

HUMEANISM AND EXCEPTIONS IN THE FUNDAMENTAL LAWS OF PHYSICS

BILLY WHEELER

Abstract. It has been argued that the fundamental laws of physics do not face a ‘problem of provisos’ equivalent to that found in other scientific disciplines (Earman, Roberts and Smith 2002) and there is only the appearance of exceptions to physical laws if they are confused with differential equations of evolution type (Smith 2002). In this paper I argue that even if this is true, fundamental laws in physics still pose a major challenge to standard Humean approaches to lawhood, as they are not in any obvious sense about regularities in behaviour. A Humean approach to physical laws with exceptions is possible, however, if we adopt a view of laws that takes them to be the algorithms in the algorithmic compressions of empirical data. When this is supplemented with a distinction between lossy and lossless compression, we can explain exceptions in terms of compression artefacts present in the application of the lossy laws.

Keywords: Laws of nature; *ceteris paribus* laws; idealization; best system account; data compression; algorithms; Humeanism; scientific fictions.

1. Introduction

It is no recent observation that scientific laws appear to have ‘exceptions’ if taken as literal descriptions of the behaviour of objects. When Galileo introduced his famous ‘law of the pendulum’ viz., *that the time taken to perform a complete oscillation depends only on the strength for gravity and length of the pendulum*, he remarked ‘I grant that these conclusions proved in the abstract will be different when applied in the concrete and will be fallacious to this extent’ (1638, p.251). The explanation he gave is familiar: all actual pendulums will be subject to certain ‘impediments’ (such as slack in the pendulum cord and friction around the pivot-point) and because of this any calculations derived from the law will fail to match the pendulum motion exactly. Despite these discrepancies scientists do not usually think these kinds of cases count as falsification of the law. Instead, it is said that the law is a ‘*ceteris paribus*’ law, and provided the practitioner is aware that a target system will — in all probability — be subject to impediments, then the law can still be used successfully for prediction and explanation.

Whilst these kinds of exceptions have not traditionally been problematic for scientists, things have been different for philosophers of science. A long-lasting intuition about laws of nature is that they should be *universally true*. A problem then emerges concerning how to interpret laws with exceptions — for according to our pre-analytic

Principia 21(3): 317–337 (2017).

Published by NEL — Epistemology and Logic Research Group, Federal University of Santa Catarina (UFSC), Brazil.

intuitions — they should not exist. One natural suggestion is to suppose that the problem only emerges for examples like Galileo's because here the law is given in a 'short-hand' form. Instead, a full statement of the law should include provisos to exclude all impediments and when this is achieved the finished law will be true and exceptionless. Unfortunately, this proposal has given rise to a famous problem known as 'Lange's Dilemma'. Consider the example:

(A₁) Taking aspirin daily reduces the risk of heart attack.

We know that this 'law' has exceptions. If a person takes aspirin daily and consumes twice as many calories as they ought to, then their chance of suffering from a heart attack will *increase* rather than *reduce*. To avoid this we could try hedging (A₁) likewise:

(A₂) Provided the individual consumes no more calories than their recommended daily allowance *and* provided they take regular exercise *and* provided they do not have a genetic predisposition toward obesity *and*, ... so on and so forth ..., then taking aspirin daily will reduce their risk of heart attack.

The inclusion of the '... so on and so forth ...' signifies that there are many potential impediments and the worry is whether this gap could ever be filled in. Marc Lange believes that we ought to be pessimistic. The number of possible impediments is too vast, he argues, and 'the only way to utter a complete law-statement is to employ some such condition as "... in the absence of relevant factors"' (1993, p.235). In other words, an exceptionless version of (A₁) should read:

(A₃) All else being equal, taking aspirin daily reduces the risk of heart attack.

This leaves us with a doubt concerning the semantic work being carried out by the 'all else being equal' proviso. A natural worry is that if 'all else being equal' means 'unless otherwise', then (A₃) will be *logically* true and not something that could be supported or disproved by empirical evidence. The problem exceptions pose is to find a way out of Lange's Dilemma — to find a reading of (A₃) that is neither false, meaningless or a tautology.

To date philosophers have attempted to solve Lange's Dilemma by either giving a *semantic* or a *metaphysical* solution. Up until now semantic solutions have been the most common. These involve formulating appropriate truth conditions for the 'all else being equal' proviso so that it does not become synonymous with the troublesome 'unless otherwise' proviso. The accounts given by Jerry Fodor (1991), Daniel Hausman (1992), Pietroski and Rey (1995) and Alan Silverberg (1996), can all be seen as contributions to this effort, and there is still a lively ongoing debate concerning what the right truth conditions for such provisos ought to be.¹

Metaphysical solutions are different from semantic solutions in that they allow for the possibility of *genuine laws* of nature with *genuine exceptions*. This undercuts Lange's Dilemma, as it implies that these laws don't need to be hedged after all, and so we can stop worrying about how the proviso ought to be interpreted. Metaphysical solutions have been less pursued: probably because they run counter to the intuition that laws are *universal* and *exceptionless*. One of the earliest writers to provide a metaphysical solution was Nancy Cartwright (1989). According to her 'the laws of electromagnetic repulsion and attraction, like the law of gravity, and a host of other laws as well, are laws about enduring *tendencies* or *capacities*' (1989, p.1, my emphasis). Cartwright's solution works because it depends on the existence of powers — a kind of irreducible dispositional property that objects can have which may or may not manifest depending on the circumstances.

Cartwright's solution to the problem of exceptions has proven popular and variations of her account have been defended by Stephen Mumford (1998) and Andreas Hüttemann (2007; 2014). This approach to the metaphysics of laws is anti-Humean in that it allows for the existence of more than just occurrent behaviour in the world. It therefore sets a challenge to those sympathetic to Humeanism to explain laws with exceptions. Only very recently has that challenge begun to be met, mainly through the work of David Braddon-Mitchell (2001), Markus Schrenk (2007; 2014) and Jonathan Cohen and Craig Callender (2009; 2010).

In this paper I evaluate to what extent a Humean metaphysical solution to the problem of exceptions in fundamental physical laws is possible. I start in section 2 by looking at the regularity view and what many take to be its most defensible variant — the Best System Account (BSA). It turns out that at least two different kinds of exception-ridden laws exist in science and in section 3 I make a distinction between *ceteris paribus* laws and ideal laws. I consider both Schrenk's and Braddon-Mitchell's variations on the BSA and argue that Braddon-Mitchell's is the more superior. However, Braddon-Mitchell's explanation in terms of 'lossy compression' can only work when the BSA is abandoned in favour of an alternative theory that takes laws to be algorithms for compressing empirical data. This theory is outlined and explained in section 4 and finally in section 5 it is applied to the most problematic laws in physics — those that have not positive instances at all.

2. The Best System Account

The BSA has its origins in the writings of John Stuart Mill (1895) and Frank Ramsey (1990), but it is with David Lewis (1973; 1980; 1986) that we find its most canonical formulation. According to Lewis a law is a general statement that fits in an appropriate way with other statements to form an axiomatic system of empirical truths.

Deductive systems can be more or less *simple* and more or less *informative* depending on the number, length and content of the axioms. According to Lewis, science values simplicity and informativeness in equal measure so that the ‘best system’ is the one that strikes as good a balance between them as truth will allow. Lewis believes that it is reasonable to suppose that the laws of nature (as opposed to what scientists currently take the laws to be) correspond to the *axioms* in the system that systematises *all* empirical truths about our universe whilst best balancing simplicity and informativeness.²

It is clear from what Lewis says that he thinks the axioms of the best system will be universally true: ‘whatever else a law may be, it is at least an exceptionless regularity’ (1986, p.45). So it does not appear at first glance that there is any room for including exception-ridden generalisations within the BSA. Despite this, Schrenk (2007; 2014) believes Lewis’ original theory can tolerate exceptions. To achieve this he utilises Lewis’ theory of counterfactuals. The truth of a counterfactual is judged, according to Lewis, by the proximity of the counterfactual world to the actual world. Proximity of worlds can be measured by a number of different criteria, but are frequently judged on the basis of the laws they share. Schrenk appeals to a remark from Lewis which appears to support the allowance of some small violations of law:

A localized violation is not the most serious sort of difference of law. The violated deterministic law has presumably not been replaced by a contrary law. Indeed, a version of the violated law, complicated and weakened by a clause to permit one exception, may still be simple and strong enough to survive as a law. (1973, p.75)

For Lewis, this means that counterfactual statements that involve just one difference to the actual world could be judged to be true (using our laws of nature) as the difference between that world and ours is only very small (our hedged law would still be a law even at the counterfactual world). For Schrenk, this means that the axioms in the best system in the nearby world can tolerate some isolated violations. Now assuming that what is good for that world is good for our own means that we can expect exceptions to laws at our world too. Where there is a singular exception, it is plausible that a statement hedged to exclude it could still turn out to be part of the best system *vis-a-vis* overall strength and simplicity, since the additional proviso ought not to complicate our system too much.

What would such a generalisation look like? Schrenk calls these hedged generalisations ‘index laws’ (2007 p.76) and they contain a clause to exclude a particular space-time region from the law. Hence if ‘Fs are Gs’ has an exception at the space-time region (x, y, z, t) then we can add the following axiom:

(Index Law) $\forall u(Fu \ \& \ \text{not} \ - \ @ (x, y, z, t) \supset Gu)$

The first thing to note about Schrenk's index laws is that they are not *ceteris paribus* laws. The proviso 'not – $@(x, y, z, t)$ ' is an unambiguous material condition and as such does not give rise to Lange's Dilemma. The second thing to note is that index laws do not really have exceptions, as these are excluded by the space-time region clause. It is questionable, therefore, whether Schrenk's account really suffices as a metaphysical solution since if this is the full version of the law then it is exceptionless after all!

There are, however, more serious concerns about Schrenk's proposal. Schrenk acknowledges that this explanation can only work when the number of exceptions to the law is exceedingly small. If the violation is 'temporally and spatially very limited' (2007, p.78) then it seems reasonable that an index law could make it into the best system. But if the number of exceptions is large (in comparison to the overall number of positive instances) then the resulting statement will be too complicated to outweigh its lack of strength. This is an important problem for Schrenk's account, for it looks like most of the cases in science philosophers are concerned with can be expected to have *more than one* exception. Even the aspirin law from above suffices to demonstrate this. 'Taking aspirin daily reduces the risk of heart attack' will no doubt have many exceptions over the total course of the universe's history. Galileo's law of the pendulum fairs even worse. This law does not just have exceptions in *some* instances — it has exceptions in *all* instances — because no real-world physical pendulum can ever be exempt from air-resistance, slack and pivot-point friction.

Schrenk attempts to explain cases where there are many exceptions by saying that in these worlds the best system would contain few — if any — axioms (2007, p.79). In other words, our intuitions would be to judge these worlds as *lawless*. But this does not square well with science as it is actually practiced — which seems quite happy to honour a generalisation the title 'law' even though it contains many exceptions.

Schrenk would not agree in calling the above examples 'exceptions'. He makes a distinction between 'apparent exceptions' and 'real exceptions':

A law might have an *apparent exception* (or falsification) because its consequent property is, while fully instantiated, counteracted or diluted by other events. Think of a multitude of forces (gravitational forces, electromagnetic forces, etc.) all acting upon a single body, an electron, for example. The law of gravitation alone will seem to have an exception if the other forces are disregarded in a prediction of the electron's trajectory. Quite a different case would be the following: God decides to switch off, for a second or two, the gravitational force attracting the electron. This, as opposed to the first case, would be an example where we could justifiably claim that there is a *real exception* to the law. (2007, p.25, my emphasis)

For Schrenk only if a generalisation has an exception that cannot be explained in the normal course of events should we consider it a real exception and as such only

those sorts of cases are the proper target for his account. Yet the distinction he makes between apparent and real exceptions is problematic for two reasons.

Firstly, this is not a distinction everyone can help themselves to — for whether one allows the existence of counteracting forces will depend very much on one’s metaphysics. Component forces are poorly understood by Humeans and the traditional tactic has been to analyse them in terms of potential behaviour. It is doubtful, therefore, whether their existence can do the work Schrenk requires of them without also accepting the existence of something like hidden powers or capacities. As a result we cannot make this distinction without abandoning Humeanism.

Secondly, even if we could make a legitimate distinction in this way, it is not entirely clear how illuminating the emerging theory of laws would be. The sort of cases that Schrenk calls ‘real exceptions’ are more akin to miracles and not the sort of things scientists deal with on a day-to-day basis. Schrenk does suggest that his ‘real exceptions’ might occur at the centre of black holes, but even he acknowledges this is controversial due to our lack of understanding of black holes and the fact that the laws themselves seem to predict their own failure under these conditions (2007, pp.54–61). Robert Kowalenko makes a similar complaint: ‘The reader may be forgiven for asking what, under these circumstances, Schrenk’s ‘laws’ have to do with the sorts of laws scientists normally talk about’ (2011, p.451). Whilst Schrenk’s version of the BSA may be able to cope with the metaphysics of extremely speculative cases, it does little to help the Humean respond to Lange’s Dilemma which is concerned, for the most part, with the laws scientists normally ‘talk about’.

In a more recent paper (Schrenk 2014) he does attempt to address these concerns by developing a version that incorporates examples of laws found in the special sciences. He provides what he calls a ‘Better Best System Account’ (BBSA). This differs from the BSA in two ways: (i) each individual branch of science, e.g. chemistry, biology, economics etc., is allowed to have its own best system using terms for kind distinctions that may not be shared by other sciences³ (2007, pp.7–9), and (ii) individual exceptions to the laws are not referred to via indices but to space-time regions that are explicitly named (2007, pp.12–13). It is the second part of Schrenk’s BBSA which he hopes can account for exceptions to traditional examples of laws.

The solution is not too dissimilar from his (2007) account in that the clause should be eliminated with material conditions — the only difference being that this time the exceptions are ruled out by explicitly *naming* them:

$$(BBSA) \quad \forall u(Fu \ \& \ \text{not} \ - (u = a_1) \ \& \ \text{not} \ - (u = a_2) \ \& \ \dots \ \text{not} \ - (u = a_3) \supset Gu)$$

In terms of its effectiveness of accounting for real-life scientific laws with exceptions I do not see how the BBSA fares any better than Schrenk’s original BSA. Recall that the reason why we feel compelled to add a *ceteris paribus* clause is because most exceptions to laws are unknown and an exhaustive list of them lies beyond our

abilities. Schrenk is not worried by this concern and claims that ignorance of the conditions does not rule out these laws existing in the best system in principle — their ‘un-namability’ is only a problem for *epistemology* and not *metaphysics* (2014, p.21). But there does remain a metaphysical worry. Schrenk’s move from excluding to explicitly naming exceptions does nothing to account for the lawhood of generalisations that fail to have *any* positive instances, such as the law of the pendulum. For such a law to be an axiom, all instances of harmonic motion would need to be named. Since this saving strategy is possible for any false empirical generalisation we might come up with, it conflates the law/accident distinction. Schrenk’s versions of the BSA and BBSA therefore fail to deliver an acceptable Humean solution to the problem of laws with exceptions.

3. Ideal Laws and *Ceteris Paribus* Laws

It is a common assumption in the debate about exceptions to laws that they can be found in most — if not all — scientific disciplines, and that their occurrence comes about in much the same way. It is the special sciences (e.g. biology, sociology, psychology, economics, etc..) where the existence of exceptions is perhaps most obvious, and it is here where the contemporary debate about interpreting *ceteris paribus* provisos began (Fodor, 1991). It could be argued that the following is a law of psychology:

(P_1) If S desires X , then S will take steps to achieve X .

It should be obvious that (P_1) will have exceptions. S will not always take steps to achieve X , especially if there are conflicting or overriding desires to do something that precludes X .

Cartwright (1983) can be credited for bringing attention to the fact that the most basic laws of physics also have exceptions. Controversially, she argued that these exceptions are no different in kind to those found in other sciences (1989; 1999). She illustrates her point using Newton’s law of universal gravitation:

Does this law truly describe how bodies behave? Assuredly not . . . it is not true that for *any* two bodies the force between them is given by the law of gravitation. Some bodies are charged bodies, and the force between them is not Gmm'/r^2 . Rather it is some resultant of this force . . . For bodies which are both massive and charged, the law of universal gravitation and Coulomb’s law (the law that gives the force between two charges) interact to determine the final force. But neither law by itself truly describes how the bodies behave. No charged objects will behave just as the law of universal gravitation says; and any massive objects will constitute a counterexample to Coulomb’s law. (1983, p.57)

According to Cartwright, the laws of physics are not fundamentally different from generalisations such as ‘aspirins relieve headaches’ in that, taken as a description of behaviour, they are strictly-speaking false. For her, both sorts of law have exceptions and they do so for the same reason: one or more interacting capacities blocks the expected behaviour. The Humean of course cannot appeal to capacities and so seems forced to acknowledge that the fundamental laws of physics are not about regularities in behaviour and therefore not candidates for genuine lawhood.

Cartwright’s assumption that exceptions to physical laws are ultimately the *same type* as those found in other sciences seems to me a mistake. In this section I present reasons for arguing that we are dealing with at least two different kinds of exception-ridden laws in the sciences.

Let us go back to Cartwright’s example ‘aspirins relieve headaches’. As an everyday rule-of-thumb this law is useful in that on some occasions taking an aspirin really does relieve a headache. However, we also know that taken as a universal generalisation this statement is false, i.e. there are some occasions when taking an aspirin does not relieve a headache. It is also unlikely that we could ever hedge this law with enough conditions to save it from falsity. We are forced to add a general *ceteris paribus* proviso that invokes the spectre of Lange’s Dilemma. Because of this, it is also unlikely that the law is completely reliable when predicting whether or not an aspirin will relieve a headache on a given occasion. As we do not know all the circumstances that prevent it from working, we can never tell whether these have all been eliminated or not in a particular instance.

If we compare the foregoing points with the laws from physics we see that things are different. Firstly, physical laws which have exceptions can never be used to make perfectly accurate predictions unless by chance. Returning to Galileo’s law of the pendulum, this law cannot be expected to correctly describe the motion of a real-world pendulum since it assumes — among other things — that the bob is a point mass and that the pendulum is swung from a frictionless pivot-conditions that cannot be physically realised. Of course we might try to approximate these conditions (and this is an important part of experiment design), but we could never instantiate them fully. Hence the law will be systematically false in all real-world cases. Secondly, unlike the aspirin example, there are many physical laws which have known ‘idealization conditions’ which if placed as a limit on the scope of the law could be used to save it from falsity. A third difference is that physical laws are ‘piecemeal improvable’. Leszek Nowak (1980) has called this method ‘concretization’ as it involves further variables and mathematical operations being added to our equation in order to accommodate the additional lines of influence. An example of this can be seen among the gas laws. These laws tell us how the pressure and temperature of a gas are related. They get progressively more complicated as their accuracy improves by incorporating further factors (e.g. going from Boyle’s law to the Ideal Gas law to van der Waals equation).

Finally, as a result of knowing the idealization conditions physicists can predict the cases when the law will fail. Ronald Laymon (1989) takes this to be an important part of the confirmation process of idealized claims and explains how they can be tested by scientists despite never being fully realised in nature.

This suggests we are dealing with two different types of exception-ridden laws in the sciences — with one group mostly proliferating in the special sciences — and the other group—the physical sciences. Let us stick with tradition and call the first group *ceteris paribus* laws, since they are the ones that need a ‘*ceteris paribus* proviso’ to be saved from falsity. For reasons that will become clear, let us call the second group of exception-ridden laws — the ones mostly found in physics — *ideal laws*. The main differences between these laws can be summarised in the following table:

	<i>Ceteris Paribus</i> Laws	Ideal Laws
Does the law have negative instances?	Yes	Yes
Does the law have positive instances?	Yes	No
Can the antecedent clause be completed with material conditions?	No	Yes
Is the law concretizable (i.e. piecemeal improvable)?	No	Yes
Does the law predict the cases when it fails?	No	Yes

Examples of physical laws that fall into the category of ‘ideal law’ include: Newton’s laws of gravity and motion, Coulomb’s law, Snell’s law of refractive index, Ohm’s law, Maxwell’s equations, Galileo’s law of freefall, Hooke’s law, the law of thermal expansion and the Ideal Gas law. I suspect numerous other examples could be added to this list.⁴

I am not the first to propose that the laws of physics and their alleged exceptions are unlike those found in other sciences. John Earman and John Roberts have argued:

In general the problem of *ceteris paribus* qualifications is distinct from the problem of *idealizations*. Often the idealization can be stated in a precise closed form, (e.g. the ideal gas law...). Here the problem is not in saying precisely what is involved in the idealization but in relating it to the world which is not ideal. By contrast, many cp laws claim to be about unidealized real world situations but make indefinite claims about these situations. (Earman and Roberts 1999, p.457)

The question that needs answering is whether exceptions pose more or less of a threat to ideal laws than they do to traditional *ceteris paribus* laws? According to Earman, Roberts and Smith (2002), physicists should not be so worried, as there is no

real ‘problem of provisos’ for these laws analogous to Lange’s Dilemma. There is only the appearance of a problem if one conflates the basic laws of physical theory with differential equations governing the evolution of a specific type of physical system. Cartwright is one of many philosophers, they claim, who makes this mistake. Their point can be explained with reference to the equation for Newton’s law of universal gravitation (UG):

$$F = -G \frac{M_1 M_2}{R^2}$$

Cartwright tells us this law ‘lies’ because it fails to represent the behaviour of physical systems that involve forces other than gravity. But according to Earman, Roberts and Smith:

UG *cannot* misrepresent the motion of body, because it says nothing specific about such temporal behaviour. Only differential equations of evolution type — which might be derivable from UG together with other considerations — can be integrated to describe the temporal motion of a body or system of bodies. UG cannot be so integrated. Thus, it cannot misrepresent temporal motion. (2002, p.286)

Their point is that Newton’s law of universal gravity is a special force law. It does not tell us anything about motion or behaviour in the sense these terms are ordinarily understood, hence it cannot ‘lie’ nor have exceptions in the way Cartwright believes. It might appear that it can if one confuses the law itself with the application of the law to a type of physical system. This kind of application requires the creation of a differential equation whose integration with other laws (such as Newton’s laws of motion) depends first upon making non-nomic assumptions about what is present in the system in question. Kepler’s laws approximate a differential equation for Newton’s law because they assume only two gravitationally bound bodies. Earman, Roberts and Smith conclude that such equations — whilst commonly referred to as ‘laws’ — are not really genuine laws of nature.

If the differential equations are not laws, then what about the simple theoretical laws from which they are derived? Smith believes these are the physical equations that really should be given the title ‘law’. It is true that these are not about behaviour ordinarily understood, but Smith believes this is an erroneous assumption made by philosophers and should be abandoned:

All of this ought to be sufficient evidence to reject the standard presumption of empiricists following Hume that laws express constant conjunctions of temporally contiguous events. Laws like Universal Gravitation say *nothing* about temporal successions by themselves. When we do have a differential equation that does tell us something about temporal behaviour, it is generally not a law because it is derived from laws and non-law ingredients. (2002, p.249)

Principia 21(3): 317–337 (2017).

Smith himself does not propose an alternative anti-Humean theory, although he does suggest something like component forces will need to be taken as the primitive subject matter for laws (2002, p.252). He therefore poses a *new* dilemma for Humeans which we might call ‘Smith’s Dilemma’: either take the fundamental equations of physical theory (such as UG) as laws or take all the differential equations that can be derived from them as laws. Both of these options appear unacceptable. If the Humean takes the fundamental equations as laws they will need to acknowledge the existence of component forces which goes beyond regularity. On the other hand, if they take the differential equations to be laws (which are about occurrent behaviour) then there will be as many different laws as there are physical systems — one for each imaginable! One assumes that the second horn of this dilemma is unacceptable because scientific practice does not consider all these equations laws — rather, the laws of nature form a very special and limited set. Smith acknowledges that it is true that many differential equations are commonly referred to as ‘laws’ (Kepler’s laws being a case in point) but puts this down to nothing more than a historical accident that ‘should not be taken seriously’ (2002, p.247).

If Humeans are to find a solution to Smith’s Dilemma it would appear that their only option is to take some of the differential equations as laws. To do this requires finding a principled way of separating the differential equations which are laws from those which are not. Smith presupposes that no such principled means can be given: but this isn’t true. As he himself remarks, there is something very different about the differential equations which are commonly referred to as ‘laws’ — they are rules which allow us to construct ‘minimally working models’ (2002, p.254). The two-body system in which there are no other forces is the minimally working model for UG (indeed it is this model that Cartwright takes UG to be ‘about’). The reason why this model is *minimal* is that it makes the fewest number of assumptions possible in order for the mathematical relationship to be realised. For each such differential equation we can assign a set of ‘minimal assumptions’ which need to be present for the law to obtain. The discovery of these minimal assumptions seems to be a crucial part of scientific practice, something which Smith has perhaps overlooked. A Humean solution to Smith’s Dilemma for physical laws should therefore start with a theory that favours these minimal equations.

4. Ideal Laws as Lossy Compression Algorithms

One potential theory is the ‘Algorithmic Theory of Laws’ which has been advocated in a number of different forms by Richard Feynman (1965), Murray Gell-Mann (1987), John Barrow (1991), Paul Davies (1995), Gregory Chaitin (2005), Seth Lloyd (2006) and Terence Tomkow (2013).⁵ Put simply, the algorithmic theory of laws claims that

the laws of nature are the algorithms (or algorithmic component) in the best possible compression of nature. Although contemporary versions of this view appeal to ideas from computer science and information theory, the basic thought behind it goes back much further. In Ernst Mach (1838-1916) we see the idea expressed in terms of ‘economy of thought’ where the scientific enterprise is compared to a kind of business that wants to maximize profit using as few resources as possible:

The communication of scientific knowledge always involves description, that is, a mimetic reproduction of facts in thought, the object of which is to replace and save the trouble of new experience. Again, to save the labor of instruction and of acquisition, concise, abridged description is sought. This is really all the natural laws are. Knowing the value of the acceleration of gravity, and Galileo’s laws of descent, we possess simple and compendious directions for reproducing in thought all possible motions of falling bodies. (1894[1943], p.193)

Although Mach himself talks about lawhood in terms of the scientific laws, we can abstract somewhat, and arrive at a theory of what the laws of nature are by themselves. This can be illustrated with the help of the following example. Consider two sequences of symbols (A) and (B):

(A) 01101001100101101001011001101001...

(B) 01011010111001010101111101001000...

At first glance both appear to be equally random and contain the same amount of information in terms of number of bits. However, the string represented by (A) is an instance of a well-known sequence in mathematics (the ‘Thue-Morse’ sequence) and can be generated by following just two rules:

(Rule-1) If ‘0’ print ‘01’

(Rule-2) If ‘1’ print ‘10’

By themselves these two rules cannot produce anything, one also needs starting conditions in the form of singular data. In this instance that corresponds to the printing of ‘0’ once. It should quickly become apparent that no matter how long the original sequence in (A) is, it is always reproducible printing ‘0’ plus the two rules above. The length simply being a matter of *how many times* Rules-1 and 2 are applied. Unlike (B), string (A) is therefore highly compressible and can be compressed to data that contains two distinct parts: (i) unstructured singular data, i.e. ‘0’, and (ii) structured rules or algorithms, i.e. ‘Rules-1 and 2’. What makes compression possible is the existence of statistical regularities in the original data: these are summarized by the algorithm which exploit it for compression gains.

Now on the assumption that our world is a regular place, it too ought to be algorithmically compressible. If it is, then the laws of nature correspond to the algorithms in its best compression. This comment requires some qualification because the world is normally thought to be made up of physical objects and processes and ‘compression’ is typically an operation on symbols or data. Scientists make observations about the goings on in the world and take recordings. These recordings or measurements form a body of data. From the standpoint of the algorithmic theory, it is this data which scientists aim to compress and uncover the laws. But of course this data is only a ‘snapshot’ of all possible data that could have been recorded. The algorithmic theory therefore abstracts from the actual recording to include all possible recordings (in the entire history of the universe). If nature is understood as a source of information and all possible recordings its data, then we can define the laws as follows:

The Algorithmic Theory of Laws: x is a law of nature if, and only if, x is an algorithm and x is a member of $d^*(w)$

Where $d^*(w)$ is the *shortest possible description* of w — the total data in our universe. The existence of such a shortest possible description is guaranteed by algorithmic information theory (Li and Vitanyi 2003). Hence the laws of nature are the algorithms which when added to singular data form the best possible compression of ‘nature’s information’.

The algorithmic theory is compatible with both a Humean and non-Humean metaphysics; but it is far more likely to be appealing to Humeans because lawhood supervenes on patterns (regularities) in the data and does not require the data being about anything like necessary connections, essences or capacities. The Humean can simply fill-in what the data is about with their preferred ontology, i.e. such as Lewisian occurrent properties at space-time points. The algorithmic theory does not demand that the data be about observable entities and processes either: indeed, if data can be about unobservables as modern science take it to be, then the laws of nature in the final ‘best compression’ will also be about unobservables. What one considers to be the data that is compressed therefore depends on one’s metaphysics. What is important, however, is that it is not *necessary* for there to be necessary connections, essences or capacities in this data, in order for there to be laws of nature.

Before looking at how this theory copes with exceptions to laws, a few remarks are worth making. Firstly, on this view lawhood is a contingent matter. Whether or not there are laws at a world depends on whether that world is compressible. Hence whilst Kepler’s laws, for example, do a good job of compressing recordings of planetary motion in our world, they might not be the algorithms in the best compression of some distant possible world where the behavior of planets is very different from ours. Another point worth remarking is that on this view lawhood remains objective: for a given body of data it is an objective fact whether it is compressible and what

its best compression is. At this point it might be argued that since data requires a linguistic representation and since this is a choice ‘made by us’, what counts as the best compression is also a choice ‘made by us’. This criticism is analogous in form to that made by Armstrong (1983, p.68) and van Fraassen (1989, p.53) against Lewis’ original version of the BSA. I have argued elsewhere that thinking about laws in terms of compression provides a solution to this problem that is unavailable to the standard BSA. I therefore refer the interested reader to (Wheeler, 2016) rather than repeat those arguments here.

How does thinking about laws as algorithms help the Humean accommodate exceptions, and in particular, exceptions to the laws in physics? David Braddon-Mitchell has previously argued (2001) that Humean laws might tolerate exceptions if we utilise a distinction, commonplace in computer science, between *lossless* and *lossy* data compression. The key difference concerns whether information is lost during the compression process. An algorithm is said to be lossless if, after the original data has been compressed, the data can be decompressed to its original form without any loss of information. By contrast, an algorithm is said to be lossy if, after compression and decompression, part of the original message is lost. Modern compression methods typically use both lossy and lossless techniques depending on how much information loss can be tolerated. For example, methods designed for *text files* are lossless. This is because natural language cannot tolerate much degradation without rendering the message meaningless. On the other hand, most *image* and *audio* files are compressed using lossy compressors. This is because the human audio-visual system can only make tone and chromatic distinctions up to a degree and typical recording techniques often supply quality of data far in excess of this. These files, therefore, can be compressed in a lossy fashion without much noticeable degradation.

Braddon-Mitchell’s proposal was to accommodate laws with exceptions by supplementing Lewis’ BSA with this distinction:

What if we could keep the same level of simplicity and strength, but get closer to a different truth-related desideratum — the whole truth — at the cost of violating nothing but the truth clause. Or else dramatically increase the level of simplicity and strength, at the cost of some violation of the nothing but the truth clause? This would be analogous to lossy data compression. We might trade off *amount* of truth and simplicity against complete accuracy. In data compression terms, we might trade off compression ratio and completeness (or extensiveness within a domain in which the algorithm is appropriate) with accuracy. (2001, p.266)

Instead of the best system being a balance of only strength and simplicity — why not allow accuracy to also be part of the mix? On some occasions, if sacrificing truth delivers significant gains in simplicity, then shouldn’t this be preferred to a perfectly accurate system that is nonetheless long and complex?

Despite Braddon-Mitchell originally intending his explanation to go with the BSA, I think it forms a more natural supplement to the Algorithmic Theory of Laws. Recall that the BSA assumes laws are axioms (that is, statements) but in computer science compression is normally thought to be achieved by a program which is a kind of *rule* or *algorithm*. It is therefore more coherent to assume that if laws are doing compressing then they are algorithms rather than statements of fact. Returning to Mach's point of science being like a business, it seems reasonable that on occasion scientists would perform a 'cost-benefit analysis' and trade off some accuracy if the gains in resources saved (i.e. simplicity) are great enough. Since on the algorithmic approach it is compression which is the overriding factor for lawhood, laws which sacrifice accuracy for compression can be laws — and powerful laws — despite having exceptions. The Humean can then explain why some laws-despite having exceptions — are still genuine laws of nature.

A strong analogy exists between idealization in science and lossy data compression which provides further support for this account. During the design of an algorithm for compression, if a programmer is aiming for a lossy compression, then their method will typically involve a data preparation stage where the values of each datum are first transformed, so that the overall data set exhibits more regularity or structure than it originally did (Salomon 2008, p.50). This 'structure-adding stage' is needed for lossy compression to be possible and once carried out is often irreversible. But this kind of data-processing is a feature of scientific representation more generally (McAllister 2003). Drawing a line of best-fit on a graph and working out the equation for that line is just one simple illustration of this.

Because the structure-adding stage leads to a modification of the original values, this has consequences for what is returned by the lossy algorithm if used to decompress (i.e. to return the original message). When there are perceptible differences between the original and decompressed files, the differences are called 'compression artefacts'. One of the most striking examples of compression artefacts comes about through the use of the JPEG compression algorithm. The reason this compressor has proven so popular is the fact it gives the user the ability to define the level of compression achieved. Images *tree1.jpg* and *tree2.jpg* show the output of JPEG for the same image although with different factors of compression:

The image in '*tree1.jpg*' has been compressed to 50% of the original whereas '*tree2.jpg*' has been compressed to only 5% of the original. Clearly the greater the compression we choose for our image the less it will resemble the original. It could be argued that an analogous situation occurs when an ideal law is used to construct a working model in science. Suppose I set up a pendulum in the laboratory and record values for its mass and length. If I then plug these values back into a model and, by measuring the time-period, use my model to calculate a value for g (the force due to gravity) I shouldn't be surprised if the value returned fails to match the real value at



(a) tree1.jpg



(b) tree2.jpg

this position. This is because the model contains a number of ‘idealization artefacts’ — such as point-masses, perfectly rigid rods and frictionless pivots. Just as I should expect from using JPEG that my image of the tree will fail to resemble in all respects the real-world tree, so too should I expect that any derivations made from using the ideal model will fail to match the real-world exactly.⁶

The strong analogy that exists between the design and application of lossy compressors, and the discovery and use of ideal laws to construct models, shows that much can be gained by thinking of laws as algorithms for compressing empirical data. I now turn to show how this equips the Humean with a solution to Smith’s Dilemma.

5. Resolving Smith’s Dilemma

In section 3 we saw that Smith introduced a new dilemma for Humeans who wish to construe the fundamental laws of physics as genuine laws of nature. Taking UG as our example, Smith argued that this law cannot be interpreted as a regularity since, by itself, UG does not describe any temporal behaviour. Instead, Smith believes, UG is really about a component force — namely, *the force due to gravity*. This is unappealing for Humeans as component forces seem to equate to Hume’s ‘secret powers’ and are not directly measurable. It is possible, Smith goes on, to turn UG into a statement about temporal behaviour by transforming it into a differential equation. However, in the process of doing this we need to make certain non-nomic assumptions about the type of physical system described, and so, the resulting equation will lack universal application. This option is problematic as there appears to be no principled way of

choosing one type of differential equation over any other and so we end up with too many objects that can be classified as a ‘law’.

The Algorithmic Theory of Laws combined with the lossy/lossless compression distinction provides the Humean with a principled way of favouring one of these equations. For each set of differential equations (that describe a particular type of physical phenomenon) there will be one which provides the rules for constructing a ‘minimal working model’. In line with Mach’s view of science as a ‘profitable business’, the minimal equation will always be the preferred choice of scientists, since it provides a description of the physical phenomenon in question using as few computational resources as possible. In reality, scientists will be willing to forgo accuracy if the gains in compression are high enough. Just by analogy with the JPEG, if a certain number of compression artefacts can be tolerated, then it is rational to prefer a lossy compression.

At this point one might want to raise a problem for this explanation. Although it provides the scientist with criteria to select the ideal law, it does not explain why similar equations designed for more complex tokens of the phenomenon are not also considered laws, given that they too provide accurate compressed description. What grounds can be given, for instance, for favouring one lossy compression that achieves 50% compression over one that provides 35%?

I think in this case the Humean should accept that both equations are ‘lawlike’ (and so share many of the typical features philosophers associate with laws). But we should make a distinction between the ‘laws *for* X’ and the ‘laws *of* X’. Many algorithms can be laws for compressing a physical phenomenon, but only one of these will be the best one. We can argue that part of the search for laws is then the search for the best possible compression algorithm (lossless or lossy) that allows us to use that law for explanation and prediction with a reasonable chance of success. Of course, what counts as a ‘reasonable chance of success’ depends on what the law is to be used for. UG might be very well and good for successfully predicting the orbit of the Moon, but it will not be the preferred choice for predicting the motion of planets subject to significant relativistic effects, such as that of Mercury.

On this point it might be further objected that the mere fact there is choice concerning the lossiness of the compression means introducing an element of subjectivity into ideal laws. On a lossless reading of the algorithmic theory there is no subjectivity problem — given any empirical data set we identify the law with the algorithm in the *best possible* algorithmic compression of that data. But with lossy compression we know we can choose the factor of compression, provided we are willing to sacrifice some quality of information. The analogy with the JPEG compressor brings this point home even more. The reason why the JPEG compressor is able to achieve impressive compression gains is because it exploits a unique fact about human perception — that our eyes are far more sensitive to *contrast* differences than they are to *colour*

differences. Yet this is an idiosyncratic fact about human perception and shows that whilst a JPEG compression factor of about 20% is perfectly acceptable for humans, this might not be the case for other animals, whose eyes cannot tolerate these kinds of chromatic artefacts.

If correct, this would seem to imply that the ‘ideal laws’ are really the laws that are ‘ideal for humans’. As such there is no objective ‘Ideal Gas law’ only ‘the law that is Ideal for what Humans want to do with Gases’. Compression factors are determined by use which in our case is determined by human interests.

In spite of this, there is one sense in which the fidelity-loss tolerated by JPEG is *unlike* that of empirical laws. The fact human eyes have more sensitive rod and less sensitive cone photoreceptor cells is a psychological and *internal* fact about humans. In this case it is not possible for any further quality to be appreciated and so on a practical level the extra information is redundant. Now compare the case with scientific laws. Given the value of the temperature and volume of a gas we can calculate its pressure. Again there is a choice. We could use the Ideal Gas law (which is computationally cheap), or we could use the van der Waal’s equation (which is computationally more expensive). But this time it is a pragmatic *external* fact that determines whether the additional quality provided would be redundant or not. For example, if we are calculating the pressure of a gas on the surface of the Sun, then clearly we should opt for the Ideal Gas law. Whereas if we are calculating the pressure of a gas in deep space, then clearly we should opt for van der Waal’s equation (since at low temperatures inter-molecular forces have greater significance). In these kinds of cases there is a sense in which the amount of information loss that can be tolerated depends on external facts about the environment that the phenomenon is in and so is not wholly determined by human interests.

This is the best explanation Humeans have for explaining why some fundamental physical laws, despite having exceptions, can be genuinely lawlike. I admit some will feel the connection between ideal laws and human interests is too close. Yet I fail to see how else they can explain the existence of these laws without straying beyond regularity and a broadly Humean metaphysics.

References

- Armstrong, D. 1983. *What is a Law of Nature?* Cambridge: Cambridge University Press.
- Braddon-Mitchell, D. 2001. Lossy Laws. *Noûs* 35(2): 260–277.
- Callender, C.; Cohen, J. 2010. Special Sciences, Conspiracy and the Better Best System Account of Lawhood. *Erkenntnis* 73(3): 427–447.
- Cartwright, N. 1983. *How the Laws of Physics Lie*. Oxford: Oxford University Press.
- . 1989. *Nature’s Capacities and Their Measurement*. Oxford: Oxford University Press.
- . 1999. *The Dappled World*. Cambridge: Cambridge University Press.

- Cohen, J.; Callender, C. 2009. A Better Best System Account of Lawhood. *Philosophical Studies* **145**(1): 1–34.
- Davies, P. 1995. Algorithmic Compressibility, Fundamental and Phenomenological Laws. In: F. Weinert (ed.) *Laws of Nature: Essays on the Philosophical, Scientific and Historical Dimensions*, pp.248–267. Berlin: Walter de Gruyter & Co.
- Dorato, M. 2005a. The Laws of Nature and the Effectiveness of Mathematics. In: G. Boniolo (ed.) *The Role of Mathematics in Physical Sciences*, pp.131–144. Amsterdam: Springer.
- . 2005b. *The Software of the Universe*. Aldershot: Ashgate.
- Earman, J.; Roberts, J. 1999. Ceteris Paribus, There is No Problem of Provisos. *Synthese* **118**: 439–478.
- Earman, J.; Roberts, J.; Smith, S. 2002. Ceteris Paribus Lost. In: J. Earman; C. Glymour; S. Mitchell (eds.) *Ceteris Paribus Laws*, pp.5–25. Dordrecht: Kluwer.
- Feynman, R. 1965. *The Character of Physical Law*. London: BBC.
- Fodor, J. 1991. You can Fool Some of The People All of The Time, Everything Else being Equal: Hedged Laws and Psychological Explanations. *Mind* **100**(1): 19–34.
- Galileo, G. 1638[1954]. *Dialogues Concerning Two New Sciences*. Translated by H. Crew and A. de Silvio. New York: Dover.
- Gell-Mann, M. 1987. Simplicity and Complexity in the Description of Nature. (M. Gell-Mann, Performer, October 1). California Institute of Technology: Pasadena.
- Hausman, D. 1992. *The Inexact and Seperate Science of Economics*. Cambridge: Cambridge University Press.
- Huttemann, A. 2007. Causation, Laws and Dispositions. In: M. Kistler; B. Gnassounou (eds.) *Dispositions and Causal Powers*, pp.207–219. Aldershot: Ashgate.
- . 2014. Ceteris Paribus Laws in Physics. *Erkenntnis* **78**(10): 1715–1728.
- Kowalenko, R. 2011. The Epistemology of Hedged Laws. *Studies in History and Philosophy of Science* **42**: 445–452.
- Lange, M. 1993. Natural Laws and the Problem of Provisos. *Erkenntnis* **38**: 233–248.
- Laymon, R. 1989. Cartwright and the Lying Laws of Physics. *The Journal of Philosophy* **86**(7): 353–372.
- Lewis, D. 1973. *Counterfactuals*. Oxford: Blackwell.
- . 1980. A Subjectivist's Guide to Objective Chance. *The University of Western Ontario Series in Philosophy of Science* **15**: 267–297.
- . 1986. Causation. In: D. Lewis, *Philosophical Papers, Volume II*, pp.159–213. Oxford: Oxford University Press.
- Li, M.; Vintanyi, P. 1993. *An Introduction to Kolmogorov Complexity and its Applications*. New York: Springer-Verlag.
- Mach, E. 1894[1943]. *Popular Scientific Lectures*. Illinois: Open Court.
- McAllister, J. 2003. Algorithmic Randomness in Empirical Data. *Studies in the History and Philosophy of Science* **34**: 633–646.
- . 2005. Algorithmic Compression of Empirical Data: Reply to Twardy, Gardner and Dowe. *Studies in the History and Philosophy of Science* **36**: 403–410.
- Mill, J. S. 1895[1974]. *A System of Logic*. London: Routledge Kegan and Paul.
- Mumford, S. 1998. *Dispositions*. Oxford: Oxford University Press.
- Nelson, R. A.; Olsson, M. G. 1986. The pendulum-Rich physics from a simple system. *American Journal of Physics* **54**(2): 112–121.

- Nowak, L. 1980. *The Structure of Idealization*. Dordrecht: Reidel.
- Pietroski, P.; Rey, G. 1995. When Other Things Aren't Equal: Saving Ceteris Paribus Laws from Vacuity. *The British Journal for the Philosophy of Science* **46**(1): 81–110.
- Ramsey, F. 1990. *Philosophical Papers*. Cambridge: Cambridge University Press.
- Saloman, D. 2008. *A Concise Introduction to Data Compression*. London: Springer-Verlag.
- Schrenk, M. 2007. *The Metaphysics of Ceteris Paribus Laws*. Lancaster: Verlag.
- . 2014. Better Best Systems and the Issue of CP-Laws. *Erkenntnis* **79**(S10): 1787–1799.
- Schurz, G. 2001. Pietroski and Rey on Ceteris Paribus Laws. *British Journal for the Philosophy of Science* **52**: 359–370.
- Silverberg, A. 1996. Psychological Laws and Non-Monotonic Logic. *Erkenntnis* **44**: 199–224.
- Smith, S. 2002. Violated Laws, Ceteris Paribus Clauses, and Capacities. *Synthese* **130**: 235–264.
- Sober, E. 1984. *The Nature of Selection*. London: The MIT Press.
- Suarez, M. 2009. *Fictions in Science*. New York: Routledge.
- Twardy, C.; Gardner, S.; Dowe, D. 2005. Empirical Data Sets are Algorithmically Compressible: Reply to McAllister? *Studies in the History and Philosophy of Science* **36**: 391–402.
- van Fraassen, B. 1989. *Laws and Symmetry*. Oxford: Clarendon Press.
- Wheeler, B. 2016. Simplicity, Language-Dependency and the Best System Account of Laws. *Theoria: An International Journal for the Theory, History and Foundations of Science* **31**: 189–206.
- Woodward, J. 2002. There is No Such Thing as a Ceteris Paribus Law. *Erkenntnis* **57**: 303–328.

BILLY WHEELER
 Sun Yat-sen University
 Department of Philosophy (Zhuhai)
 wheeler@mail.sisu.edu.cn

RECEIVED: 09/07/2017

ACCEPTED: 21/11/2017

Notes

¹For a critical discussion of semantic solutions see Earman and Roberts (1999), Schurz (2001) and Woodward (2002).

²Lewis' final theory is bit more complicated than this in that he allows for the possibility that some systems might be 'tied' for best balance. To accommodate probabilistic laws, Lewis also adds a third desired feature which he calls 'fit', so that laws are the axioms in a system best balancing simplicity, strength and *closeness to the actual world history*. For further details, see Lewis (1994).

³This part of Schrenk's version of the BSA is also adopted by Callender and Cohen (2009; 2010).

⁴I treat it here as an empirical/historical fact that most ideal laws are to be found in the physical sciences and *ceteris paribus* laws in the special sciences. I see no analytical reason for thinking things could not be vice-versa. In addition, there are grounds for thinking that Darwin's law of natural selection is best classified as an ideal law as its idealization conditions are clearly statable (Sober 1984, p.38).

Principia **21**(3): 317–337 (2017).

⁵Algorithmic approaches have also been discussed by Mauro Dorato (2005a; 2005b) and James McAllister (2003). See McAllister (2003; 2005) for criticism of the view and Twardy, Gardner and Dowe (2005) for a defence.

⁶The account offered here therefore also provides a novel explanation for the existence of ‘fictions’ in scientific theories. See Mauricio Suarez (2009) for the context to the wider debate on the function of fictions in science.