

# Laws of Nature

Peter Smith

April 24, 2006

## 1 Preamble

Where to begin? I'll take three books from my shelves. First, now nearly forty years old, a little book of television lectures by the great physicist Richard Feynman, *The Character of Physical Law*. He talks about the laws of motion, the inverse square law of gravitation, conservation laws, symmetry principles and the various ways these all hang together. Feynman obviously takes it that it is a prime aim of science to discover such laws. But what are laws? He writes – and this is about his one and only shot at a characterization at the level of abstraction that we might think of as philosophical –

There is a rhythm and a pattern behind the phenomena of nature which is not apparent to the eye, but only to the eye of analysis; and it is these rhythms and patterns which we call Physical Laws.

Which isn't, I think we have to agree, terrifically helpful. It perhaps suggests a rather Humean, 'regularity', view of laws. Laws describe regularities and patterns to be found in the world: and what makes a regularity a law is that it picks up on patterns underlying the messy surface phenomena, deeper patterns that give us understanding and enable us in various ways to predict and control the world. But it could hardly be said that Feynman is wedded to Humeanism (and he's pretty contemptuous of philosophy in any case).

Second, I'll take down a much more recent book – Roger Penrose's massive, thousand page, *The Road to Reality*, which is subtitled 'A Complete Guide to the Laws of the Universe'. I haven't read it from cover to cover: but as far as I can see Penrose says even less than Feynman about what we mean when we say that some proposition expresses a law. From time to time he talks of laws governing the universe or shaping the universe – as if the universe is in some sense answerable to the laws (so the laws aren't just interesting regularities – as on the Humean view – but in some way underpin and are responsible for the regularities). But if there is very occasional anti-Humean rhetoric in Penrose, I doubt that it is intended to be taken any more seriously than (for example) his vivid talk about photons 'choosing' its route through a beam-splitter.

If there is a big conceptual question that bugs Penrose, it is the same one that bugs Feynman: it isn't about the metaphysical status of laws but about the why the fundamental physical laws should be, as they seem to be, essentially *mathematical* in character. So the familiar philosophers' debates between Humeans and anti-Humeans about the nature of laws just passes them by – it doesn't impact at all on their conceptions of the scientist's task in discovering the laws. Nor is it at all clear why it *should*. And I guess what goes for Feynman and Penrose, goes for nearly every working scientist: they happily proceed seemingly without engaging with the issues – and without even saying anything that would weigh on one side or the other – in philosophers' debates about the laws.

How worrying should that be for philosophers of science? You might say: it's not worrying at all – science is one thing, philosophy of science is something else, with its own problematic and its own methods. Or, on the other side, you might say: this kind of 'disconnect' between the philosophers and the physicists should make us a bit suspicious about whether there is anything of *real* interest at stake in the standard philosophers debates about laws – perhaps idle wheels are spinning. Well, we certainly can't take a view on that now, not before looking at some aspects of the philosophers' debates. I'm just flagging up the worry (but it seems to me a real one).

Penrose (like Feynman) assumes, take it as evident that, the prime task of science is to discover the laws. Most philosophers assume that too. Here's Hugh Mellor:

Certified laws of nature are the primary products of scientific thought and observation.

So the title of the third book from my shelves, by Ronald Giere, is very striking. Giere is a significant philosopher of science – and one whose work I find very congenial (the views about the nature of dynamical models in my *Explaining Chaos* could reasonably be thought of as a souped-up version of Giere's). The title of his 1999 collection of essays, however, is *Science without Laws*. But he surely can't be claiming that scientists don't do, or shouldn't do, what Feynman and Penrose's books reveal them as doing – that is looking for, and discovering, general principles describing what happens. So what on earth can be going on?

Well, we need to make a key distinction – let's say for the moment that laws in the anodyne, physicists', sense are propositions like Newton's laws of motion (or their relativistic versions), the law of conservation of energy, the law that boson number is conserved in particle physics, Snell's law in optics, Maxwell's laws in electrodynamics and that sort of thing. We'd better be cautious in saying that such propositions are strictly true: so let's borrow Mellor's word and call them 'certified' (leaving it open just what certification amounts to). Then *of course*, science looks for certified 'laws' in the anodyne sense of looking for propositions as scientifically successful as the familiar laws of motion, the conservation laws, and so on. Giere wouldn't dissent. What Giere's beef is about, in the title essay of his collection, is whether such certified propositions should be thought of as capital-'L' Laws in some weightier, philosophers', sense – a sense tied up with notions of being (for a start) true, absolutely universal generalizations, perhaps with some modal oomph. Capital-'L' Laws in the philosophers' sense, according to Giere, are myths that play no part in the real business of scientific theory construction.

Well, again, we can't take a view on that just yet, not before thinking some more about what loading might get built into the philosophers' notion of Capital-'L' Laws and whether that loading is problematic. But if Giere is right, that would certainly help explain how Penrose can talk about the laws of nature for a thousand pages without seemingly engaging with the philosophical debates about capital-'L' Laws. If, as Giere thinks, the laws (in the anodyne, working-scientists' sense) are simply ill-represented as being capital-'L' Laws, then no wonder there is a disconnect between scientists' concern with the small-'l' laws and typical philosophers' debates about the mythical(?) capital-'L' Laws.

## 2 One standard view of capital-L 'Laws'

I'll outline one core package of philosophical assumptions about laws of nature. I'm assuming that you will have already encountered some of the debates, and so I hope that you'll recognize the package as containing extremely familiar lines of thought. You'll probably spot straight away that there is a once-popular line that doesn't fall into the standard package, and I'll mention that in just a moment. But you'll also recognize that the familiar debate between latter-day Humeans like

David Lewis and anti-Humeans like David Armstrong or Michael Tooley takes place against the background of a shared commitment to the core package I'm about to describe.

One preliminary remark. Talk of laws on nature can oscillate between talk of true law-propositions and talk of the worldly facts that make those propositions true. For the moment, we should be picky, and keep proposition talk and truth-maker talk clearly separated. (Once we've reminded ourselves of the points at which it is important to keep these distinguished, we can later fall back into talking of laws of nature, letting context and common sense decide whether propositions or facts are in question.)

So here's the package of assumptions. I'm not aiming for a minimal set of independent assumptions here – some redundancy is fine if it helps characterize the position in a more accessible way.

1. (a) Law-propositions aim directly to describe facts in the world.  
(b) Law-propositions are either true or false, and if true, are contingently true.  
(c) True law-propositions are made *true* by worldly facts that obtain independently of the thoughts and activities of scientific enquirers (cf. ethical truths).
2. True law-propositions have the status of *laws* independently of the thoughts and activities of scientific enquirers. [Contrast 1(c): that claim was about truth, this one is about the status of being a law.]
3. True law-propositions not only tell us what always does happen in various circumstances, but also imply truths about what *will* happen if we were to change the world in certain ways, and what *would* have happened had the world been changed in other ways. As the point is often put: law-like generalizations are non-accidental and sustain counterfactuals.
4. Law-propositions are inductively confirmable by their instances (cf. 'gruesome' propositions), have predictive force, and figure in explanations of phenomena within their scope.
5. (a) It is possible to get to know that various propositions, *if* true, have the status of laws.  
(b) It is possible to get to know that various law-propositions indeed *are* true.
6. It is a primary task of science to discover true law-propositions as partially characterized in (1) to (5).

We can take (1) to (5) as a partial characterization of a philosophical conception of capital-'L' Laws; and then (6) is the claim which Giere rejects, namely that science is about discovering capital-'L' Laws. But Giere and a few other dissidents like van Fraassen aside, for most philosophers of science, the key philosophical issue about laws is to explicate how this core package of intuitive presumptions can possibly be true together – and in particular to give a coherent story that comfortably reconciles our *ontology* of laws (our story about the sorts of facts that make law-propositions true) and our *epistemology* of laws (our story about how we get to know those facts).

Five initial comments on elements of the core package may help to fix ideas.

- i. The thesis (3) is typically read as the claim that, if it is a law that *As* are *Bs*, then this entails that all *As* are *Bs*. But note that this claim that law-propositions entail universal generalizations is compatible with both the Humean view that law-propositions just *are* universal generalizations (generalizations that have some additional property of "loveliness" that suffices to distinguish them from merely accidental generalizations), and with the view that

law-propositions have some other logical shape – e.g. express relations between universals. Nothing in (3) begs the question for or against the Humean position, nor does it beg the question for or against Armstrong’s view that the law-proposition that *As* are *Bs* claims that a relation of necessitation holds between the universal *A* and the universal *B*.

- ii. Similarly, nothing in (1) begs the question for or against the standard contemporary Humean, and nothing begs the question for or against the Armstrongian. The Humean and the friend of universals have different views as to the nature of the worldly facts which make law-propositions true when they are (corresponding to their different views about the logical shape of those propositions); but they agree that there are such truth-making facts.
- iii. Now to mention a familiar, once-popular, position that isn’t an elaboration of the core package. A *subjective* Humean view says that it is in a broad sense subjective (not whether a law proposition is true but) whether a given true proposition should count as a law. Consider, for example, the line that a true generalization is a law just if it is relied on in certain ways by scientists in their activities. This line obviously *would* be incompatible with (2). So subjective Humeans – Ayer for example – don’t sign up to the full standard package.
- iv. However, what is usually thought to be the strongest modern Humean position – i.e. the Ramsey/Lewis view that the laws are the true generalizations that are “lovely” because they are integral to the best systematization of the truths into a deductive system – *is* compatible with (2). Or at least, it is compatible to the extent that the notion of ‘best systematization’ can be made independent of judgements of scientific enquirers. (Of course, a familiar complaint against the Ramsey-Lewis view is that the idea of ‘best systematization’ is too tinged with subjectivity. But the fact that the complaint is made, and is assumed to be an objection, is witness to the fact that most philosophers indeed assume that (2) should be true. And the fact that defenders of the Ramsey-Lewis view usually think that they should endeavour to fight off the objection is witness to the fact that latter-day Humeans also mostly accept (2).)
- v. One could in theory combine epistemic pessimism – the view that we can’t (or can’t often) get to know laws – with the view that the business of science is to get to know laws. We’d then have to conclude that science wasn’t, or wasn’t often, successful in its aims. Which wouldn’t exactly be an attractive position to end up. Those who want to say that discovering laws of nature is of the essence of the scientific enterprise usually also want to be epistemic optimists too (despite Popper!). That puts us in a familiar bind. On the one hand, on the Humean view that the laws are just “lovely” generalizations, we might hope that the epistemology of laws won’t be especially problematic (or at least, no more problematic than the epistemology of reliable inductive reasoning generally). However, while the Humean’s thin ontology might be epistemically comforting, there are familiar worries about whether the Humean can ever get the ties between laws and counterfactuals right (I’ll come to these in a moment). Certainly, those like Armstrong in the anti-Humean camp think these worries are insuperable. Their rival story about the sort of facts that make law-propositions true is supposed to help us out precisely here. But then the anti-Humean’s richer ontology gives us epistemic problems. To take Armstrong’s version, relations of necessitation between universals are supposedly one thing, and mere general correlations are something else. Our empirical knowledge of world at best gives us direct knowledge of the general correlations (and reflection perhaps gives us the organization of those correlations into a Ramsey-Lewis system); yet how does that show that genuine relations between universals hold, rather than – so to speak – just the unsupported

Humean web? (The pattern is entirely familiar: easy epistemology but a reduced ontology inadequate to its explanatory role vs. a richer ontology giving rise to epistemic problems.)

So much then for some initial thoughts and clarifications. The point to emphasize is perhaps simply this: what I have called the standard core package of philosophical assumptions is the framework within which most contemporary Humean vs. anti-Humean debates get fought out. Some might think that the inconclusiveness of those debates is perhaps some reason for reassessing the standard core package in one way or another. But more on that later.

### 3 Trouble for (‘objective’) Humeans

I certainly don’t propose to follow all the ins and outs of the debates between (objective) Humeans and anti-Humeans pursued within the framework of the standard package. If, for a start, you want to look at a fair-minded but powerful composite critique of the Humean side and a vigorous defence of one kind of anti-Humean, see Armstrong’s characteristically clear *What is a Law of Nature?* And for a more recent, again accessible, textbook treatment, see Stathis Psillos, *Causation and Explanation*.

However, I will start by considering one line of anti-Humean argument. The basic thought has been pressed by Michael Tooley and John Carroll, but I’ll try to develop the argument from scratch in my own way. I find this argument interesting and rather compelling. Psillos has a shot at criticism, but his counters are surprisingly feeble. So I find myself having to give the game to the Tooley/Carroll argument – somewhat reluctantly, for it goes against my austere Humean instincts: so if you can come up with a better refutation than Psillos’s, then that would be good.

A preliminary comment. The argument we’ll be considering is a modal argument – the basic thought is that on a Humean view something is impossible which, intuitively, certainly seems possible. For vividness, I’ll explain the argument using ‘possible world’ talk: so I’ll be saying something like “on a Humean view it makes no sense to say that there is a pair of distinct possible worlds which differ thus and so; but intuitively that does make perfectly good sense.” But the use of ‘possible world’ talk rather than more common-or-garden modal talk is – as far as I can see – inessential to the argument (it is just for vividness, and I’m most certainly not intending to presuppose a Lewisian realism about worlds).

Here, then, is the anti-Humean argument I want to consider, in bare bones form:

1. On a Humean view, possible worlds which match on the facts about what occurs within those worlds – and hence on the generalizations which hold in those worlds – must match on the laws. (The laws supervene on the true generalizations.)
2. But worlds which match on what actually happens within them can differ in respect of their laws.
3. So the Humean view is unacceptable.

Premiss (1) is often taken to be pretty much a definition of what makes a view Humean in one standard sense. (I guess it shouldn’t need saying, but I’ll say it all the same: some think that the so-called Humean view may no more be David Hume’s view than the Holy Roman Empire was Holy, Roman, or an Empire. But I’m certainly not going to fuss about absolutely standard nomenclature.) John Earman for example calls assent to (1) ‘the empiricist loyalty test on laws’. And take, for example, the Ramsey-Lewis version of the Humean view. The laws, in the sense of the true law-propositions, are those generalizations that feature appropriately in the best systematization (or all

the best systematizations) of the generalizations that are true of the world. If two possible worlds match on the generalizations that are true of them, then they will match on the best systematization of those generalizations, and so must match on the laws.

Now, in fact I suspect that there is a version of *subjective* Humeanism which doesn't need to accept (1). I'll talk about this later. So maybe our argument isn't an argument against every form of Humeanism: but since (1) is accepted by most contemporary Humeans – and is treated as definitive of their position – our argument is at least an argument against standard modern versions that accept what I called the core package.

So Premiss (2) is the substantive one. Here's a story to support it.

Consider a world  $W_1$  in which there are five different kinds of ultimate particle, say the  $\alpha$ ,  $\beta$ ,  $\gamma$ ,  $\delta$ ,  $\epsilon$  particles. And suppose that, as it happens, in the course of the history of the world, an  $\alpha$  and  $\beta$  never, ever, get to interact. The other 14 types of interaction ( $\alpha - \alpha$ ,  $\alpha - \gamma$ ,  $\dots$ ,  $\delta - \epsilon$ ,  $\epsilon - \epsilon$ ) all occur, governed by the fourteen general laws  $L_{\alpha-\alpha}$ ,  $L_{\alpha-\gamma}$ ,  $\dots$ ,  $L_{\delta-\epsilon}$ ,  $L_{\epsilon-\epsilon}$ . Those fourteen fundamental laws of nature are, we'll suppose, entirely independent of each other – there is no Grand Unified Theory which further systematizes them. They are ultimate brute laws.

However, the story continues, surely there could *also* be another uninstantiated fundamental law governing this world – a brute basic law as unrelated to the other fourteen as they are to each other – which dictates what *would* have happened had an  $\alpha$  and  $\beta$  got near enough to interact. By hypothesis, the facts of about the interactions which actually take place in the world which fix the truths about the generalizations true of the world will at best fix (on a Humean story) the other fourteen laws. On the Humean view, there is no generalization and no law to govern the fifteenth interaction. Yet surely there could be such a law. And there could be another possible world  $W_2$  with the same occurrences – so again where  $\alpha$  and  $\beta$  particles never, ever, get to interact – but where the brute law governing their interaction would be different, and something different would have happened if things had gone in such a way that the particles do interact.

So we have just perfectly coherently imagined two possible worlds where the overall laws are different but where everything that happens in the worlds is the same.

That seems to establish Premiss (2), and so the anti-Humean argument is apparently secured.

To resist the argument, the Humean will need to dig his heels in and insist that in our world  $W_1$  without  $\alpha$ - $\beta$  interactions there simply can't really be any fact about what would have happened if an  $\alpha$  particle had encountered a  $\beta$  particle. But how plausible is that? Here's just one way of bringing out the unattractiveness of the Humean response.

Consider a world  $W_3$  which is very like our world  $W_1$  except that here all fifteen inter-particle interactions do occur, and take place in law-governed ways. So there is a law  $L_{\alpha-\beta}$  governing this world. Now make the following three suppositions – which all seem to be untendentiously possible:

- i. The other fourteen laws governing inter-particle interactions are shared by  $W_1$  and  $W_3$ .
- ii. Some of these other shared laws governing particle interactions are chancy.
- iii. We said, things in fact developed in  $W_3$  so that there were  $\alpha$ - $\beta$  interactions. But compatibly with the chancy laws, things *could* have so developed that in fact there were no  $\alpha$ - $\beta$  interactions at all. To take a vivid case, let's suppose that the existence of  $\alpha$ - $\beta$  events in  $W_3$  crucially depends on just one chance event.

Now, on the position of the Humean who digs in his heels who says that there can be no fundamental, underived, law governing  $\alpha$ - $\beta$  interactions if there are no such interactions, it seems that we would have to say the following.

1. There is a law  $L_{\alpha-\beta}$  in  $W_3$ .
2. But had a nomically possible, quite ordinary, chance event in  $W_3$  not happened so that we were in a world like  $W_1$  without  $\alpha$ - $\beta$  interactions then  $L_{\alpha-\beta}$  would not have been a law.
3. So whether  $L_{\alpha-\beta}$  is a law in a world can depend on an ordinary chance events.

But isn't this upshot weird? Again I find that counter-intuitive (if I understand the idea of capital-'L' laws at all). We might say: isn't the very idea of a law the idea of – so to speak – a fixed framework within which a world unfolds (maybe indeterministically if some of the laws are chancy laws); and with different initial conditions, or with different chancy happenings, worlds sharing the same laws can unfold differently. A different permissible set of initial conditions or a different chance occurrence in accordance with the laws shouldn't be able to delete a law!

## 4 Responses

The defender of Humean supervenience could of course continue to dig in his heels. But he's going to have to be careful. For the Humean can't get off the hook by rejecting *all* uninstantiated laws. Take Newton's first law in the form 'the motion of a particle, upon which no forces act, must be uniform and in a straight line' (I've lifted that statement straight from Penrose, so I can't be accused of cheating.) Then we have every reason to think that the law, thus stated, is vacuous, since every particle is subject to gravitational forces. (And if you protest and say that what is meant by 'no forces' is 'no net forces' then it is still the case that, for all we know, the law is vacuous).

Now, of course, Newton's first law is derivable as a special case from the second law (the law  $f = ma$  yield  $a = 0$ , so uniform straight motion, when  $f = 0$  for massive particles). So the Humean can allow Newton's first law even if vacuous because it derivable from a law which in turn implies lots of instantiated laws of the shape  $f_1 = m_1 a_1$ ,  $f_2 = m_2 a_2$ , etc. In a phrase, Newton's first law is a law because – although it is uninstantiated – it keeps good company, and falls out of the best Ramsey/Lewis system of generalizations. In our example of the  $\alpha$ - $\beta$  interactions, according to the Humean, there isn't good enough company for a vacuous generalization about such interactions to keep to count as a law. And I suppose the fundamental Tooley/Carroll complaint is: why should the amount of company a generalization keeps make it a law or keep it from being one? isn't it a more intrinsic matter than that? So which the Humean says 'no, it isn't'.

Interestingly, Psillos does not consider the straight bullet-biting response, but offers two other defenses of Humean supervenience against the Tooley/Carroll type of argument. Actually, Psillos fumbles the version of the argument he reports. Carroll's version has a world  $W_1$  with  $X$ -particles and  $Y$ -fields, and a law about how passing through a  $Y$ -field affects and  $X$ -particle – although, as it happens in  $W_1$  they never do interact. And then another world  $W_2$  with  $X$ -particles and  $Y$ -fields, and a different law about how passing through a  $Y$ -field affects and  $X$ -particle – although again, as it happens, just as in  $W_1$  they never do interact. We might suppose then that the happenings in  $W_1$  and  $W_2$  are the same, though the laws are different. Oddly, Psillos reports this example as involving the idea that 'in both  $W_1$  and  $W_2$  there never have been, and never will be, fundamental particles of type  $X$ , nor fields of type  $Y$ '. But that's just a misreading: however let that pass.

Psillos's first defence on behalf of Humean supervenience is as follows:

The [Ramsey/Lewis] approach does allow for uninstantiated laws, provided that their inclusion in the best system enhances its strength, without detracting from its simplicity. In the thought-experiment we are discussing, there are two mutually incompatible uninstantiated laws that hold in two different but Humean-indistinguishable, worlds. However, the Humean could argue, had these laws been instantiated they would have given rise to different regularities. That is, there would have been an actual (Humean) difference in the two worlds. Hence, which of the two laws holds would be determined by non-nomic facts in worlds  $W_1$  and  $W_2$ .

But that is just thumpingly point-missing. Of course, had there been interactions between  $X$ -particles and  $Y$ -fields, there would have been Humean differences between the relevant worlds. But then we would have been in *different* possible worlds – no longer in the original possible worlds  $W_1$  and  $W_2$ , but in different worlds  $W_4$  and  $W_5$  where  $X$ -particles and  $Y$ -fields meet up after all. The point remains that  $W_1$  and  $W_2$  are Humeanly indiscernible, yet it seems we can coherently think of them as having different laws: Psillos's remark doesn't affect that point at all.

His second defence on behalf of Humean supervenience borrows a thought from Armstrong (who, though of course an arch anti-Humean, doesn't much like the Tooley/Carroll argument). Here's Psillos:

For [Armstrong], uninstantiated laws are 'concealed counterfactuals' of the form: 'if, contrary to fact, certain sorts of things existed, then these things would obey a certain law'. The truth or falsity of such counterfactuals 'depends wholly upon actual, that is instantiated, laws.' So ... [Humean supervenience] is not threatened by Carroll's counter-examples.

Why is the Humean supposed to be off the hook? OK, if we assume that the only true uninstantiated laws and those that are grounded in instantiated laws, then there can't be Carroll-type laws which by hypothesis are uninstantiated and ungrounded in instantiated laws. But one man's modus ponens is another man's modus tollens. The Carroll response should be that the intuitive possibility of his cases undermines Armstrong's assumption.

Of course, Armstrong's assumption doesn't come from nowhere – it is supposedly unwritten by his general theory of universals and the nature of laws. But two points about that. First, it is in fact unclear whether his treatment of uninstantiated laws is in fact dictated by his general theory of laws. And second, even if it were, the *Humean* could hardly appeal to Armstrong's anti-Humean account of laws to defend his position, as he shouldn't believe Armstrong's premisses.

The two considerations that Psillos adduces therefore don't dent the Tooley/Carroll argument. So far, then, I still find that argument very telling.

## 5 Summary so far

Here's one kind of Humean view about what makes for a law of nature:

1. A law is a true universal generalization with some additional feature of 'loveliness'.
2. Whether a universal generalization has some this additional feature of 'loveliness'
  - (a) is an objective matter, independent of the thoughts and activities of scientists;

- (b) obtains in virtue of the some non-nomic (metaphysically light-weight) facts about this and other generalizations.

We noted that the Ramsey/Lewis account of laws is an ‘objective’ Humean position of this type. For on that picture, a law is essentially (1) a true generalization which fits in the right way into the best system of true generalizations, where (2a) what counts as the best system depends on (supposedly) objective features of economy and power. Evidently, the Ramsey/Lewis position respects (2b), and so underpins the closely related Humean supervenience thesis: possible worlds which match on the non-nomic facts about what occurs within those worlds must match on the laws (for evidently, if worlds match on the non-nomic facts, they match on the true brute generalizations, and match too on the best system of these generalizations).

If we are going to accept (2) while adhering to the Humean thought that there is no more to the world than what is fixed by the arrangement of physical properties across space and time, it is difficult to see how we could do any better than the Ramsey/Lewis picture. The only factual ingredients available to the Humean, over and above the fact that *this* generalization is true, would seem to be facts about how the given generalization fits together with *other* generalizations in a predictively powerful network (and the point of laws, for the Humean, presumably just is to give us a lot of predictive power, so what else but to consider how generalizations can be put together in a simple by highly informative package?). So the Ramsey/Lewis picture – already prefigured in Mill – will surely be the ‘objective’ Humean’s best shot.

So far, then, I’ve considered *one* line of attack on this ‘objective’ Humean picture, based on pumping the anti-supervenience intuition that there can be possible worlds which agree on the non-nomic facts – agree on what generalizations obtain – but disagree on the uninstantiated laws. I’ll not now say more about that particular argument (which does, though, seem to me to be a powerful one). Here are three more arguments against the Humean view in its Ramsey/Lewis version.

## 6 Three more arguments

1. *Accidents again* Suppose a world with five types of elementary particles – the  $\alpha$ ,  $\beta$ ,  $\gamma$ ,  $\delta$ ,  $\epsilon$  particles. Suppose that all fifteen types of interaction between particles is exemplified in the world, and suppose that in particular the following generalization is true of the world: every time an  $\alpha$  encounters a  $\beta$  they interact and they decay into (say) a pair of  $\gamma$  particles. But suppose too that not a great number of  $\alpha$ ’s in fact encounter a  $\beta$ . Still, the generalization is true, and although it doesn’t have many instances, it would seem that the generalization could belong with the generalizations about the other fourteen types of fundamental interaction, as one of the generalizations that belongs to a deductive system with optimal strength and simplicity. (Maybe, for example, this generalization – when added to the rest – gives the collection of generalizations about the fifteen types of interaction a simplifying symmetry.) So the generalization would then be a law on the Ramsey/Lewis view.

But for all that, couldn’t it still just be a cosmic coincidence that  $\alpha/\beta$  interactions result in a pair of  $\gamma$  particles? Surely so.

To press the point, couldn’t the law instead be that  $\alpha/\beta$  interactions have a certain low chance of producing a pair of  $\gamma$  particles, and it just so happens that the actual world realizes over its history the very low probability sequence of fifty  $\alpha/\beta$  interactions resulting in a pair of  $\gamma$  particles? (Someone wins the lottery, however improbably: in some possible worlds, the very unlikely combination of low-chance but law-permitted events happens – and we can’t rule out being in the unlikely world as a metaphysical impossibility.)

So, the argument goes, the Ramsey/Lewis view would deem to be laws what don’t have to be

laws, namely cosmic accidents that involve sufficiently interesting, basic, etc. events occurring in a neat enough pattern to get recorded in our best systematized set of generalizations about the world.

This is a kind of converse to our earlier argument. Previously we said that we could have true (because uninstantiated) generalizations which are laws – but which *aren't* on the Ramsey-Lewis view. Now we are saying that we could have true (and instantiated) generalizations which aren't laws because cosmic coincidences – but which *are* laws on the Ramsey-Lewis view.

2. *Counterfactuals* Laws sustain counterfactuals (that, we said last time, is agreed on almost all sides). And it is plausible to claim that whatever it is that makes the laws truly laws also should make true the true counterfactuals the laws support.

Can the Ramsey/Lewis theory account for that? Suppose in the actual world *All Fs are followed by Gs* is a true generalization, and belongs to the best system of generalizations. Then supposedly this in itself should imply that, had some object which isn't an *F* been an *F*, then it would have been followed by a *G*. But why? Consider three possible worlds:

- (a)  $W_a$ , the actual world where (let's suppose) it is a chancy matter whether an event is an *F*-event, but all *F*s that do happen are followed by *G*s.
- (b)  $W_b$ , a possible world which is just like  $W_a$  up to the occurrence of some pre-cursor event that *could* have led to an *F*-event in  $W_a$  but by chance *doesn't*: but in  $W_b$ , the pre-cursor event *is* followed by a new *F*-event, which is then followed by a *G*-event (as is every *F*-event in this world). This then leads to ramifying differences between  $W_b$  and the actual world  $W_a$ .
- (c)  $W_c$ , a possible world which is just like  $W_a$  up to the occurrence of some pre-cursor event that *could* have led to an *F*-event in  $W_a$  but by chance *doesn't*: but in  $W_c$ , the pre-cursor event is, by chance, followed by a new *F*-event, but this is *not* followed a *G*-event. In fact the only difference in events between  $W_c$  and the actual world  $W_a$  is the occurrence of that isolated *F*-event instead of a not-*F*-event. So all the other *F* events are followed by *G* events, but it isn't a *universal* law any more.

Now,  $W_c$  is much more like  $W_a$  as far as what happens in it is concerned. And ultimately, on a Humean picture, non-nomic facts about what happens in a world are the only facts about the world. So the natural Humean line would surely be to say that the more similar world to  $W_a$  is  $W_c$  where it is no longer the case that all *F*s are followed by *G*s. And if we evaluate the counterfactual 'if this non-*F* event had in fact been *F* it would have been followed by a *G*' by looking at close worlds, then the counterfactual comes out *false*.

So it seems that the Ramsey-Lewis view of laws by itself doesn't ensure that laws sustain counterfactuals.

Now, the issues here should actually be familiar ones. In a famous passage in his paper 'Counterfactual Dependence and Time's Arrow', David Lewis addresses our very question of what makes for closeness of worlds when assessing counterfactuals. On his theory of counterfactuals, we evaluate 'if *A* had happened, then *C* would have happened' by considering what happens in close worlds where *A* does happen. And he suggests various factors that we rank worlds by for closeness, and stipulates that it of greatest important to match worlds by their sharing laws and overall 'of little or no importance to secure approximate similarity of particular fact, even in matters that concern us greatly'.

To put flesh on the bare bones, suppose in  $W_a$  Bush doesn't press the button, but if he had

pressed the button there would have been a nuclear holocaust. In world  $W_b$ , he presses the button, the laws still hold, and so the electricity flows through the wires and the rockets fire etc: this is a world where there is a holocaust, and so  $W_b$  rapidly becomes very unlike our world  $W_a$  in matters of particular fact. In world  $W_c$ , he again presses the button and there is a small local miracle that stops the current so that things continue just as if he hadn't pressed it (perhaps Bush's memory needs to be wiped too, etc.). This is very like the actual world in matters of particular fact but the same laws no longer strictly hold. If we evaluate counterfactuals by looking at close worlds, then to get the intuitively correct value for 'if Bush had pressed the button there would have been a nuclear holocaust', we need to say  $W_b$  is closer than  $W_c$ , and so we need to treat matching on laws as more important than matching on particular facts.

*This* closeness ranking delivers the intuitively correct results, for sure, because we've cooked it up to make things come out right. But the question I'm raising is: on a Ramsey-Lewis view of laws, *why* should matching on laws be thought of as a more important similarity relation between worlds than matching on particular facts. To repeat, the story is that laws are just nice generalizations, and generalizations are just made true but the patterning of particular facts. For the Humean, the truths about the world – including the truths about the laws – are all made true by the piling up of particular facts. So you might have supposed that for the Humean the *natural* ordering on worlds is by their similarity at the level of the truth-making facts in the world, i.e. at the level of particular facts.

So the situation is this. Lewis gets things about right for counterfactuals: but this isn't due to the Ramsey-Lewis theory of laws and the general Humean metaphysics that motivates it. Rather, it is because of a seemingly quite independent add-on thesis about what makes for closeness of worlds which is expressly designed to tie laws and counterfactuals together. But that isn't an *explanation* of why counterfactuals hold true in virtue of the laws: rather it is just assumed that counterfactuals and laws need to hang together and this fact is used in *stipulating* the needed shape of the closeness relation of worlds.

3. *More on explanation* Having raised an issue for the Humean about explanation, we can press another, simpler, worry. Suppose I ask why it is true that this  $F$  is a  $G$ , and you respond because it is a fundamental law that all  $F$ s are  $G$ s. That seems some sort of explanation of the particular correlation – it isn't a fluke, or a 50-50 shot, it is governed by a law. However, the argument goes, on a Humean view, there is no more to being a law than being a nice brute generalization – and a brute generalization that all  $F$ s are  $G$ s owes its truth to this  $F$  being a  $G$ , that  $F$  being a  $G$ , that other  $F$  being a  $G$ , and so on. If anything, on the brute generalization view, the explanation seems to go the other way: the generalization holds for no other reason than because its instances hold. So, the argument concludes, the Humean Ramsey/Lewis view must deprive laws of their explanatory ooph.

These objections to the Ramsey/Lewis view aren't supposed to be novel – relations of them are variously pressed by anti-Humeans like Armstrong. But I do think that they are pretty telling against the Ramsey/Lewis view. Is that the end of Humeanism? *Not necessarily. Maybe the Humean has gone off-track in going down the 'objective' route of endorsing (2)?*

## 7 Getting the Humean out of trouble?

Let's try to dig down a bit and see what is doing the damage. Look again at our original characterization of Lewis's kind of Humeanism:

1. A law is a true universal generalization with some additional feature of ‘loveliness’
2. Whether a universal generalization has some this additional feature of ‘loveliness’
  - (a) is an objective matter, independent of the thoughts and activities of scientists;
  - (b) obtains in virtue of the some non-nomic, metaphysically light-weight facts about this and other generalizations.

What our arguments show is that there seems to be some sort of three-way tension between some intuitive thoughts about laws, and (2a), and (2b). Suppose we think that (2a) it is indeed some kind of objective fact that a given generalization makes the grade to be a law, some extra fact over and above the fact that the generalization is plain true. And suppose we want to adhere to intuitions like the following:

- i. An uninstantiated generalization can still be a law.
- ii. Even if a true generalization belongs to a best system, it can still be an accident, i.e. not a law.
- iii. Laws have a special explanatory role.

Then – reformulating those intuitions in terms of the thought that lawhood is an extra fact, over and above the mere truth of a generalization – we’ll be looking for extra facts of the right kind. And what our anti-Humean arguments are intended to show is that the materials allowed by (2b) aren’t enough to give us those extra facts. So – adhering to intuitions plus the assumption that lawhood is a matter of fact – we are invited by Armstrong and others to drop (2b) and cook up the needed extra law-making facts from some more metaphysically heavyweight ingredients.

But suppose that, instead, we drop the assumption (2a) – or at least give it a more deflationary reading. Suppose, in other words, we go back to an older style of Humeanism (maybe Hume’s own) which denies that there is some extra *fact* that makes a law a law. Maybe by (so to speak) getting off the bus earlier, the *radical* Humean won’t end up in a bad part of town.

Now, if we don’t say that laws are generalizations with some additional factual something about them (like holding their own in a nice systematization of generalizations), then what should we say about laws? Here’s Goodman:

I want to emphasize the Humean idea that rather than a sentence being used for prediction because it is a law, it is called a law because it is used for prediction.

So a law is a generalization that is treated in a certain way by scientists. Or here is Ramsey making the same claim in a late paper (in which he backtracks from the best-system Humean story which Lewis develops),

Laws form the system with which the speaker meets the future.

Here writing a little later is Ayer,

I propose to explain the distinction between generalizations of law and generalizations of fact . . . by the . . . method of analysing the distinction between treating a generalization as a statement of law and treating as a statement of fact.

Which hints at the idea (to put it suggestively) that a law is a generalization to which scientists have a certain attitude.

Of course, this older, pre-Lewis, style of Humeanism is well-known to Armstrong and other anti-Humeans, who lambast the position as involving a hopeless kind of subjectivism. If we say that being a law is a matter of being a true generalization which is *treated* as a law, then that seems to rule out the very coherence of the idea of there being unknown laws (which plainly, if unknown, aren't *treated* as laws). But it is only becoming scientific modesty to suppose that, like our predecessors, there are scientific laws we don't know. And – what Armstrong takes to be the most fundamental difficulty – if our attitudes are not just to be arbitrary, surely we ought to treat one generalization as a law (whatever that involves) and another as a non-law for some *reason*, on the basis of some discriminating fact that differentiates the two. But then, surely this discriminating *fact* must be what matters for lawhood: and hence our *attitude* to the fact should drop out of the story after all.

It might be tempting to rush to assume that a radical Humeanism that rejects (2a) is hopeless. But let's go more slowly.

First, consider a parallel, comparing radical Humeanism about *laws* with Humeanism about *morals*. So, corresponding to the idea that the scientific laws are laws because of our attitudes to them is the idea that the moral laws are laws because of our attitudes to them. And corresponding to the idea that something is a scientific law because we use it for prediction (and not vice versa) is the idea that something is a moral law because we use it to guide our life.

Now, the troubles with crude expressionism about morality are well-known to every beginning ethics student. For a start – to echo the objection to the roughly analogous view about scientific laws – if we say that being a moral law is a matter of being *treated* as a law, then that seems to rule out the very coherence of the idea of there being unknown moral laws. But it surely it is only becoming modesty to suppose that, like our predecessors, there are moral things we are getting badly wrong (as they got things wrong about slavery, say). However, while the troubles with *crude* expressionism are well known, there are sophisticated kinds of Humean expressionism about morality which are more resistant to criticism.

So that suggests a question: can we find a sophisticated version of the radical Humean thought that being a law is a matter of being *treated* as a law which stands to crude Humeanism as a sophisticated version of ethical expressionism stands to the crude version?

We find a sketch of a positive answer in one of Simon Blackburn's papers, 'Morals and Modals', collected in his *Essays on Quasi-Realism*. The position, it is fair to say, isn't developed particularly far, but we do get some very intriguing pointers. In the rest of these notes, let me say a little about this line and how it might be pushed forward.

## 8 Another summary: towards subjective Humeanism

We noted that, on Lewis's version of Humeanism, a law is a true universal generalization  $G$  with some additional feature of 'loveliness', where 'loveliness' is an *objective* matter, i.e. something independent of the thoughts and activities of scientists, and which obtains in virtue of some (metaphysically light-weight) *facts* about  $G$  and other generalizations. We argued that the trouble with this view is that there don't seem to be enough facts to go around on the Humean story: that is to say, the Humean doesn't have enough facts at his disposal to make true various claims about laws that *seem* to be true. The anti-Humean will conclude that there are more facts than the Humean dreams of, e.g. perhaps facts about relations of necessitation among universals.

Of course, the defender of Lewis's position can fight back and claim that the intuitions that seemingly conflict with the Humean position should be rejected. But then, if the Humean view is presented as an *analysis* of our intuitive notion of law, we at least need an explanation of why so

many people find the counter-Humean intuitions rather compelling – not just in an off-the-cuff, first guess, sort of way but in reflectively articulating the constraints that they think a philosophical account should meet. And it isn't clear what the Lewisian can say here. If, alternatively, Lewis's position is presented not as an analysis but as a sanitized *replacement* for our pre-philosophical notion of law, then we need an argument that Lewis's position really is the last, best hope for the Humean.

Now I've argued that something like the Lewis view of 'loveliness' – of what differentiates laws from mere accidental generalizations – is indeed the only option *for the Humean who assumes that 'loveliness' is an out-there, worldly, feature of laws*. But what about the more radical Humean (we might plausibly say, the more Hume-like Humean) who denies that 'loveliness' is an 'objective' matter, and who takes the lead from Ramsey. A lawlike proposition, Ramsey says in 'General Propositions and Causality', is distinguished from a mere generalization 'in that we trust it to guide us in a new instance'. Attributions of lawhood, then, express this attitude of trust (and Ramsey notes, 'Many sentences express cognitive attitudes without being propositions' – i.e. without expressing facts).

How is this expressivism about lawhood to be developed (if it isn't going to be, not just radically Humean, but also radically implausible)? The obvious way forward is to aim to echo the best kinds of expressivism elsewhere, e.g. expressivism about ethics in its 'projectivist'/'quasi-realist guise'. So let's reflect (necessarily briefly) on this.

## 9 The projectivist move: a first shot

The sophisticated expressivist or 'projectivist' picture about morals has the following structure:

1. It starts with an account of the attitudes, the emotional responses that various actions and situations can produce in us, and the dispositions to action these responses might engender.
2. It treats our moral judgements as giving voice to these attitudes, emotional responses, dispositions, etc. . . .
3. Then – the difficult bit – it needs to explain why our judgements (although expressive) have the *look* of straight-forward propositions, and why our ordinary talk of such propositions as true or false, as being known or unknown etc., is in order, and why they embed in more complex contexts just as propositions do.

The key Humean thought here, as Blackburn puts it, is that what starts life as a non-descriptive psychological state ends up expressed, thought about, and considered 'in propositional form'.

How can this be right? Take a claim like 'In my view, fox hunting is wrong, but I could be mistaken about that' or 'Slavery would still be wrong, even if we all thought it wasn't.' These claims – which surely make perfectly good sense – seem at first blush to be troublesome for the expressivist. Their very sense seems to require a gap in principle between values and feelings. This might encourage a realist view about values, which certainly would give us a gap, though realism leaves us with some nasty puzzles: what kind of properties can ethical values be, somehow out there independently of us yet commanding certain reactions from us, and what kind of epistemic access can we have to them? The projectivist by contrast aims to make sense of these troublesome claims without abandoning the basic Humean picture that morals are ultimately linked to sentiment rather than a mysterious kind of non-natural property. On this view, 'In my view, fox hunting is wrong, but I could be mistaken about that' just expresses a *moral* stance, an openness to the possibility of improving and refining one's understandings both of the practice and what it involves and of the

human sentiments that surround it. The claim is – to quote Blackburn – ‘part of good ethical thought rather than a metaphysical insight’. Similarly ‘Slavery would still be wrong, even if we all thought it wasn’t’ just expresses the possibility of a regrettable coarsening of our moral sensibilities.

To push the point a bit further, compare ‘Apple pie would still be yummy even if we all thought it wasn’t’. That’s pretty difficult to construe: there’s no more to being yummy than typically eliciting the reaction *yum!*. Thinking of something as yummy may indeed just be to project onto the thing our reaction *yum!*. However – the argument goes – we don’t have the same kind of difficulty in construing ‘Slavery would still be wrong, even if we all had still been thinking it wasn’t’. Does that contrast show that thinking of something as wrong commits us to reaction-independent objective properties out there? The Humean will say, however, that the difference in status between ‘yummy’ and ‘wrong’ lies not (so to speak) at the worldly end but at our end: in other words, it isn’t a difference in the kind of properties ‘out there’, but a difference in the complexity and structuring of our attitudes, and not least in the importance we place on sharing moral attitudes. And while there is little nuance or openness to reflective control to the *yum!* reaction, our moral take on the world involves a whole web of interrelated attitudes, and attitudes to attitudes, with room for examining relationships between these attitudes, bringing one to bear on another, and changing them accordingly. And – the basic projectivist story goes – because our web of attitudes is complex enough to make room for the thought that we might come to regret one of our attitudes, we allow for the possibility of ‘error’. That is why, when we ‘spread’ our attitudes onto the world, we do so in such a way that it takes on (much of) the logical form of straight fact-stating discourse.

Note, the expressivist who gives an explanation of how it can be that our moral discourse *looks* realist (but isn’t), agrees that there are moral values and that murder is wrong – indeed – being a good deflationist about the truth-predicate – he will agree that it is true that murder is wrong. He can even say that it is an ethical ‘fact’ that murder is wrong so long as the notion of ‘fact’ is here just being used as in ordinary, street-level claims, without any special philosophical loading. Like Berkeley, the expressivist can speak with the vulgar: though of course he denies that it is a Fact (with a capital ‘F’) that murder is wrong, if capital ‘F’ facts are existing states of affairs built out of Humean-approved entities and their properties.

So much for morals, now for laws. The sophisticated expressivist (or ‘projectivist’) picture about laws will have a similar structure:

1. It starts with an account of the attitudes and the dispositions to action that scientists may have vis-à-vis various generalizations – they use them in predictions, deploy them in explanations, rely on them in planning their practical engagements with the world, etc.
2. It treats our judgement that some true generalization is a law as giving voice to these attitudes (which explains for a start why laws are particularly related to predictions and explanations).
3. Then – the difficult bit – it needs to explain why our judgements of lawhood (although expressive) have the form of propositions.

As for the ethical case, we can ask: how can this be right? Take a claim like ‘In my view, it is a law that  $\alpha/\beta$  interactions produce two  $\gamma$  particles, but I could be wrong about that (maybe it is just an accident)’. That claim – the sort of thing which our earlier anti-Ramsey/Lewis arguments took to make perfectly good sense – seems at first blush to troublesome for the expressivist: it looks like a claim about the independence of lawhood from our attitude that the generalization is a law. It looks like a claim of a kind of metaphysical independence of a further fact about lawhood (a further fact over and above the truth of a generalization). Or what about ‘Although there haven’t

been any  $\alpha/\beta$  interactions [so we don't have any special attitudes to any particular generalizations about such interactions] there could well be a basic law about what would have happened had there been an interaction'. That was the modal intuition underlying our Tooley/Carroll argument in Section 3, and again it seems to imply the independence of facts about laws from facts about attitudes (since we don't have attitudes to unknown, unexemplified generalizations).

The sophisticated expressivist (or 'projectivist'), however, aims to make sense of the troublesome claims without abandoning the basic Humean picture that lawhood is ultimately linked to how we treat a proposition rather than a mysterious kind of extra property. On this approach, 'In my view, it is a law that  $\alpha/\beta$  interactions produce two  $\gamma$  particles, but I could be wrong about that (maybe it is just an accident)' just expresses a becomingly modest scientific attitude, a recognition (perhaps) that one has had to take a stance when evidence is weak or an openness to reconsider one's attitudes again if further relevant evidence can be brought to bear. To echo Blackburn, this would be 'part of good scientific thought rather than a metaphysical insight'. Again, 'Although there haven't been any  $\alpha/\beta$  interactions there could well be a basic law about what would have happened had there been an interaction' reveals — as a first shot — a suitably modesty in the face of ignorance, combined with a recognition that generalizations about  $\alpha/\beta$  interactions can be the sort of thing that it is appropriate to treat in the law-like way if and when we can get hold of them. Again, part of good scientific thought rather than a metaphysical insight.

It is a tough question how far this radical Humean line can be pushed. That's my next topic. But the summary initial claim that I am making is this. On a Lewis-style 'objective' Humean picture, being a law is a matter of some fact's obtaining (a capital 'F' fact, if you like), over and over what makes a true generalization true. But the Humean story of the capital 'F' facts that go to make up the world don't seem to supply the right kind of facts to do lawmaking thus conceived (or at least not if we want to adhere to various basic intuitions about laws). So, the best bet for the Humean is to be more radical — arguably, to be more Humean — and to deny that being a law is a matter of some capital 'F' fact obtaining over and above those that make generalizations true. And the suggestion is that the radical Humean can still adhere to the troublesome basic intuitions about laws, but treated as expressions of good scientific practice, not as metaphysical insights. So maybe the Humean in radical guise lives to fight another day.

## 10 Laws and provisos

To develop the radical Humean story, we need three ingredients —

1. a story about what, exactly, we are supposed to be taking attitudes towards in making law claims (the objects of our attitudes);
2. a story about what exactly the attitudes are that we are expressing in making law claims (the nature of the attitudes);
3. a story about why an expression of these attitudes takes on the look of something propositional.

Let's start with the first.

You might think that we needn't dwell on this. Isn't the answer just obvious? What we regard as lawful are various universal generalizations, plain and simple.

But in fact, it isn't so plain, and it isn't so simple. For a start, not all propositions that are termed 'laws' by scientists are regarded as *universal* generalizations. Take Snell's Law or refraction (I'm here borrowing an example from Marc Lange's book *Natural Laws in Scientific Practice*) —

this is the law that when a beam of light passes from one medium to another, the angle of incidence  $i$  and angle of refraction  $r$  are related by  $\sin i / \sin r = \mu$ , for a constant  $\mu$ . However, this in fact only holds provided that temperature and pressure fall into ‘normal’ ranges (whatever that means), when there is no magnetic or electrical potential difference across the boundary between the media, when the media lack double refractivity, and so on. Snell’s Law has to be understood subject to *provisos*. Moreover, it isn’t the case that the provisos are, so to speak, fixed in advance, and that we can determinately list them all. To put it another way, Snell’s law holds *other things being equal*; but what makes for other things not being equal is pretty open-ended.

Another example: Galileo’s law of falling bodies: falling bodies accelerate to Earth with acceleration  $9.8\text{m/s}^2$ . Provided that we are not dealing with feathers where air resistance matters. And providing of course there aren’t collisions, parachutes, updrafts, magnetic fields, etc. etc.

Now, there are laws that – by contrast – do hold (we believe) without provisos: for example, nothing can be accelerated to a superluminal velocity; or for any electron and any proton, their charges are equal and opposite. And one stern line of thought might be that it is these so to speak *iron laws* – absolutely exceptionless laws – which are the genuine items. And the likes of Snell’s Law and Galileo’s Law – which have to be hedged around with provisos – are second-rate affairs, not kosher laws at all. (An aside: to talk of exceptionless, iron, laws is *not* the same thing as talking about deterministic laws. There could be exceptionless but chancy laws – for example that any elementary particle of a certain type has a 50% chance of decaying in the next second.)

The trouble with taking the stern line and downgrading provisoed laws is the obvious one that science will be left with precious few philosopher-approved top-grade laws – at best, the fundamental laws of physics. And we don’t know many of those (our best going theory of the very small – quantum mechanics – and our best going theory of the very large – general relativity – can’t be consistently combined, so we *know* they can’t both be the last word as they stand). So we are in danger of ending up with no known top-grade laws.

So let’s consider being less stern. Snell’s Law is entirely typical of run-of-the-mill laws: so let’s learn to live with provisos. But there is a puzzle here. Snell’s Law,  $\sin i / \sin r$  is constant, other things being equal, is not vacuous. But the statement ‘ $\sin i / \sin r$  is constant except when it isn’t’ *is* vacuous. So how do we understand the ‘other things being equal’ proviso so the law isn’t vacuous? More generally, how should we best think of scientific theories to accommodate the fact that most laws have to be hedged about with provisos. How do hedged generalizations work in theories? And how do they get applied to the world?

## 11 Laws and models

Let me tell you a story . . .

I remember, as a smallish boy visiting the Science Museum, being fascinated by an orrery – a mechanical model of the solar system. It was a beautifully constructed model, with little brass globes representing the sun and the planets and the moon, and these globes were on spindles attached to a complex geared mechanism. Wind a handle at the side, and the globes made their way around their orbits. Mercury got up a fairly brisk pace around the sun; while the outermost planet (Uranus, I guess) made hardly discernible progress, however furiously you wound.

Some features of the orrery were intended to be representational. The relative speeds of the orbiting globes as you wound the handle was evidently supposed to approximate the relative speeds of the planets in their orbits; the relative sizes of the brass globes indicated the approximate relative sizes of the planets. Other features were not representational. For example, the ratio of a globe’s diameter to the diameter of its orbit was not to be taken seriously; the planets were not were not

depicted on the same scale as their orbits. (Maybe there was a label on the exhibit, explaining this: but much, no doubt, was taken for granted.)

Later, I got to use what I suggest we can think of as ‘abstract orreries’, mathematical models rather than mechanical ones. Newtonian celestial mechanics teaches us how to construct a simple first orrery, by solving a bunch of problems for single bodies orbiting a fixed gravitational source. That’s pretty easy. Then the fun really starts when we try to improve things, using (say) classical perturbation theory, to take into account the action of Jupiter on the other planets and so on, to give us a series of ever better abstract orreries. As with the mechanical model, a number of features of these abstract orreries are supposed to be representational, others not. The abstract orreries deal in ‘point masses’; that obviously isn’t supposed to mean that planets are point-sized. On the other hand, the orbits of the planetary points in the abstract orreries *are* supposed to represent, or approximately represent, the course of the planets in the heavens.

And how do they do that? Well, a neat way of looking at things is this. Think of the real-world planetary motions being plotted on a space-time diagram; and then run the Newtonian orrery in the same abstract four-dimensional space. Start off our abstract orrery from the plotted positions the planets have at some time  $t_0$ . Then we (so to speak) turn the handle – run the orrery for a suitable time-lapse – and the orbits in the orrery close-track the plotted actual orbits. The longer they close-track, the better the orrery.

Now, this talk of abstract orreries merely adds a touch of colour to what is surely the natural account of what is happening when we do classical celestial mechanics. We build some mathematical models. We present these as being similar in their behaviour to the worldly phenomena being modelled. And similarity is (at least as a first approximation) a matter of tracking the time-evolutions of the relevant quantities.

What goes for celestial mechanics goes for many other areas of mathematical physics too. It’s (so to speak) abstract orreries all the way. In the case of classical optics, for example, we have models of light rays reflecting and refracting. We can think of these models as abstract geometric structures. And the model is a good one if the geometric structures in the model are suitably similar to a plot of the trajectories of actual light rays.

According to a still-prevalent line, a theory is a set of sentences, and virtue in a theory is a matter of the truth of the sentences in the set. Against this, however, I’d canvass a different proposal: certain kinds of theory, at any rate, are better regarded as involving instructions for building mathematical models, models which are then claimed to track the world in certain respects. And virtue in a theory is (at least in part) accuracy in the models. That suggestion, of course, isn’t radically new. The nearest version to the one I favour is to be found in the work of Ronald Giere, though when it comes to details I would still spin the idea differently even from him: see my book *Explaining Chaos*. But our purposes here, those differences perhaps don’t matter, and I won’t dwell on them.

On this ‘modelling’ view of science the role of laws is to give us instructions for model-building. For example, Newton’s Laws of motion tell us how to build classical dynamical models of the world. The laws of classical optics tell us how to build models of optical phenomena. And the role of any provisos that attach to laws is to indicate limits of application of the models they generate. Take Galileo’s Law again. In building some simple dynamical models, we use the ‘clean’, unqualified, no-fuss, generalization  $a = 9.8\text{m/s}^2$ . But then the question arises about the range of application of such models. We have some paradigm examples of where the models work well – for a start, when dropping objects in the sorts of conditions which we encounter in the school laboratory. We have some conception of a range of cases where the law doesn’t apply. And the thought is that the law applies to cases sufficiently like the paradigms and sufficiently unlike the cases where there are disturbing factors. And what makes for ‘likeness’ here? Undoubtedly it’s a fuzzy matter, and

judgements may always be subject to revision. But that's the case with lots of our judgements: we can't immediately conclude that there is no content to the idea.

So the suggestion is that a prime task of science is model-building; many generalizations we regard as laws are propositions we use in model construction; provisos with which we hedge around laws give indications of the range of application of the resulting models. So here's a Humean proposal: *to say that a generalization is a law is (as a first shot) to recommend it for model-building*. To accept a generalization as a law is to be prepared to deploy it for model building – so as Ramsey put it, 'Laws form the system with which the speaker meets the future'.

## 12 An aside on Giere

Here's Ronald Giere talking about Newtonian mechanics:

What one learns about the world is not general truths about the relationship between mass, force, and acceleration, but that the motions of a vast variety of real-world systems can be successfully represented by models constructed according to Newton's principles of motion. And here 'successful representation' does not imply an exact fit, but at most a fit within the limits of what can be detected using existing experimental techniques. The fact that so many different kinds of physical systems can be so represented is enough to justify the high regard these principles have enjoyed for three hundred years. Interpreting them as universal laws laid down by God or Nature is not at all required.

This is from his article entitled 'Science without Laws of Nature'. So officially he ends up in a very different position from the expressivist Humean: Giere officially rejects the idea of laws (though he talks about *principles* of motion) while the Humean aims to explicate law-talk, though in a way free from metaphysical excess. But I suspect that the real disagreement is small. For Giere and my Humean both reject the idea of what I called capital 'L'-laws, where capital-'L' lawhood has to involve there being some fact of the matter differentiating laws from mere generalizations. And both Giere and my Humean say that what is going on when we deploy what scientist's call laws (with, so to speak, a small 'l') is that we are deploying principles we use in model-building. The difference between them is that Giere seems to think that law-talk is too infected with non-scientific, metaphysical, pictures of lawhood, and needs eliminating. My Humean is more conciliatory, and thinks that law-talk, properly understood, is in good order as it is – we just need to give a deflationary account of it, and recognize that its use is to express endorsement of model-building policies rather than as reporting some kind of special factual status.

## 13 Another summary: the tasks ahead

We said that developing a radical Humean story about laws will involve three components

1. a story about what, exactly, we are supposed to be taking attitudes towards in making law claims (the *objects* of our attitudes);
2. a story about what exactly the attitudes are that we are expressing in making law claims (the nature of the *attitudes* themselves)
3. a story about why an expression of these attitudes takes on the *look* of something propositional.

We've now outlined a story of the first type, and hinted at the corresponding story of the second type:

1. Following a not-unfamiliar 'modelling model' of science, I suggested that certain kinds of theory should be regarded as involving instructions for building mathematical models, models which are then claimed to track the world in certain respects. The relevant generalizations which we regard as laws – the objects of our attitudes – are propositions we use in model construction; the provisos with which we hedge laws then give indications of the range of application of the resulting models.
2. So here's a corresponding proposal about the attitude being expressed by saying some generalization is a law: as a first shot, law-claims recommend a proposition for model-building purposes.

Our tasks now are to develop these first two stories, and to give some idea of how to begin the needed third story.

## 14 More on using models in explanations

Let me start by saying a bit more about how we construct and use mathematical models in science (actually, I think what I am going to say applies to scientific models in a broader class than just 'mathematical models' strictly so called: but to keep things under control, we'll concentrate on the narrower class).

It is surely all-important that construction principles for models are not (as it were) 'disposable', i.e. we don't use-once-and-then-throw-away. The whole project of science is to provide us with principles for model-construction that can be used and re-used to provide different particular models for use on different particular occasions. For example, Newton's laws can be used and reused in modelling the behaviour of this particular pendulum, the trajectory of that particular canon-ball, and the motion of those planets. And take a particular pendulum: then we can apply the same principles to model its behaviour in various circumstances, as we let it swing at various angles, or as we alter the weight of the bob, or the length of the string, or the downward force on the bob (e.g. by adjusting the magnetic attraction on an iron bob).

Now, our general principles for building models of the behaviour of a variety of pendulums produce a range of particular models that reveal that changing the angle of swing doesn't change the period or the mass of the bob, while changing the length of the string or the force acting on the bob does. In other words, keeping the downward forces equal, a pendulum  $A_1$  having period  $\omega_1$  while pendulum  $A_2$  has period  $\omega_2$  goes with  $A_1$  having length  $l_1$  while  $A_2$  has length  $l_2$  – but change of period doesn't co-vary with change of weight of the bob. Or to put it another way, our principles for building pendulum models show that a given pendulum has period  $\omega_1$  rather than  $\omega_2$  goes with its having length  $l_1$  rather than  $l_2$ . But that kind of *contrastative* statement is – on one plausible view – of the essence of *explanations*. Which suggests the following thought: *principles for model-building are principles we can use in building models that generate (contrastative) explanations*. But we've already suggested that proposing a generalization as a law is recommending it as a principle for model-building: so if model-building ipso facto goes with generating what look like contrastative explanations, then we can see that proposing a generalization as a law goes with taking the generalization as apt for various explanatory purposes.

Let's press that point rather further. We will continue to take it that explanations are paradigmatically contrastative (so the characteristic form of a fully articulated explanation statement is ' $P$  rather than  $Q$  because  $A$  rather than  $B$ '). Therefore, to generate explanations, we need to be able

to generate not only conditionals of the form ‘if  $A$  is true, then  $P$  is true’ with antecedents true-in-the-actual-world, but also possible-world conditionals of the form ‘had  $B$  been true, then  $Q$  would have been true’. By contrast, to generate *predictions* about what *will* happen, we strictly speaking only need to be able to generate the conditionals with true-in-the-actual-world antecedents: for to be able to predict what will happen, given that initial conditions  $A$  are actually exemplified, we will only need to invoke the conditional ‘if  $A$  is true, then  $P$  is true’ with true antecedents. And in principle, knowing the predictive conditionals doesn’t require knowing the explanation-relevant counterfactuals.

Now of course, *in practice*, it is just the *same* model-building devices which we use to generate actual-world and possible-world conditionals. Indeed, since we usually don’t know the future, we usually don’t know which conditionals will turn out to be which. I’ll pick up this key point in just a moment. But (to press the point of principle about the prediction/explanation contrast) imagine we are given – by a benevolent demon, if you like – the deductive system of actually-true generalizations which ‘achieves a best combination of simplicity and strength’. Then we have access to an ideally powerful tool for predictive purposes. However it seems that we needn’t thereby have been given everything we would want for contrastative explanations. For we needn’t have sufficient generalizations to use in building models which tell us, counterfactually, what *would* have happened in various non-actual circumstances. To use an example I’ve used before, it could be that one of the generalizations that ‘achieves a best combination of simplicity and strength’ is that all  $\alpha$ - $\beta$  interactions produce a pair of  $\gamma$  particles. But – if there are only a few such interactions – then (our intuition was) the truth of this generalization could still just have been a cosmic coincidence. And if that is just a coincidence, it wouldn’t provide an explanation of the production of a pair of  $\gamma$  particles. To put it the other way about, if the occurrence of the  $\alpha$ - $\beta$  interaction is to be the explanation of the production of a pair of  $\gamma$  particles then we need to be in a position to assert the contrastative fact that, had the interaction *not* taken place, then we *wouldn’t* have got those  $\gamma$  particles. But our benevolent demon needn’t, it seems, give us enough for that. So – assuming that the laws are the explanation-generating generalizations – our benevolent demon (although he gives us, by hypothesis, all the ‘Lewis-laws’) doesn’t give us all the laws.

How then does our projectivist Humean regard this apparent gap between the genuine explanatory laws and the predictive ‘Lewis laws’? His story could be something like this:

In reality, of course, there is no benevolent demon. We have only partial knowledge of the world, and in particular very partial knowledge of the future. And the world is a messy and complex place. To get any predictive purchase in a reasonably economical way, we have to use idealized mathematical models of portions of the world (often isolated systems, often ‘specially prepared’ to be insulated from much of the causal hurly-burly). These models are generated by portable principles we can use and re-use – models which (as we’ve seen) will therefore generate actual-world and possible-world conditionals together and hence yield the kind of contrastive propositions characteristic of explanations. So, for us in our non-demonic epistemic predicament, we can’t separate getting a predictive grip on the world from a modelling approach which also generates explanations. Hence, for us, the search for control over the world is tied up with the search for explanation-yielding models. Hence, no doubt, our deeply entrenched (very likely, innately wired in) habit of looking for explanations. And there’s more: it seems (and you can understand how there could be Darwinian reasons for this) that we are habitually optimistic about the availability of explanations: we seem to habitually think that, with due endeavour and ingenuity, models should be forthcoming. Or as we’d put it – given our projective ways – we suppose there should always *be* explanations, even

if we don't yet know them. This is a habit of mind which we can't shake off – even when we try to confront in imagination the extreme case of being possession of demonic knowledge of all the 'Lewis-laws': even then, we still can't help thinking there are more explanations to be had (even though there are no more [Humean-approved] facts for these explanations to call on). But applied at the extreme, our habit of thought is inappropriate (though far too engrained for us to be comfortable with the thought that we should drop it).

The line of thought here is an intriguing inversion of the anti-Humean's line of argument. The latter argues that Lewis's predictively-ideal best system of true generalizations doesn't truly fix the explanatory laws: and the moral drawn is that our explanatory, law-seeking, task is to dig deeper than mere prediction (even ideal prediction), and to reveal something of the organization of the world beyond the predictively true generalizations. Our Humean by contrast thinks of model-building (along with contrastative explanation) as something we go in for because of our epistemic *limitations*: model-building doesn't exemplify a cognitive task deeper than prediction, but rather is how we cope with our limited successes at making predictions.

## 15 Our attitudes to laws

The suggestion is – roughly put – that to endorse some principle as a law is to ratify it as appropriate for use in model-building.

If the models fail badly, then we will of course withdraw our ratification. On the success side, we can have models that seem to work perfectly as far as we can tell (rare indeed!); models that work perfectly well in certain cases, but fail in other applications within their originally intended domains of application (maybe models incorporating Snell's law); models that never work perfectly but across their intended domain of application work to a high level of approximation (Galileo's principle about unimpeded falling bodies). I suspect that in some contexts of discussion, depending on the contrast class we have in mind, we will call Galileo's principle (for example) a law, in other contexts we will deny that it is a law ('strictly speaking', as we say). For the Humean this isn't wavering about the factual status of Galileo's principle (is it a genuine relation between universals?), but – plausibly enough – it is a context-shifted change in the strength and scope of the ratification.

But – an objector might protest – why should we say that endorsing a generalization as a law is ratifying it (expressing a recommendation) rather than simply asserting: this generalization *works* (in the model-building enterprise)? After all – the objector will continue – we'll only want to ratify a law if it works: so there will this much objective basis for the reaction we are giving voice to (according to the Humean). So why not say that in asserting that something is a law, we are asserting that this objective basis holds, rather than expressing an attitude?

The projectivist Humean, I take it, will want to hear more about this 'objective basis'. To say that some proposition is a law is, presumably, not to say that the *actual* models that we have so far built with it are very successful (we might be struggling with extracting usable predictions from an abstract law, as in the early days of General Relativity, or we having troubles with auxiliaries, etc. etc.: and, on the other side, we might have models that, so to speak, work by accident and whose principles fail of application when pushed further). The objector's idea, presumably, is more that the models using the law, suitably developed by ideally rational enquiries, would be successful. But now the Humean will worry about the factual content of *that* idea, and give a projectivist account of the content of that counterfactual supposition. Asserting that counterfactual just is expressing our trust in the law.

Is there a parallel here between the debate between a Humean projectivist about morals and an ‘ideal observer’ theory? On the ideal observer theory, what is right is what an ideally disinterested observer of impartial sympathies, etc. etc. would judge to be right: the Humean needn’t so much object to that idea as to the suggestion that this is where a theory of our moral sentiments can bottom out – for what sort of fact is it that an ideal observer would judge thus and so? No capital-‘F’ fact, out there in the world, the Humean will say: rather it’s a projection of the moral reactions we find we approve of most in ourselves (at our most unself-centered and sympathetic).

## 16 Projectivist semantics?

There’s a remaining issue: we asked what story we might give about why an expression of these attitudes takes on the *look* of something propositional. I wish I had something useful to say here, but I’m not confident of the shape of the debate here.

Sometimes it seems as if the projectivist with respect to a given class of assertion aims (or takes himself as having to aim) to (1) give an account of the nature of the attitudes expressed by ‘atomic’ assertions of that kind, and then (2) explain more about the complexity of those attitudes, and the attitudes that we have to them, etc., and thereby (3) warrant our use of discourse with the look of assertoric, fact-stating, discourse. (Recall Blackburn’s discussion of how we might have the complex attitude (H) *Hooray to (Boo to foxhunting) ⇒ Boo to encouraging your children to hunt*), where ‘ $\Rightarrow$ ’ indicates a suitable coupling relation between attitudes: he then suggests that while the atomic ‘Foxhunting is wrong’ expresses moral disapproval of hunting, the conditional ‘If foxhunting is wrong, then encouraging your children to hunt is wrong’ might be taken to express attitude (H).)

But there’s another way of proceeding (or is it a different way?). We might – so to speak – look at the language-game of evaluation as a whole. In a Sellarsian spirit we might give ‘language entry’ and ‘language exit’ rules attaching to atomic sentences (rules like *you can assert ‘Foxhunting is wrong’ is wrong to express disapproval and if you come to endorse ‘Foxhunting is wrong’ then disapprove of foxhunting*). And then we describe the transition rules between complex sentences (the usual kind of inference rules governing connectives etc.). We don’t aim to give a deeper warrant for the language-game – in Wittgensteinian spirit: the game is played, and our moral life involves playing that game. And we don’t call on the projectivist story to warrant the shape of the game: rather, we say – playing the game that way just is ‘what is meant by “projecting” attitudes onto the world’ (which is also something Blackburn says).

Is there a tension here? On second thoughts, I think not: but in any case, I take the second line as really the more fundamental, and the first as providing an additional gloss. How does this apply to the law case? Well, that’s something to explore another time: there’s work to do – but the prospects for subjective Humeanism, I submit, are far from hopeless.