is needed because there is another set of states, the so-called *impure states* (or *mixtures* or *statistical operators* or *density matrices*), for which not all of these assertions hold. Roughly speaking, one can think of an impure state as the sort of state one would use to make predictions if one were unsure which pure state a system is in. If, for example, one knew that a single electron were either in $|z\uparrow>$ or in $|z\downarrow>$ but were unsure which, and assigned a 50 per cent chance to each possibility, then one could calculate in the usual way expectations for various measurements. And one would find that in such a state there is a 50 per cent chance for an up or a down outcome for spin measurements in any direction.[3] There is a determinate mathematical procedure for deriving the mixed state of each component of a composite system from the state of the whole, when that state is entangled. The resulting mixed state will make all the right statistical predictions about the component, for all possible measurements performed solely on it.

So why not just say that when a pair of electrons is in the Singlet State each electron is in the appropriate mixed state, and thereby recover Separability? The problem for this approach arises when we consider the $m = 0$ Triplet state. If one makes a z-spin measurement on an electron in the $m = 0$ Triplet State, there is a 50 per cent chance of up and 50 per cent of down, and similarly for measurements in any direction. *The mixed states assigned to component electrons in the Singlet State are identical to the mixed states assigned to component electrons in the m = 0 Triplet State.* So if Separability holds, then since each component of a pair of particles in the Singlet State is in exactly the same spin state as each component of a pair in the $m = 0$ Triplet State, and since the spatio-temporal relations between the members of the pair could be identical, the Singlet State would have to be identical to the $m = 0$ Triplet State. Separability holds that the global physical state of a system supervenes

[3] The 'roughly' in the foregoing characterization of mixed states is required because the statistical operators or density matrices are less discriminating than states of epistemic uncertainty. If I know that a particle is definitely either in $|z\uparrow>$ or in $|z\downarrow>$, but am completely unsure which, then I am in a different epistemic state from knowing that the particle is either in $|x\uparrow>$ or in $|x\downarrow>$ and being completely unsure which. But both of these epistemic states would yield exactly the same statistical predictions for all observables, and both correspond to the same statistical operator.

on the local states of the parts plus their spatio-temporal relations, and in this case the states of the parts and the spatio-temporal relations (which play no role in the spin state in any case) are identical.

At this point the reader might well wonder why the Singlet State and the $m = 0$ Triplet State aren't identical. After all, it follows from what has been said that no measurements on a component electron of a pair in the Singlet State can distinguish it from a component of a pair in the Triplet. Taken individually, parts of Singlets act just like parts of $m = 0$ Triplets. Furthermore, the only difference between the Singlet and the $m = 0$ Triplet is a minus sign, which seemingly becomes irrelevant when calculating probabilities, since the coefficients in the superposition are squared. But matters are not so simple. Although no *local* measurement on a *single* electron can distinguish the Singlet from the $m = 0$ Triplet, a *global* measurement on the whole composite system can.

We have already noted that if z-spin measurements are made on both electrons, then the predictions for the Singlet and $m = 0$ Triplet States are the same: 50 per cent up-down and 50 per cent down-up. It is here that the minus sign disappears in the squaring. But what if we measure the *x*-spins of both particles?

To get the answer, we must express the Singlet and $m = 0$ Triplet in terms of x-spin. This is possible since the z-spin states can themselves be written as superpositions of x-spin, as already noted. Starting with our original definition of the Singlet and $m = 0$ Triplet, and substituting

$$|z\uparrow> = \frac{1}{\sqrt{2}}|x\uparrow> + \frac{1}{\sqrt{2}}|x\downarrow> \text{ and}$$
$$|z\downarrow> = \frac{1}{\sqrt{2}}|x\uparrow> - \frac{1}{\sqrt{2}}|x\downarrow>$$

we can derive:

$$\begin{aligned}\text{Singlet} = {} & \frac{1}{\sqrt{2}}\left(\frac{1}{\sqrt{2}}|x\uparrow>_1 + \frac{1}{\sqrt{2}}|x\downarrow>_1\right)\left(\frac{1}{\sqrt{2}}|x\uparrow>_2 - \frac{1}{\sqrt{2}}|x\downarrow>_2\right) \\ & -\frac{1}{\sqrt{2}}\left(\frac{1}{\sqrt{2}}|x\uparrow>_1 - \frac{1}{\sqrt{2}}|x\downarrow>_1\right)\left(\frac{1}{\sqrt{2}}|x\uparrow>_2 + \frac{1}{\sqrt{2}}|x\downarrow>_2\right) \\ = {} & \frac{1}{2\sqrt{2}}(|x\uparrow>_1|x\uparrow>_2 - |x\uparrow>_1|x\downarrow>_2 + |x\downarrow>_1|x\uparrow>_2\end{aligned}$$

$$-|x\downarrow>_1|x\downarrow>_2) - \frac{1}{2\sqrt{2}}(|x\uparrow>_1|x\uparrow>_2 + |x\uparrow>_1|x\downarrow>_2$$
$$-|x\downarrow>_1|x\uparrow>_2 - |x\downarrow>_1|x\downarrow>_2)$$
$$= \frac{1}{2\sqrt{2}}(2|x\downarrow>_1|x\uparrow>_2 - 2|x\uparrow>_1|x\downarrow>_2)$$
$$= \frac{1}{\sqrt{2}}(|x\downarrow>_1|x\uparrow>_2 - |x\uparrow>_1|x\downarrow>_2)$$

$$\text{m} = 0 \text{ Triplet} = \frac{1}{\sqrt{2}}\left(\frac{1}{\sqrt{2}}|x\uparrow>_1 + \frac{1}{\sqrt{2}}|x\downarrow>_1\right)\left(\frac{1}{\sqrt{2}}|x\uparrow>_2 - \frac{1}{\sqrt{2}}|x\downarrow>_2\right) + \frac{1}{\sqrt{2}}\left(\frac{1}{\sqrt{2}}|x\uparrow>_1 - \frac{1}{\sqrt{2}}|x\downarrow>_1\right)\left(\frac{1}{\sqrt{2}}|x\uparrow>_2 + \frac{1}{\sqrt{2}}|x\downarrow>_2\right)$$
$$= \frac{1}{2\sqrt{2}}(|x\uparrow>_1|x\uparrow>_2 - |x\uparrow>_1|x\downarrow>_2 + |x\downarrow>_1|x\uparrow>_2$$
$$-|x\downarrow>_1|x\downarrow>_2) + \frac{1}{2\sqrt{2}}(|x\uparrow>_1|x\uparrow>_2 + |x\uparrow>_1|x\downarrow>_2$$
$$-|x\downarrow>_1|x\uparrow>_2 - |x\downarrow>_1|x\downarrow>_2)$$
$$= \frac{1}{2\sqrt{2}}(2|x\uparrow>_1|x\uparrow>_2 - 2|x\downarrow>_1|x\downarrow>_2)$$
$$= \frac{1}{\sqrt{2}}(|x\uparrow>_1|x\uparrow>_2 - |x\downarrow>_1|x\downarrow>_2)$$

The simple minus sign distinguishing the two states when written in terms of z-spin now looms large: it implies different cancellations when converting to x-spin. And that in turn implies quite different statistics for pairs of x-spin measurements. If the particles are in the Singlet State and x-spin is measured on each, then they are certain to give opposite results, with a 50 per cent chance of up-down and 50 per cent chance of down-up. But the same measurement on the m = 0 Triplet State is certain to yield the *same* result for both electrons: half of the time both will be up and half of the time both down. A single global measurement of x-spin is guaranteed to distinguish a pair of particles in the Singlet State from a pair in the m = 0 Triplet, so the two states cannot be identified.

The difficulty facing Separability is now inescapable. Consider two pairs of electrons, one in the Singlet and the other in the m = 0 Triplet State, such that the spatio-temporal relations within each pair are identical (e.g. in each pair the electrons are 5 meters apart). Can one analyze the global physical state of each pair into local physical states of each part taken individually plus the spatio-temporal relations? Evidently not. For what is the local physical state of each electron?

If only pure states are allowed as possible physical states, then none of the electrons has any local state, i.e. a state that can be specified without reference to the other member of the pair. One could say that the state of one electron is, for example, being part of a pair of electrons which are in the Singlet State, but that is evidently not a purely local matter. And if none of the electrons has a local intrinsic spin state, the global state of each pair cannot supervene on the local states of the parts plus space-time relations.

If mixed states are allowed as possible physical states of systems, then the problem is still not solved. Each electron can now be assigned its own intrinsic local state, but all four electrons are assigned exactly the same state. So if the global spin state of the system supervenes on the local intrinsic states of the parts plus space-time relations, the two pairs must be in identical spin states, which they are not. Either way you cut it, Separability fails. The upshot is that no physical theory that takes the wavefunction seriously can be a Separable theory. If we have reason to believe that the quantum theory, or any extension of it, is part of a true description of the world, then we have reason to believe the world is not Separable.[4]

2. LEWIS'S REACTION AND THE MOTIVATION FOR SEPARABILITY

Lewis is aware that the quantum theory poses a threat to Separability, and says he is prepared to take the consequences:

> Is this [namely Humean Supervenience] materialism?—no and yes. I take it that materialism is metaphysics built to endorse the truth and descriptive completeness of

[4] There is one final move that can be made to save Separability, namely demand that one's account of space itself be altered so that the physical state of the world be Separable in the space. Barry Loewer suggests this in 'Humean Supervenience' (1996), recommending that the 'fundamental space' of quantum mechanics be taken to be configuration space, rather than space-time. This constitutes the ultimate elevation of Separability as a regulative principle, rather than an empirical theory, and urges even more strongly the question of motivation.

physics more or less as we know it; and it just might be that Humean supervenience is true, but our best physics is dead wrong in its inventories of the qualities. Maybe, but I doubt it. Most likely, if Humean supervenience is true at all, it is true in more or less the way that present physics would suggest ...

Really, what I uphold is not so much the truth of Humean supervenience but the *tenability* of it. If physics itself were to teach me that it is false, I wouldn't grieve.

That might happen: maybe the lesson of Bell's theorem is exactly that there are *physical* entities which are unlocalized, and which therefore might make a difference between worlds ... that match perfectly in their arrangements of local qualities. Maybe so. I'm ready to believe it. But I am not ready to take lessons in ontology from quantum physics as it now is. First I must see how it looks when it is purified of instrumentalist frivolity, and dares to say something not just about pointer readings but about the constitution of the world; and when it is purified of doublethinking deviant logic; and—most of all—when it is purified of supernatural tales about the power of observant minds to make things jump. If, after all that, it still teaches nonlocality, I shall submit willingly to the best of authority. (Lewis 1986a, p. xi)

If we take Lewis at his word, then we should abandon Separability (and hence his version of Humean Supervenience) forthwith. For one *can* see how quantum physics looks when purified of instrumentalism, and quantum logic, and consciousness-induced wave collapse. This has been done in several quite different ways: in David Bohm's so-called ontological interpretation (see e.g. Bohm and Hiley 1993), in the (mind-independent) spontaneous collapse theories of Ghirardi, Rimini, and Weber (1986) and of Philip Perle (1990), even in the Many Minds theory of David Albert and Barry Loewer (see Albert 1992). These theories all have fundamentally different ontologies and dynamics, but all agree that the physical state of the world is not Separable, for they all take the wavefunction seriously as a representation of the physical state. This is not to say that Non-Separability is absolutely forced on us by empirical considerations: it would not be impossible to construct a Separable physics with the same empirical import as the present quantum theory.[5] But no one is trying to do it, and there seems to be no reason to start: the quantum theory (in a coherent formulation) is elegant, simple, and

[5] In this regard, one must carefully distinguish Separability from relativistic Locality, i.e. from the claim that the physical state at any point of space-time is determined or influenced only by events in its past light-cone, or, colloquially, that no influence travels faster than light. Bell's theorem shows that certain empirical predictions of the quantum theory, namely violations of Bell's inequality for events at spacelike separation, cannot be recovered by any local theory. They could, however, be recovered by a Separable theory which contains superluminal or backward causal connections. See my 1994 for all the grisly details.

empirically impeccable. Lewis would not elevate his preference for Separable theories into some a priori constraint which could dictate to physics, as the quote shows. Given the definition of materialism cited above, contemporary materialism (i.e. metaphysics built to endorse the approximate truth and descriptive completeness of contemporary physics) must deny Separability.

This leaves us with two questions. First, what drew Lewis to Separability in the first place? Since the doctrine appears to be false, we ought to consider carefully the grounds upon which it was thought to be established, or at least rendered plausible. Second, and more importantly, what of Physical Statism? This second component of Humean Supervenience remains as yet untouched by any criticism, and one could continue to insist upon it even while abandoning Separability. Perhaps the physical state of the universe does not supervene on the local intrinsic states of its point-like parts together with spatio-temporal relations, but yet the 'modal properties, laws, causal connections, chances' (ibid. 111) all are determined by the non-Separable total physical state of the universe. Perhaps. But our suspicions have been rightly aroused. The considerations in favor of Humean Supervenience already led us astray with respect to Separability, so why think they are likely to be any more reliable with respect to Physical Statism? Before we can even begin to take up this question, we must answer the first: what considerations seemed to support Separability in the first place?

Fortunately, the answer to this question is clear, simple, and intelligible. It has, indeed, already been stated. Lewis wants a metaphysics built to endorse the ontology of physics. And, as the quotation from Einstein above forcefully illustrates, *classical* physics *is* Separable. Classical mechanics and field theory do postulate that the physical state of the whole universe is determined entirely by the spatio-temporal dispositions of bodies, their intrinsic physical properties (such as charge and mass), and the values of fields at all points in space through time. Taking one's ontology from classical physics does entail Separability. But the advent of the quantum theory, as we have seen, has superseded that argument; it is irreparably damaged, and Lewis has nothing more to say.

Perhaps, though, Einstein does. At the end of his discussion, Einstein suggests that Separability is a kind of a priori constraint on any comprehensible and empirically verifiable physics (NB: not an a priori constraint on how the world must be, but how it must be if we are to know it through empirical procedures, a truly Kantian theme). Einstein writes that if Separability 'were to be completely abolished, the idea of the existence of (quasi-) enclosed systems, and thereby the postulation of laws which can be checked empirically in the accepted sense, would become impossible' (Born 1971). It is hard to respond directly to this claim, since no further explanation or justification is offered, but if 'completely abolished' means just 'denied', then the quantum theory itself stands as a refutation of the claim. Quantum theory has both been formulated and rigorously tested despite the centrality of non-Separable elements in its ontology. Whatever Einstein had in mind, he had to be wrong.

(If by 'completely denied' Einstein merely means that any empirically testable theory must postulate *some* local intrinsic physical states, but not that the total physical state of all systems is Separable, then he would be anticipating John Bell's call that one carefully consider what the 'local beables' of a theory are, i.e. the objectively existing quantities which '(unlike the total energy, for example) can be assigned to some bounded space-time region' (Bell 1987, p. 53). One could try to make an argument that a physical theory with *no* local beables cannot be brought into correspondence with our experience of the world, but even this weaker claim may face serious obstacles.)

So no credible motivation for Separability exists in the face of the existence and empirical testability of quantum physics. What of Physical Statism?

3. THE MOTIVATION AND STATUS OF PHYSICAL STATISM

Why believe that the 'modal properties, laws, causal connections, chances' all supervene on the total physical state of the universe, that there could not be two possible worlds which agree on their total physical state but disagree on some of these? If the motivation for Separability is to be found in the ontology of classical physics, and if Lewis's materialism is just metaphysics tailored to endorse the (approximate) truth and descriptive completeness of physics as we know it, then one would first seek the motivation for Physical Statism also in the practices of the physicists (and other natural scientists).

What has physics to say about modal properties, laws, causal connections, and chances? The topic is a large one, so I will treat much of it summarily. Physicists do make assertions about what is physically possible and impossible. Cosmologists, for example, regard both open and closed universes as physically possible, and study the features of each. How they manage this is relatively clear: they have the field equations of General Relativity, and regard the physically possible universes as models of these laws (together, perhaps, with some other conditions, e.g. the absence of closed timelike curves, but this is controversial). So once we have an account of physical law, the account of physical possibility is near to hand. Similarly, physicists are happy to evaluate counterfactuals, so long as they are precisely enough stated. This means that the counterfactual condition must be specified completely enough to delimit a class of models in which the condition is satisfied, and the class is coherent enough that in all (or in most, by some natural measure) the consequent of the counterfactual has the same truth value. So, for example, one asks: what would the Earth be like if there were no moon? *What If the Moon Didn't Exist* by Neil Comins (1995) is devoted to just this question. But the question is not yet precise enough: in what way would history be different so that there is no moon? Are we to imagine the moon simply popping out of existence now, or the particles which formed the moon failing to coalesce and forming a ring instead, or . . . ? The moon formed from debris spewed out from the Earth after a collision with a large planetesimal. So Comins frames one of his counterfactuals this way:

> The planetesimal that created the moon traveled trillions of miles over millions of years before hitting Earth. It also swept by other planetesimals as well as by the planets Mars and Venus. Its orbit was altered by the gravitational force from each body it encountered. As a result of all these variations in its path, the planetesimal finally ended up striking the earth. But it need not have met this fate.
>
> If that planetesimal had formed in an orbit different from its actual path by only a few inches, it would not have struck the earth. Over the planetesimal's lifetime the difference between the true orbit and any other path would have been amplified by the gravitational attractions it experienced passing near other bodies. This amplification effect, discovered in the 1980s, stems from a branch of mathematics called chaos. Had it begun in a slightly altered orbit, the planetesimal would easily have been twenty-five thousand miles to one side of its true path by the time it reached the earth in that last, fateful orbit. That change, absolutely minuscule in astronomical terms, would have prevented the collision.
>
> Even such a near miss between earth and another body is no minor event. As it passed by, the planetesimal would be whipped into a dramatically different orbit by the gravity of the nearby earth. Depending on its new course, the planetesimal might eventually strike the sun, Jupiter, or another body, or leave the solar system forever. (ibid. 6)

The treatment of the counterfactual conditional in this passage is transparent. We are to consider a solar system just like ours (with respect to the positions and velocities of matter) save for the placement of a single planetesimal at the time of its formation. The laws of gravity are now used to determine how things would have evolved, in particular they are used to show how the difference in position would be amplified through time, enough to avoid collision with the Earth. The counterfactual has not been precisely enough specified to determine a unique further trajectory for the planetesimal, as different fates befall it in different models that meet the stated conditions.

There is more to be said about this treatment of counterfactuals (and I have tried to say some of it in my 'A Modest Proposal on Laws, Counterfactuals, and Explanations' (Chapter 1, this volume)), but it is relatively clear that once one has the laws (such as that of gravitation) in hand, the treatment of physical possibility and of counterfactuals is relatively straightforward.

If one provides a counterfactual analysis of causation (à la Lewis 1986a, chapter 21) then causal claims supervene on the counterfactuals, which are in turn underwritten by the laws of nature.

Chance also appears in physics, in two guises. There are chances that derive from stochastic dynamical equations, as in quantum theory with wave collapse. These chances are to be found written into the fundamental dynamics themselves. Then there are chances that are sometimes associated with deterministic systems, such as Buffon's calculation of the chance that a thrown needle will fall so as to intersect one of a set of parallel lines. These chances derive from a presumably natural probability measure over the possible initial conditions for the set-up. What makes such a measure natural is a somewhat vexed question, but physicists certainly take the matter to be determined (so far as it is) by other physical facts: there is never any thought of two possible worlds which agree in all their laws and in the total physical state, but disagree nonetheless on which measure over initial conditions is natural. In some cases the considerations are straightforward. Buffon, for example, uses a measure which is isotropic and homogeneous because space itself, and the various factors which influence the needle, are posited as isotropic and homogeneous. For example, if the needle were magnetized, and the lines ruled in a north–south direction, the calculation would be incorrect. In other cases, such as statistical thermodynamics, justification of a natural measure is much more subtle and difficult. But what is clear in these deterministic cases is that the chances are only as objective as the naturalness of the measure, and that in turn must be defended on other physical grounds.

So given the total physical state of the world and the laws of nature, it looks promising, and demonstrably in accord with actual scientific practice, to regard physical possibility, counterfactuals, causal connections, and chances to be fixed, insofar as they are objective at all. But what of the laws themselves?

Explicating these other notions in terms of the laws and physical state is not sufficient for Physical Statism: the laws in turn must be shown to supervene on the total physical state of the world. Is there anything in the practice of physics, classical or contemporary, which suggests that the laws themselves are determined by the total physical state?

In short, the answer is no. It matters not whether one starts with Newton, who, in the *Principia*, simply announces his three laws of motion after giving the definitions of various terms, or whether one turns directly to any contemporary textbook on quantum theory, which will treat, e.g., the Schrödinger equation as a fundamental dynamical principle. Physicists seek laws, announce laws, and use laws, but they do not even attempt to analyze them in terms of the total physical state of the universe or anything else. (One may, of course, attempt to explicate one law as a consequence or approximate consequence of another, as when showing Kepler's laws to be approximate consequences of Newton's laws of motion and gravitation in a particular situation, but that is not an attempt to analyze lawhood *per se*.) Unlike reductive analyses of possibility, causality, and chance, reductive analyses of law are not endorsed by scientific practice.

Indeed, scientific practice seems to preclude such an analysis. As we have seen, physical possibility is easily understood in terms of the models of the laws of physics. Let us suppose (and how can one deny it) that every model of a set of laws is a possible way for *a world governed by those laws* to be. Then we can ask: can two different sets of laws have models with the same physical state? Indeed they can. Minkowski space-time, the space-time of Special Relativity, is a model of the field equations of General Relativity (in particular, it is a vacuum solution). So an empty Minkowski space-time is one way the world could be if it is governed by the laws of General Relativity. But is Minkowski space-time a model *only* of the General Relativistic laws? Of course not! One could, for example, postulate that Special Relativity is the complete and accurate account of space-time structure, and produce another theory of gravitation, which would still have the vacuum Minkowski space-time as a model. So under the assumption that no possible world can be governed both by the laws of General Relativity and by a rival theory of gravity, the total physical state of the world cannot always determine the laws. The only way out is either to assert that empty Minkowski space-time must be governed by *both* sets of laws, since it is a model of both, or (a more likely move) that it can be governed by *neither* set of laws, since neither is the simplest account of space-time structure adequate to the model (the simplest account is just Special Relativity). But how can one maintain that the General Relativistic laws cannot obtain in a world that is a model of those laws, and hence allowed by them? The necessity of distinguishing the physical possibilities (i.e. the ways the world could be given that a set of laws obtains in that world) from the models of the laws signals a momentous shift from philosophical analyses that follow scientific practice to analyses that dictate it.

The situation is even worse for probabilistic laws. Consider a law that assigns a probability to any given event, say the decay of a radioactive atom. The models of such a law will include worlds where every decay event assigned a non-zero probability occurs. Hence that set of models will be *identical*, with respect to the non-probabilistic facts, to the models of a law that assigns a different probability to the event. A law that assigns radium a half-life of thirty years does not rule out every atom of radium decaying in fifteen years, it simply apportions such an eventuality a very low probability. Again, since different laws share the same models, either the laws cannot supervene on the matters of particular fact or else some models of the laws cannot be regarded as physical possibilities relative to those laws. But this last option leads to worse problems. Given a particular initial state, a probabilistic law allows many possible eventualities, and assigns each a probability (let us assume finite models). The sum of the probabilities of the various models is unity: consider, for example, a sequence of a thousand flips of a fair coin. If not all of these models are physical possibilities relative to the law, i.e. worlds where the law can hold, then a law will assign a non-zero probability to its own failure, and the sum of the probabilities of the evaluations consistent with the law will not be unity. These are indigestible consequences.

The supervenience of law on physical state, then, is not only not assumed in scientific practice, it runs contrary to that practice. We should need powerful reasons to pursue such a philosophical analysis of laws. Since Lewis does not provide one, we must seek elsewhere.

4. WHY WAS HUME HUMEAN?

The natural place to begin a search for motivations for Humeanism is Hume. Hume's reasons are quite clear, and also completely outdated, so they need little detain us.

Hume believed that every simple idea had to have been copied from a simple impression, either of perception or reflection. This raised a problem

for the notion of cause and effect, since that concept included the idea of necessary connection, and the necessity of any connection between empirical events is not itself accessible to observation. Insofar as causation is analyzed by nomic subsumption, this raises a parallel problem for laws. Hume's solution was twofold: first to trace the original of the idea of necessary connection to an impression of reflection that accompanies the transition of the mind from one idea to another arising from habit, second to reduce the objective (mind-independent) conditions for causal connection to patterns of succession among event types. Hume did not see how the very idea of a cause, or a law, could ever arise in the mind if it were not somehow reducible to perceptible events. There is little need to delve more deeply into Hume's motivations, since the empiricist theory of ideas is no longer defended anywhere.

The empiricist account of concepts did not go easily. It arose again in the Logical Empiricism of the early half of the twentieth century, albeit transposed from an analysis of the *content of ideas* to an analysis of the *truth conditions of sentences*. In short, Carnap is just Hume warmed over and updated. Compare the following passages, the first from Carnap's 'The Elimination of Metaphysics through Logical Analysis of Language':

> (Meaningful) statements are divided into the following kinds. First there are statements which are true solely by virtue of their logical form ... They say nothing about reality. The formulae of logic and mathematics are of this kind. They are not themselves factual statements, but serve for the transformation of such statements. Secondly there are the negations of such statements ('*contradictions*'). They are self-contradictory, and hence false by virtue of their form. With respect to all other statements the decision about truth or falsehood lies in the protocol sentences. They are therefore (true or false) *empirical statements* and belong to the domain of empirical science. Any statement one desires to construct which does not fall within one of these categories becomes automatically meaningless. Since metaphysics does not want to assert analytic propositions, nor to fall within the domain of empirical science, it is compelled to employ words for which no criteria of application are specified and which are therefore devoid of sense, or else to combine meaningful words in such a way that neither an analytic (or contradictory) statement nor an empirical statement is produced. In either case pseudo-statements are the inevitable product.
>
> Logical analysis, then pronounces the verdict of meaninglessness on any alleged knowledge that pretends to reach above or behind experience. This verdict hits, in the first place, any speculative metaphysics, any alleged knowledge by *pure thinking* or by *pure intuition* that pretends to be able to do without experience. But the verdict equally applies to the kind of metaphysics which, starting from experience, wants to acquire knowledge about that which *transcends experience* by means of

special *inferences* (e.g. the neo-vitalist thesis of the directive presence of an 'entelechy' in organic processes, which supposedly cannot be understood in terms of physics; the question of the 'essence of causality', transcending the ascertainment of certain regularities of succession; the talk of the 'thing in itself') ...

Finally, the verdict of meaninglessness also hits those metaphysical movements which are usually called, improperly, epistemological movements, that is *realism* (insofar as it claims to say more than the empirical fact that the sequence of events exhibits a certain regularity, which makes the application of the inductive method possible) and its opponents: subjective *idealism*, solipsism, phenomenalism, and *positivism* (in the earlier sense). (Carnap 1959, pp. 76–7)

The second, more familiar, from Hume's *Inquiry*:

All the objects of human reason and inquiry may naturally be divided into two kinds, to wit, 'Relations of Ideas' and 'Matters of Fact'. Of the first kind are the sciences of Geometry, Algebra, and Arithmetic, and, in short, every affirmation which is either intuitively or demonstratively certain ...

Matters of fact, which are the second objects of human reason, are not ascertained in the same manner, nor is our evidence of their truth, however great, of a like nature with the foregoing. The contrary of every matter of fact is still possible, because it can never imply a contradiction and is conceived by the mind with the same facility and distinctness as if ever so conformable to reality ...

When we run over our libraries, persuaded of these principles, what havoc must we make? If we take in hand any volume—of divinity or school metaphysics for instance—let us ask, *Does it contain any abstract reasoning concerning quantity or number?* No. *Does it contain any experimental reasoning concerning matter of fact and existence?* No. Commit it then to the flames, for it can contain nothing but sophistry and illusion.

Aside from the fact that Hume is more succinct and elegant a stylist, the doctrines are nearly identical. Just like Hume, the positivists based their justification for the supervenience of law on patterns of observable events on a *semantic* thesis: any non-analytic claims that go beyond what can be empirically justified are meaningless. If one accepts this constraint, then the notion of law used by the physicists is indeed in trouble: since there can be observationally identical models of different sets of laws, the claim that a certain law obtains must go beyond what can be observed. The positivists had either to rework the notion of a law or abandon it altogether.

But no one is a positivist any more, and the shortcomings of verificationist theories of meaning are so well known as not to bear repeating. It is odd, then, that contemporary philosophers should flock to the banner of Hume.

Lewis does not announce himself a positivist, and presumably would be embarrassed at the association. The semantic theory that underlies Hume's own views has been thoroughly discredited. Why should one have 'Humean scruples' any more?

5. OTHER POSSIBILITIES

Justifications for Humean Supervenience can be divided into four categories: semantic, epistemological, methodological, and prejudicial. Semantic considerations, as we have just seen, tie the very meanings or truth conditions of sentences to matters of observable fact in such a way that any attempt to even make a claim that goes beyond those matters of fact becomes mere *flatus vocis*. Epistemological arguments can grant the *meaning* of claims that do not supervene on matters of particular observable fact, but still insist that such claims can never be *known*, or perhaps even *reasonably believed*, and so are practically idle. Methodological claims invoke some widely accepted methodological principle that implies a preference for Humean over non-Humean theories, usually on the basis of some sort of parsimony. Finally, the prejudicial stance simply declares facts that go beyond the totality of local, particular, and non-nomic facts (plus space-time) to be weird or spooky or strange. These categories are obviously somewhat arbitrary, and there is much overlap between them. The positivists, for example, consciously conflated matters of semantics and epistemology. But these four categories provide a reasonable framework for our enquiry.

Lewis, as we have seen, pegs his defense of Humean Supervenience to materialism, i.e. physicalism, but that provides no motivation for the reduction of laws to something else. Hume's and Carnap's semantic views are thin reeds, to say the least. Here is what John Earman has to say when the question of justification comes up (see quotation cited above):

The well-known motivations for (E1) fall into two related categories. There is ontological argument, intuition, and sloganeering ('The world is a world of occurrent facts'), the three often being hard to distinguish. Then there are epistemological arguments and threatenings, the most widely used being the threat of unknowability, based on two premises: we can in principle know directly or non-inferentially only (some subset of) occurrent fact, what is underdetermined by everything we can know in principle is unknowable in principle .... The argument connects back to the ontological if we add the further premise that what isn't knowable in principle isn't in principle. (Earman 1984, p. 195)

Earman's characterization is telling in at least three ways. The first is that Earman himself, although sympathetic to the empiricist position, does not directly endorse any of these motivations. Second is the introduction of the term 'occurrent' into the discussion (but see also how Earman finesses the question of the meaning of 'occurrent' in the passage cited earlier). If only local matters of particular non-nomic fact, and logical combinations of them, are 'occurrent', then accepting anything that goes beyond these is accepting something non-occurrent. But 'non-occurrent' sounds suspiciously like 'not really happening', or perhaps 'not really there'. The exact meaning of 'occurrent', and the appropriate conditions for its use, deserve more notice.

The standard means of explicating the notion of an occurrent property is to contrast it with a *merely* dispositional one, fragility being the usual example. It is true that the window is fragile, and true that it is massive, but the ontological status of massiveness and fragility are quite different. Massiveness is just a matter of how the window *is*, while fragility is a matter of *how it would behave under certain circumstances*. One then asserts that in order to be legitimate, the truth conditions for claims about dispositional properties must ultimately reduce to matters of occurrent fact: there are no 'free floating' dispositions. This line becomes a bit embarrassing if dispositional accounts of apparently occurrent properties are offered (e.g. if having a mass is analyzed in terms of how an object would behave if subject to a force), but let's leave this wrinkle aside.

Insofar as 'non-occurrent' carries the implication 'not really there in its own right', then the use of it *at the beginning of an ontological discussion* is obviously question-begging. I believe that there are laws at every point of spacetime, and that the laws cannot be reduced to facts about the total physical state of the universe. In some straightforward sense, I therefore believe that the laws are occurrent. If someone replies that on their understanding of 'occurrent', laws just aren't the sorts of thing that *could* be occurrent, then I would respond that on their understanding of 'occurrent', the inference from 'non-occurrent' to 'in need of reduction' is invalid.

What is the real problem with fragility? Fragility is not as 'real' a property as mass for the simple reason that we take it to supervene on fundamental physical and chemical properties. Once you have specified the exact physical composition of the window, *and given the physical laws*, it follows by analysis that the window will break if a rock is thrown at it with sufficient force. Fragility is not a *fundamental* physical property, in that two pieces of glass cannot be physically identical save for their fragility: if one is fragile and the other isn't, then there must be some other difference between them at the level of their physical composition. Mass and charge *are* fundamental physical properties in that particles *can* differ solely in mass (e.g. electrons and muons) or solely in charge (e.g. electrons and positrons) without there being any further physical difference which accounts for this. An ontology that accepts fragility at a fundamental level, alongside mass and charge, is both too bloated and in danger of contradiction. Too bloated because all of the behavior of the window can be predicted and explained (we believe) from its description in terms of mass, charge, etc. In danger because it suggests that two windows could agree in all other physical respects but differ in fragility, which we take to be false. It is a further notable fact that fragility, unlike mass and charge, does not figure in the fundamental laws of physics.

The reduction of fragility, though, is of necessity a reduction to both fundamental physical state *and law*. And none of the reasons for believing that fragility can be explicated in terms of physical state and law can be used to argue that law itself can be explicated in terms of physical state alone. As we have seen, while science is in the business of explaining things like the fragility of a piece of glass, it is not in the business of giving reductive analyses of the laws of nature.

The third revealing aspect of Earman's survey of motivations is the introduction of epistemology, together with the qualifier 'in principle'. On the assumption that all atomic observations are observations of matters of local, particular, non-nomic fact and their space-time relations, the non-supervenience of laws on the total physical state of the universe implies that there could be two universes that are observationally identical, yet differ with respect to their laws. What it would 'feel like' to live in either of these universes would be exactly the same, yet the facts which obtain in them would differ. Hard positivists would deny that any such seeming difference could be real, or the words used to describe the apparent differences meaningful, but we have let hard positivism go its way. Still, there is something slightly spooky about facts that go beyond the evidence. Let's see if we can exorcize this ghost.

Let us first recognize that the existence of facts which are not determined by the complete totality of all observations, past, present, and future, is commonplace. No doubt, Socrates had a blood type. Also, no doubt, one could not deduce that blood type from a complete catalogue of every observation that has been or will be made. No test of the requisite type was ever made on Socrates, and, doubtless, no remains that could be identified as his will be subject to such a test. Socrates' blood type is now, and will always be, beyond our ken.

Of course, only a lunatic would conclude that the ontological status of Socrates' blood is thereby affected, that he (miraculously) had no blood type at all, or an indeterminate one. No doubt there was a fact of the matter about his blood, despite our irremediable inability to know it. So the question arises: what difference does it make to ontology that a particular fact cannot be deduced from actual evidence?

This question is likely to strike one as trifling. After all, as Earman's quote indicates, the issue should not be what is, or was, or will be known, but rather *what is knowable in principle*. Socrates never had his blood tested but he *could* have had it tested (in some sense of 'could'), and the test *would have* revealed the blood type. Socrates' blood type in not ontologically worrisome because it was knowable in principle.

This sort of response is so common and well entrenched that it takes an effort to see how utterly bizarre it is. There are straightforward epistemological problems about Socrates' blood type: we don't know what it was and we never will. The question is whether this epistemological problem should have any implications for ontology at all. The commonsense answer is no. And the philosopher wishing to maintain the commonsense response, but still desiring to link ontology to epistemology, makes the link not to *actual* evidence but to *merely possible* evidence. But if I am somehow worried about the ontological status of Socrates' blood type for epistemological reasons, how will the invocation of *counterfactual assertions about merely possible observations* allay those worries? Is anyone really supposed to say: 'Well, I was initially concerned about how "Socrates had blood type O" could have a determinate truth value, but my fears have been allayed by giving truth conditions in terms of counterfactuals, for I have no qualms about accepting that the counterfactuals have definite (but unknown) truth values'!? This has everything exactly backwards: we think that there is a determinate (but unknown) fact about how such tests would have come out exactly because we think there is a determinate (but unknown) fact about what the blood type was, and that the testing procedures would have revealed it. Relying on the counterfactuals to somehow validate the use of plain indicatives ('Socrates' blood type was O') is both baroque and self-defeating.

On any view, the laws of nature are not knowable *in fact* in the sense that they do not follow deductively from all the observations there have been, are, and will be. On my telling, they are not knowable even *in principle* in that they do not follow deductively from everything in the history of the universe that *could* have been observed. What follows? If severe ontological worries follow unknowability in principle, it is hard to see why the same worries won't accompany unknowability in fact, and that fate strikes most of the events in the history of the universe. Indeed, knowability in fact is already a philosopher's fantasy: there is no catalogue of observations being recorded in the epistemologist's book of life, such that all actual observations are accessible to us. The practical problem of epistemology is inference from data we have to hand, and from this essentially all of the universe, including the recent past and all of the future, is underdetermined. If ontology follows epistemology this far, there will be precious little of existence left to us at all.

There is a milder form of epistemological upset than unknowability in principle. This form admits that we cannot be constrained in our ontology only to those things that follow from our observations, so ampliative inferences must be allowed. But still, we should distinguish those ampliative inferences that can be (in some instances) checked from those which cannot. Induction, for example, is always a leap beyond the known, but we are constantly assured by later experience that we have landed safely. Inference from observations to laws of nature, on the other hand, can never be later vindicated since all the data give us are patterns of particular matters of fact.

This is a serious concern, and a complete answer to it is beyond our present scope. Suffice it to say that this sort of skeptical worry about ampliative inferences also strikes theoretical entities such as quarks, which are never directly observed. Suspending belief about unobservable entities is a time-worn strategy, albeit one uniformly rejected by actual scientific practice. And if one is agnostic about things like electrons and quarks, not to mention quark color and flavors, then the question of laws of nature, as understood in scientific practice, is already decided, since the laws of fundamental physics are couched in those terms. But this is not the sort of view that a scientific realist such as Lewis would endorse.

Our epistemological access to most matters of particular, local, non-nomic fact already demands ampliative inferences that can never be fully vindicated by experience. That is why an *epistemological* defense of Humean Supervenience cannot be satisfied merely with a reduction of laws to total physical state, or total local physical state, but must insist on *total observed physical state*. But this so handcuffs ontology that few would be willing to accept it. Thus, as cited above, Earman writes:

> There may also be general occurrent facts (I think there are), but these are also parasitic on the singular occurrent facts. Conservative empiricists may want to

restrict the antecedent of (E1) so as to range only over observable facts while more liberal empiricists may be happy with unobservable facts such as the fact that quark q is charming and flavorful at t. In this way we arrive at many different versions of the loyalty test, one for each persuasion of empiricist. (ibid.)

Lewis would clearly be a liberal empiricist on this telling, but just what sort of an empiricist is a liberal empiricist at all? Quarks and their flavors, as well as wavefunctions, are neither the data of experience nor constructs from them. Of course, we only believe in quarks and wavefunctions on the basis of observations, not a priori, but the same is true of laws of nature, such as the relativistic field equations. Lewis's penchant for the singular, local, and non-nomic cannot be given an epistemological foundation, for attempts to appeal to epistemology will banish a lot more than non-supervenient laws.

There is one final methodological consideration that might be used to defend Lewis's version of Humean Supervenience: Ockham's Razor. Any physicalist already admits the total physical state of the world into his or her ontology, and thereby also anything that supervenes on the total physical state. On the general principle that less is more, why not try to live as economically as possible, preferring not to go beyond the total physical state unless pushed?

The question is how much of a push is needed. While the Razor is demonstrably good methodological advice in some circumstances (in particular, when it councils higher credence to explanations which posit a single cause to multiple events that occur in a striking pattern over explanations invoking coincidental multiple causes), it is hardly universally accepted. Very few physicists, for example, still pursue pure relationist theories which eschew space-time for spatio-temporal relations between material objects. The relationist ontology is strictly a subset of the usual physical ontology, since those who countenance space-time also accept that there are material objects with spatio-temporal relations, they simply do not think that to be all there is.[6] But for all that, physicists are not rushing to replace field theory with action-at-a-distance particle theories, in hopes of reducing ontology. The relationist theories are just too complicated and contrived.

[6] One can argue for an understanding of spatio-temporal structure in which there are no spatio-temporal relations, if by 'relations' one means something whose existence only requires the existence of the relata. On this understanding, no spatio-temporal relations can be 'skimmed off' the substantivalist ontology. See 'Suggestions from Physics for Deep Metaphysics', Chapter 3, this volume.

Similarly, one could try to live within the confines of the total physical state of the world, and find good enough substitutes for laws that can be defined in those terms, but why do so? One would take fewer chances thereby (though barely fewer, since the total physical state is itself largely unknown), but at what price? One could also refuse to assert or deny counterfactual claims in general, and thereby take fewer risks of being wrong, but it seems like a poor and miserly existence.

Lewis's own stated motivation is robust and healthy. Ontology is the general account of what there is, and our knowledge of what there is is grounded in empirical science, not in a priori speculation or prejudice. Philosophical accounts that force upon us something that physics rejects ought to be viewed with suspicion. But equally suspect are philosophical scruples that rule out what physics happily acknowledges. As we have seen, contemporary physics posits physical facts that are Non-Separable. What grounding could a preference for Separability have to suggest that we ought to warp either the physics itself, or our account of space to accommodate Separability? And physics has always postulated that there are laws without suggesting that they supervene on or reduce to matters of particular non-nomic fact. Hume, armed with an empiricist semantics, had reason (by his own lights) to be worried. But we no longer accept Hume's account of concept formation and its allied account of linguistic meaning and truth. So the question remains: why be Humean?