Laws of nature

The notion of a law of nature is fundamental to science. In one sense this is obvious, in that much of science is concerned with the discovery of laws (which are often named after their discoverers hence Boyle's law, Newton's laws, Ostwald's law, Mendel's laws, and so on). If this were the only way in which laws are important, then their claim to be fundamental would be weak. After all, science concerns itself with much more than just the uncovering of laws. Explaining, categorizing, detecting causes, measuring, and predicting are other aims of scientific activity, each quite distinct from law-detecting. A claim of this book is that laws are important because each of these activities depends on the existence of laws. For example, take Henry Cavendish's attempt to measure the gravitational constant G using the torsion balance. If you were to have asked him what this constant is, he would have told you that it is the ratio of the gravitational force between two objects and the product of their masses divided by the square of their separation. If you were then to ask why this ratio is constant, then the answer would be that it is a law of nature that it is so. If there were no law of gravitation, then there would be no gravitational constant; and if there were no gravitational constant there would be nothing that counts as the measurement of that constant. So the existence and nature of a law was essential to Cavendish's measuring activities. Somewhat contentiously, I think that the notion of explanation similarly depends upon that of law, and even more contentiously I think that causes are also dependent on laws. A defence of these views must wait until the next chapter. I hope it is clear already that laws are important, a fact which can be seen by considering that if there were no laws at all the world would be an intrinsically chaotic and random place in which science would be impossible. (That is assuming that a world without laws could exist - which I am inclined to think doubtful, since to be any sort of thing is to be subject to some laws.) Therefore, in this chapter we will consider the question: *What is a law of nature?* Upon our answer, much that follows will depend.

Before going on to discuss what laws of nature are, we need to be clear about what laws are not. Laws need to be distinguished from statements of laws and from our theories about what laws there are. Laws are things in the world which we try to discover. In this sense they are facts or are like them. A statement of a law is a linguistic item and so need not exist, even though the corresponding law exists (for instance if there were no people to utter or write statements). Similarly, theories are human creations but laws of nature are not. The laws are there whatever we do -- one of the tasks of the scientist is to speculate about them and investigate them. A scientist may come up with a theory that will be false or true depending on what the laws actually are. And, of course, there may be laws about which we know nothing. So, for instance, the law of universal gravitation was discovered by Newton. A statement of the law, contained in his theory of planetary motion, first appeared in his *Principia mathematica*. In the following we are interested in what laws are -- that is, what sort of thing it was that Newton discovered, and not in his statement of the law or his theories containing that statement.[8] Furthermore, we are not here asking the question: Can we know the existence of any laws? Although the question “What is an X?” is connected to the question “How do we know about Xs?”, they are still distinct questions. It seems sensible to start with the former -- after all, if we have no idea what Xs are, we are unlikely to have a good answer to the question of how we know about them.

1. Minimalism about laws - the simple regularity theory

In the Introduction we saw that the function of (Humean) induction is to take us from observations of particular cases to generalizations. It is reasonable to think that this is how we come to know laws -- if we can come to know them at all. Corresponding to a true inductive conclusion, a true generalization, will be a *regularity*. A regularity is just a general fact. So one inductive conclusion we might draw is that all emeralds are green, and another is that all colitis patients suffer from anaemia. If these generalizations are true, then it is a fact that each and every emerald is green and that each and every colitis sufferer is anaemic. These facts are what I have called regularities. The first view of laws I shall consider is perhaps the most natural one. It is that laws are just regularities.

This view expresses something that I shall call *minimalism* about laws. Minimalism takes a particular view of the relation between a law and its instances. If it is a law that bodies in free fall near the surface of the Earth accelerate towards its centre at 9.8 ms^{-2}, then a particular apple falling to the ground and accelerating at 9.8 ms^{-2} is an *instance* of this law. The minimalist says that the law is essentially no more than the collection of all such instances. There have been, and will be, many particular instances, some observed but most not, of objects accelerating towards the centre of the Earth at this rate. The law, according to the minimalist, is simply the regular occurrence of its instances. Correspondingly, the statement of a law will be the generalization or summary of these instances.

Minimalism is an expression of *empiricism*, which, in broad terms, demands that our concepts be explicable in terms that relate to our experiences. Empiricist minimalism traces its ancestry at least as far back as David Hume. By defining laws in terms of regularities we are satisfying this requirement (as long as the facts making up the regularities are of the sort that can be experienced). Later we shall come across an approach to laws that is not empiricist.

The simplest version of minimalism says that laws and regularities are the same. This is called the *simple regularity theory* (SRT) of laws.

SRT: It is a law that Fs are Gs *if and only if* all Fs are Gs.

---

[8] It must be admitted that scientists' (and philosophers') usage is a bit loose in this regard. We do talk about Kepler's laws and the ideal gas laws, and in doing so we must be talking about certain statements or theories -- for there are no such laws, these theories being only close approximations to the truth.

While the SRT has the merit of simplicity it suffers from the rather greater demerit that it is false. If it were true, then all and only regularities would be laws. But this is not the case.

The most obvious problem is that the existence of a simple regularity is not sufficient for there to be a corresponding law, i.e. there are simple regularities that are not laws. A criticism is also made that it is not even necessary for there to be a regularity for the corresponding law to exist. That is, there are laws without the appropriate regularities.

Regularities that are not laws

I will start with the objection that being a regularity is not sufficient for being a law. Consider the following regularities:

(a) All persisting lumps of pure gold-195 have a mass less than 1,000 kg.

(b) All persisting lumps of pure uranium-235 have a mass of less than 1,000 kg.9

Both (a) and (b) state true generalizations. But (a) is accidental and (b) is law-like. It is no law that there are no lumps of the pure isotope of gold -- we could make one if we thought it worth our while. However, it is a law that there are no such lumps of uranium-235, because 1,000 kg exceeds the critical mass of that isotope (something less than a kilogram) and so any such lump would cause its own chain reaction and self-destruct. What this shows is that there can be two very similar looking regularities, one of which is a law and the other not.

This is not an isolated example. There are an indefinite number of regularities that are not laws. Take the generalization: all planets with intelligent life forms have a single moon. For the sake of argument, let us imagine that the Earth is the only planet in the universe with intelligent life and that there could exist intelligent life on a planet with no moons or more than one. (For all I know, these propositions are quite likely to be true.) Under these circumstances, the above generalization would be true, even though there is only one instance of it. But it would not be a law; it is just a coincidence. The point here is that the SRT does not distinguish between genuine laws and mere coincidences. What we have done is to find a property to take the place of F which has just one instance and then we take any other property of that instance for G. Then "All Fs are Gs" will be true. And, with a little thought, we can find any number of such spurious coincidental regularities. For most things that there are we could list enough of their general properties to distinguish one thing from everything else. So, for instance, with a person, call her Alice, we just list her hair colour, eye colour, height, weight, age, sex, other distinguishing features, and so on in enough detail that only Alice has precisely those qualities. These qualities we bundle together as a single property F. So only Alice is F. Then choose some other property of Alice (not necessarily unique to her), say the fact that she plays the oboe. Then we will have a

---

[9] Examples of this sort are to be found in the writings of Carl Hempel and Hans Reichenbach.

true generalization that all people who are F (i.e. have fair hair, green eyes, are 1.63 m tall, weigh 59.8 kg, have a retrousse nose, etc.) play the oboe. But we do not want to regard this as a law, since the detail listed under F may have nothing whatsoever to do with an interest in and talent for playing the oboe.

The minimalist who wants to defend the SRT might say that these examples look rather contrived. First, is it right to bundle a lot of properties together as one property F? Secondly, can just one instance be regarded even as a regularity? (If it is not a regularity then it will not be a counterinstance.) However, I do not think that the minimalist can make much headway with these defences.

To the general remark that the examples look rather contrived, the critic of the SRT has two responses. First, not all the cases are contrived, as we can see from the gold and uranium example. We can find others. One famous case is that of Bode's "law" of planetary orbits. In 1772, J. E. Bode showed that the radii of known planetary orbits fit the following formula: $0.4 + 0.3 \times 2n$ (measured in astronomical units) where $n = 0$ for Venus, 1 for the Earth, 2 for Mars, and so on, including the minor planets. (Mercury could be included by ignoring the second term, its orbital radius being 0.4 astronomical units.) Remarkably, Bode's law was confirmed by the later discovery of Uranus in 1781. Some commentators argued that the hypothesis was so well confirmed that it achieved the status of a law, and consequently ruled out speculation concerning the existence of a possible asteroid between the Earth and Mars. Bode's "law" was eventually refuted by the observation of such an asteroid, and later by the discovery of the planet Neptune, which did not fit the pattern. What Bode's non-law shows is that there can be remarkable uniformities in nature that are purely coincidental. In this case the accidental nature was shown by the existence of planets not conforming to the proposed law. But Neptune, and indeed Pluto too, might well have had orbits fitting Bode's formula. Such a coincidence would still have been just that, a coincidence, and not sufficient to raise its status to that of a law.

Secondly, the critic may respond that the fact that we can contrive regularities is just the point. The SRT is so simple that it allows in all sorts of made-up regularities that are patently not laws. At the very least the SRT will have to be amended and sharpened up to exclude them. For instance, taking the first specific point, as I have stated it the SRT does not specify what may or may not be substituted for F and G. Certainly it is an important question whether compounds of properties are themselves also properties. In the Alice example I stuck a whole lot of properties together and called them F. But perhaps sticking properties together in this way does always yield a new property. In which case we might want to say that only uncompounded properties may be substituted for F and G in the schema for the SRT.

However, amending the SRT to exclude Fs that are compound will not help matters anyway, for two reasons. First, there is no reason why there should not be uncompounded properties with unique instances. Secondly, some laws do involve compounds -- the gas laws relate the pressure of a gas to the compound of its volume and temperature. To exclude regularities with compounds of properties would be to exclude a regularity for which there is a corresponding law. To the second point, that the regularities constructed have only one instance, one rejoinder must be this: why cannot a law have just one instance? It is conceivable that there are laws the only instance of which is the Big Bang. Indeed, a law might have no instances at all. Most of the transuraniurn elements do not exist in nature and must be produced artificially in laboratories or nuclear explosions. Given the difficulty and expense of producing these isotopes and because of their short half-lives it is not surprising that many tests and experiments that might have been carried out have not been. Their electrical conductivity has not been examined, nor has their chemical behaviour. There must be laws governing the chemical and physical behaviour of these elements under circumstances which have never and never will arise for them. There must be facts about whether nobelium-254, which is produced only in the laboratory, burns in oxygen and, if so, what the colour of its flame is, what its oxide is like, and so forth; these facts will be determined by laws of nature, laws which in this case have no instances.

So we do not want to exclude something from being a law just because it has few instances, even just one instance, or none at all. At the same time the possibility of instanceless laws raises another problem for the SRT similar to the first. According to the SRT empty laws will be empty regularities - cases of "all Fs are Gs" where there are no Fs. There is no problem with this; it is standard practice in logic to regard all empty generalizations as trivially true.[10] What is a problem is how to distinguish those empty regularities that are laws from all the other empty regularities. After all, a trivial empty regularity exhibits precisely as much regularity as an empty law.

Let us look at a different problem for the SRT. This concerns functional laws. The gas law is an example of a functional law. I t says that one magnitude -- the pressure of a body of gas -- is a function of other magnitudes, often expressed in a formula such as:

$$P = kT/V$$

where P is pressure, T is temperature, V is volume, and k is a constant. In regarding this as a law, we believe that T and V can take any of a continuous range of values, and that P will correspondingly take the value given by the formula. Actual gases will never take all of the infinite range of values allowed for by this formula. In this respect the formula goes beyond the regularity of what actually occurs in the history of the universe. But the SRT is committed to saying that a law is just a summary of its instances and does not seek to go beyond them. If the simple regularity theorist sticks to this commitment, the function ought to be a partial or gappy one, leaving out values that are not actually instantiated. Would such a gappy "law" really be a law? One's intuition is that a law should cover the uninstantiated values too. If the SRT is to be modified to allow filling in of the gaps, then this needs to be justified. Furthermore, the filling in of the gaps in one way rather than another needs justification. Of course there may well be a perfectly natural and obvious way of doing this, such as fitting a simple curve to the given points. The critic of the SRT will argue that this is justified because this total (non-gappy) function is the best explanation of the instantiated values. But the simple regularity theorist cannot argue in this way, because this argument accepts that a law is something else other than the totality of its instances. As far as what actually occurs is concerned, one function which fits the points is as

[10] According to this view, the claim that all living dodos live in Australia is true; and it is also true that all living dodos live in Switzerland. Perhaps it is odd to regard such statements as true, but if we do not regard them as such we cannot count any of the empty generalizations about the artificial elements as being true -- but some must be, namely those that correspond to laws.

good as any other. From the SRT point of view, the facts cannot decide between two such possible functions. But, if the facts do not decide, we cannot make an arbitrary choice, say choosing the simplest for the sake of convenience, as this would introduce an arbitrary element into the notion of lawhood. Nor can we allow all the functions that fit the facts to be laws. The reason for this is the same as the reason why we cannot choose one arbitrarily and also the same as the reason why we cannot have all empty generalizations as laws. This reason is that we expect laws to give us determinate answers to questions of what would have happened in counterfactual situations -- that is situations that did not occur, but might have.

Laws and counterfactuals

Freddie's car is black and when he left it in the sun it got hot very quickly. The statement "Had Freddie's car been white, it would have got hot less quickly" is an example of a counterfactual statement. It is not about what did happen, but what would have happened in a possible but notactual (a counter-to-fact) situation (i.e. Freddie's car being white rather than black). The counterfactual in question is true. And it is true because it is a law that white things absorb heat less rapidly than black things. Laws support counterfactuals.

We saw above that every empty regularity will be true and hence will be a law, according to the SRT. This is an undesirable conclusion. Counterfactuals help us see why. Take some property with no instances, F. If we allowed all empty regularities to be laws we would have both law 1 "it is a law that Fs are Gs'' and law 2 "it is a law that Fs are not-Gs". What would have happened if *a*, which is not F had been F? According to law 1, *a* would have been G, while law 2 says *a* would have been not-G. So they cannot both really be laws. Similarly, we cannot have both of two distinct functions governing the same magnitudes being laws, even if they agree in their values for actual instances and diverge only for non-actual values. For the two functional laws will contradict one another in the conclusion of the counterfactuals they support when we ask what values would P have taken had T and V taken such-and-such (non-actual) values.

Counterfactuals also underline the difference between accidental and nomic regularities. Recall the regularities concerning very large lumps of gold and uranium isotopes. There are no such lumps of either. In the case of uranium-235, there could not be such lumps, there being a law that there are no such lumps. On the other hand, there is no law concerning large lumps of gold, and so there could have been persisting 2,000 kg lumps of gold-195. In this way counterfactuals distinguish between laws and accidents.

Some philosophers think that the very fact that laws support counterfactuals is enough to show the minimalist to be wrong (and the SRT supporter in particular). The reasoning is that counterfactuals go beyond the actual instances of a law, as they tell us what would have happened in possible but non-actual circumstances. And so the minimalist must be mistaken in regarding laws merely as some sort of summary of their actual instances. This argument seems powerful, but I think it is not a good line for the anti-minimalist to pursue. The problem is that counterfactuals are not any better understood than laws, and one can argue that our understanding of counterfactuals is dependent on our notion of law or something like it, in which case corresponding to the minimalist account of laws will be a minimalist account of counterfactuals.[11] You can see that this response is plausible by considering that counterfactuals are read as if there is a hidden clause, for instance "Freddie's car would have got hot less quickly had it been white and everything else *and been the same as far as possible*". (Which is why one cannot reject the counterfactual by saying that had Freddie's car been white, the Sun might not have been shining.) The clause which says that everything should be the same as far as possible requires among other things, like the weather being the same, that the laws of nature be the same. So one can say that laws support counterfactuals only because counterfactuals implicitly refer to laws. Counterfactuals therefore have nothing to tell us about the analysis of laws. Consider the fact that laws do not support all counterfactuals -- particularly those counterfactuals relating to the way things would be with different laws. For instance one could ask how quickly would two things have accelerated towards one another had the gravitational constant G been twice what it is. The actual law of gravitation clearly does not support the correct answer to this counterfactual question.

Laws that are not regularities: probabilistic laws

So far there is mounting evidence that being a simple regularity is not sufficient for being a law. There are many regularities that do not constitute laws of nature:

(a) accidental regularities
(b) contrived regularities
(c) uninstantiated trivial regularities
(d) competing functional regularities.

The natural response on the part of the minimalist is to try to amend the SRT by adding further conditions that will reduce the range of regularities which qualify as laws. So the original idea is maintained, that a law is a regularity, and with an appropriate amendment it will now be that laws are a certain sort of regularity, not any old regularity.

This is the line I will examine shortly. But before doing so I want to consider an argument which suggests that being a regularity is not even sufficient for being a law. That is, there are law's that are not simple regularities. If this line of thinking were correct, then it would be no good trying to improve the SRT by adding extra conditions to cut down the regularities to the ones we want, as this would still leave out some laws that do not have a corresponding regularity.

The problem concerns probabilistic laws. Probabilistic laws are common in nuclear physics. Atomic nuclei as well as individual particles are prone to decay. This tendency to decay can be quantified as the probability that a nucleus will decay within a certain

[11] Such accounts are provided by J. L. Mackie and David Lewis, whose views on counterfactuals are, despite superficial dissimilarites, closely related.

period. (When the probability is one-half, the corresponding period of time is called the *half-life*.) So a law of nuclear physics may say that nuclei of a certain kind have a probability p of decaying within time t. What is the regularity here? The SRT, as it stands, has no answer to this question. But an answer, which satisfies the minimalist's aim of portraying laws as summaries of the individual facts, is this. The law just mentioned will be equivalent to the fact that, of all the relevant particles taken together, a proportion p will have decayed within t. (Note that we would find out what the value of p is by looking at the proportion that decays in observed samples. Another feature that might be included in the law is *resiliency*, i.e. that p is the proportion which decays in all appropriate subpopulations, and not just the population as a whole.)

The problem arises when we consider each of the many particles individually. Each particle has a probability p of decaying in time t. This is perfectly consistent with the particle decaying well before t or well after t. So the law allows for any particle to decay only after time t^* which is later than t. And what goes for one particle goes for all. Thus the law is consistent with all the particles decaying only after t^*, in which case the proportion decaying by t is zero. This means that we have a law the instances of which do not form the sort of regularity the minimalist requires. The minimalist requires the proportion p to decay by t. We would certainly expect that. But this is by no means necessary. While it is extremely unlikely that no particle will decay until after t, it is not impossible. Another, less extreme, case would be this. Instead of the proportion p decaying within t, a very slightly smaller proportion than p might decay. The chance of this is not only greater than zero, i.e. a possibility, but may even be quite high.

The minimalist's guiding intuition is that the existence and form of a law is determined by its instances. Here we have a case where our intuitions about laws allow for a radical divergence between the law and its instances. And so such a law is one which could not be a regularity of the sort required by the minimalist. This would appear to be a damning argument against minimalists, not even allowing for an amendment of their position by the addition of extra conditions. Nonetheless, I think the argument employed against the minimalist on this point could be seen as simply begging the question. The argument starts by considering an individual particle subject to a probabilistic law. The law might be that such particles have a half-life of time t. The argument points out that this law is consistent with the particle decaying at time t^* after t. This is all true. The argument then proceeded to claim that any collection could, therefore, be such that all its particles decay after t, including the collection of all the particles over all time.

The form of this argument is this: what is possible for any individual particle is possible for a collection of particles; or, more generally, if it is possible that X, and it is possible that Y, and it is possible that Z, and so on, then it is possible that X and Y and Z, etc. This form of argument certainly is not logically valid. Without looking outside, I think to myself it is possible that it is raining and it is possible that it is not raining, but I certainly should not conclude that it is possible that it is both raining and not raining. More pertinent to this issue is another counter-example. It is possible for one person to be taller than average, but it is not possible for everyone to be taller than average. This is because the notion of an average is logically related to the properties of a whole group. What this suggests is that the relevant step in the argument against the minimalist could only be valid if there is a logical gap between a law and the collection of its instances. But this is precisely what the minimalist denies. For, on the contrary, the minimalist claims that there is no gap; rather there is an intimate logical link in that the probabilistic law is some sort of sophisticated averaging out over its instances. So it turns out that the contentious step in the argument is invalid or, at best, question begging against the minimalist.

Still, even if the argument is invalid, it does help alert our intuitions to an implausibility with minimalism. We may not be able to prove that there can be a gap between the chance of decay that each atom has and the proportion of particles actually decaying -- but the fact that such a gap does seem possible is evidence against the minimalist.

2. The systematic account of laws of nature

We have seen so far that not all regularities are also laws, though it is possible for the minimalist to resist the argument that not all laws are regularities. What the minimalist needs to do is to find some further conditions, in addition to being a regularity, which will pare down the class of regularities to just those that are laws.

Recall that the minimalist wants laws simply to generalize over their instances. Law statements will be summaries of the facts about those instances. If this is right we should not expect every regularity to be a law. For we can summarize a set of facts without having to mention every regularity that they display, and an efficient summary would only record sufficient regularities to capture all the facts in which we are interested. Such a summary would be *systematic*; rather than being a list of separate regularities, it would consist of a shorter list of interlocking regularities, which together are enough to encapsulate all the facts. It would be as simple as possible, but also as strong as possible, capturing as many facts and possible facts as it can.

What I mean here by "capturing" a fact is the subsuming of that fact under some law. We want it to be that laws capture all and only their instances. So if the fact is that some object a has the property G, then this would be captured by a law of the form all Fs are Gs (where a is F). For example, if the fact we are interested in is the explosive behaviour of a piece of lithium, then this fact can be captured by the law that lithium reacts explosively on contact with water (if in this case it was in contact with water). If the fact is that some magnitude M has the value m then the fact is captured by a law of the form $M = f(X, Y, Z \dots)$ (where the value of X is x, the value of Y is y, etc., and $m = f(x, y, z \dots)$. An example of this might be the fact that the current through a circuit is 0.2 A. This fact can be captured by Ohm's law $V = IR$, if for instance there is a potential difference of 10 V applied to a circuit that has a resistance of 50 W. As we shall see in the next chapter, this is tantamount to saying that the facts in question have explanations. If a fact has an explanation according to one law, we do not need another independent law to explain the same fact. Let it be the case that certain objects, a, b, c, and d are all G. Suppose that they are also all F and also the only Fs that there are. We might think then that it is a law that Fs are Gs and that this law explains why each of the objects is G. But we could do without this law and the explanations it furnishes if it is better established that there is a law that Hs are Gs and *a* is H, and that there is a law that Js are Gs and that *b* is J, and that there is a law that Ks are Gs and that *c* is K, and so on. In this case the proposed law that Fs are Gs is redundant. We can instead regard the fact that all Fs are Gs as an accidental regularity and not a law-like one. If we symbolize "*a* is F" by "F*a*", then as Figure 1.1 shows, we could organize the facts into an economical system with three rather than four laws.

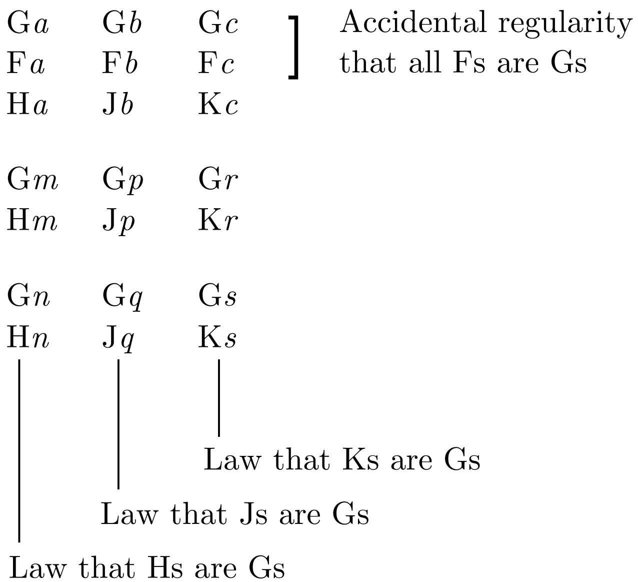

Figure 1. 1

This organizing of facts into an economical system allows us to distinguish between accidental regularities and laws. This was one of the problems facing the SRT. The systematic approach does better. It also easily accommodates the other problem case, functional laws. The simplest way of systematizing a collection of distinct points on a graph is to draw the simplest line through them. And we will only want one line, as adding a second line through the same points will make the- system much less simple without adding to the power of the system, because the first line already captures the facts we are interested in. Thus in a systematic version of minimalism we will have non-gappy functional laws that are unique in their domain.

These are the virtues of the account of laws first given by Frank Ramsey and later developed by David Lewis.[12] The idea is that we regard the system as axiomatized, that is to say we boil the system down to the most general principles from which the regularities that are laws follow. A collection of facts may be axiomatized in many (an infinite number of) ways. So the appropriate axiomatization will be that which, as discussed above, is as simple as possible and as strong as possible. The strength of a

[12] See Lewis, *Counterfactuals*, pp. 72-77. Lewis refers (p. 73) to an unpublished note by F. P. Ramsey.

proposition can be regarded as its informativeness. So, considering (a) all emeralds are green, (b) all emeralds are coloured, and (c) all South American emeralds are green, (a) is more informative and so stronger than both (b) and (c).

Formally, the systematic account says:

> A regularity is a law of nature if and only if it appears as a theorem or axiom in that true deductive system which achieves a best combination of simplicity and strength.

Remember that this is a "God's eye" view, not ours, in the sense that we do not know all the facts that there are and so do not know in that direct way what the best axioniatization of them is. What we are saying is that this is what a law of nature is. Our ignorance of the best axiomatization of everything ignorance of the laws of nature. Of course we are not entirely ignorant of them. If science is any good we know some of them, or approximations to them at any rate. And it can be said in support of Ramsey and Lewis, that the sorts of thing we look for in a theory which proposes a law of nature are precisely what they say something ought to have in order to be a law. First we look to see whether it is supported by evidence in the form of instances, i.e. whether it truly is a regularity. But we will also ask whether it is simple, powerful, and integrates with the other laws we believe to exist.

The systematic view is not entirely without its problems. First, the notion of *simplicity* is important to the systematic characterization of law, yet simplicity is a notoriously difficult notion to pin down. In particular, one might think that simplicity is a concept which has a significant subjective component to it. What may appear simple to one person might look rather complex to another. In which case the concept of law would also have a subjective element that conflicts with our intuition that laws are objective and independent of our perspective. Another way of looking at the problem of simplicity is to return to Goodman's puzzle about the concept "grue". The fact that we can replace the simple looking "X is grue" by the complex looking "either X is green and observed before midnight on 31 December 2000 or X is blue and not observed before midnight on 31 December 2000", and vice versa, shows that we cannot be sure of depending merely on linguistic appearance to tell us the difference. The simple law that emeralds are green appears to be the same as the complex law that emeralds are grue, if observed before midnight on 31 December 2000, or bleen if not observed before midnight on 31 December 2000. Without something like a solution to Goodman's problem we have no reason to prefer one set of concepts to another when trying to pin down the idea of simplicity.

The second problem with the systematic view is that as I have presented it, i.e. it presumes there is precisely one system that optimally combines strength and simplicity. For a start it is not laid down how we are supposed to weigh simplicity and strength against one another. We could add more laws to capture more potential facts and thus gain strength, but with a loss in simplicity. Alternatively, we might favour a simpler system that has less strength. Again there is the suspicion that the minimalist's answer to this problem may be objectionably subjective. Even if there is a clear and objective way of balancing these two, it may yet turn out that two or more distinct systems do equally well. So which are our laws? Perhaps something is a law if it appears in any one of the optimal systems. But this will not do, because the different systems might have conflicting laws that lead to incompatible counterfactuals. On the other hand, we may be more restrictive and accept as laws only those regularities that appear in all the optimal systems.[13] However, we may find that there are very few such common regularities, and perhaps none at all.

Basic laws and derived laws

It is a law-like and true (almost) generalization that all objects of 10 kg when subjected to a force of 40 N accelerate at 4 m s^{-2}. But this generalization would not feature as an axiom of our system, because it would be more efficient to have the generalization that the acceleration of a body is equal to the ratio of the resultant force acting upon it and its mass. By having that generalization we can dispense with the many generalizations that mention specific masses, forces, and accelerations. From the overarching generalization these more specific ones may be derived. That is to say, our system will be axiomatic; it will comprise the smallest set of laws from which it is possible to deduce all other laws. This shows an increase in simplicity -- few generalizations instead of many -- and in strength, as the overarching generalization will have consequences not entailed by the collection of specific generalizations. In the example there may be some values of the masses, forces, and accelerations for which there are no specific generalizations, just because nothing ever both had that mass and was subject to that force.

This illustrates a distinction between more fundamental laws and those derived from them. Most basic of all are those laws that form the axioms of the optimal system. If the laws of all the sciences can be reduced to laws of physics, as some argue is true, at least for chemistry, then the basic laws will be the fundamental laws of physics. If, on the other hand, there is some field the laws of which are not reducible, then the basic laws will include laws of that field. It might be that we do not yet actually know any basic laws -- all the laws we are acquainted with would then be derived laws.

Laws and accidents

I hope that it should now appear that the systematic view of laws is a considerable improvement on the SRT. While it still has problems, the account looks very much as if it is materially adequate. An analysis of a concept is *materially adequate* if it is true that were anything to be a case of the concept (the *analysandum* -- the thing to be analysed) it would also be a case of the analysis, and vice versa. So “female fox” is a materially adequate analysis of “vixen” only when it is true that if something were a female fox it would also be a vixen, and if something were a vixen it would also be a female fox. The systematic account looks to be materially adequate because the

---

[13] This is the view favored by David Lewis. See his *Counterfactuals*, p. 73.

regularities that are part of or consequences of the optimal system are precisely the regularities we would identify as nomic regularities (regularities of natural law).

In this section I want to argue that this is an illusion, that it is perfectly possible that systematic regularities fail to correspond to the laws there are. And in the next section we will see that even if, as a matter of fact, laws and systematic regularities coincide, this is not because they are the same things. The point will be that there are certain things we use laws for, in particular for *explaining*, for which we cannot use regularities, however systematic. And towards the end of this chapter I shall explain what sort of connection there is between laws and systematic regularities, and why therefore the systematic account looks as if it is materially adequate.

Earlier on we looked at the argument that probabilistic laws were a counter-example to minimalism. The idea was that there could be a divergence between the law and the corresponding regularity. Although the argument is not valid, I suggested that our intuitions should be against the minimalist. More generally, I believe that our intuitions suggest that there could be a divergence between the laws there are and the (systematic) regularities there are.

Returning to the simple regularity theory for a moment, one reason why this theory could not be right is that there could be a simple regularity that is merely accidental. The fact that there is regularity is a coincidence and nothing more. These instances are not tied together by a law that explains them all. To some extent this can be catered for by the systematic account. For, if the regularity is a coincidence, then we might expect to find that the events making up this regularity each have alternative explanations in terms of a variety of other laws (i.e. each one falls under some other regularity). If this is the case, then the systematic account would not count this regularity as a law. For it would add little in the way of strength, since the events that fall under it also fall under other systematic regularities, while adding it would detract from the simplicity of the overall system.

This raises a difficult issue for the systematic theorist, and indeed for minimalists more generally. Might not this way of excluding coincidental regularities get it wrong and exclude genuine laws while including accidents? It might be that there is a large number of laws each of which has only a small number of instances. These laws may throw up some extraordinary coincidences covering a large number of events. By the systematic account the coincidence would be counted as a law because it contributes significantly to the strength of the system, while at the same time the real laws are excluded because their work in systematizing the facts is made redundant by the accidental regularities. Figure 1.2 shows a world similar to the one discussed a few pages back. It has the same laws, but the additional facts mean that the accidental regularities now predominate over the genuine laws -- the optimal systematization would include the regularity that Fs are Gs but exclude the law that Hs are Gs. If such a world is possible, and I think our intuitions suggest it is, then the systematic regularity theory gets it wrong about which the laws are. In this way the systematic theory fails to be materially adequate.

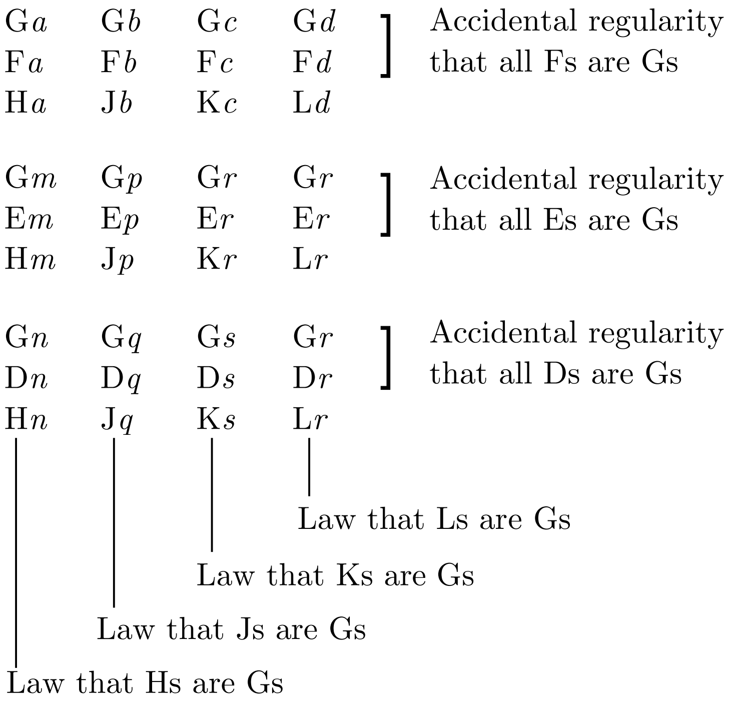

Figure 1.2

Laws, regularities, and explanation

I have mentioned that there are various things we want laws to do, which they are not up to doing if they are mere regularities. We have already looked at a related issue. One thing laws are supposed to do is support counterfactuals. Some opponents of the regularity theory think that this is sufficient to show that minimalism is mistaken, but I think that this is a difficult line to pursue (even if ultimately correct). More promising is to argue that laws cannot do the job of *explaining their instances* if they are merely regularities. Another issue is whether we can make sense of induction if laws are just regularities.

The key to understanding why a mere regularity cannot explain its instances is the principle that something cannot explain itself. The notion of explanation is one we will explore in detail in the next chapter. However, for the time being, I hope it can be agreed that for any fact F, F does not explain F -- the fact that it rained may explain the fact that the grass is damp, but it does not explain why it rained.

As a law explains its instances, it explains all of them and, therefore, it explains the fact that is the uniformity consisting of all of them. But a regularity cannot do this because, according to the above-stated principle, it cannot explain itself. To spell this out in more detail, let us imagine that there is a law that Fs are Gs. Let there in fact be only four Fs in the world, *a*, *b*, *c*, and *d*. So the generalization "all Fs are Gs" is equivalent to

(A) (Fa & Ga) & (Fb & Gb) & (Fc & Gc) & (Fd & Gd) & (nothing is F other than a, b, c, or d).

The conjunction of all the instances is

(B) (Fa & Ga) & (Fb & Gb) & (Fc Gc) & (Fd & Gd)

(A) and (B) are identical except for the last part of (A), which says that nothing is F other than a, b, c, or d. Let us call this bit (C)

(C) nothing is F other than a, b, c, or d

So (A) = (B) & (C).

Thus, to say (A) explains (B) is the same as saying (B) & (C) explains (B). But how can that be? For clearly (B) does not explain (B), as already discussed. So if (B) & (C) explains (B) it is by virtue of the fact that (C). However, (C) says that nothing is F other than a, b, c, or d. That other things are not F cannot contribute to explaining why these things are G. Therefore, nothing referred to in (A) can explain why it is the case that (B).

That we have considered a law with just four instances is immaterial. The argument can be extended to any finite number of instances. Many laws do have only a finite number of actual instances (biological laws for example). Nor does moving to an infinite number of instances change the nature of the argument. Even if it were thought that the infinite case might be different, then we could give essentially the same argument but, instead of considering the conjunction of instances, we could focus on just one instance. Why is a, which is an F, also a G? Could being told that all Fs are G explain this? To say that all Fs are Gs is to say if a is an F then a is a G, and all other Fs are Gs. The question of why a, which is an F is also a G is not explained by saying that if a is an F then a is a G, because that is what we want explained. On the other hand, facts about other Fs even all of them, simply do not impinge on this F.

It should be noted that this objection covers, in effect, all the forms of the regularity theory, not just the simple regularity theory. For more sophisticated versions operate by paring down the range of regularities eligible as laws and exclude those that fail some sort of test. But they still maintain that the essence of lawhood is to be a regularity, and the relation of explanation between law and instance is still, according to the regularity theorist, the relation between a regularity and the instance of the regularity. However sophisticated a regularity theory is, it cannot then escape this criticism. For instance, the deductive integratibility required by Lewis and Ramsey does not serve to provide any more unity to a law than is provided by a generalization. That the generalization is an axiom or consequence of the optimal axiomatic system does nothing to change the fact that it or the regularity it describes cannot explain its instances.

A case that seems to go against this may in fact be seen to prove to rule. I have an electric toaster that seems to have a fault. I take it back to the shop where I bought it, where I am told "They all do that, sir". This seems to explain my problem. Whether or not I am happy, at least I have had it explained why my toaster behaves the way it does. However, I think that this is an illusion. Being told that Emily's toaster does this, and Ned's and Ian's too does not really explain why my toaster does this. After all, if Emily wants to know why her toaster behaves that way, she is going to be told that Ned's does that too and Ian's and so does mine. So part of the explanation of why mine does this is that Emily's does and part of the explanation of why Emily's does this is that mine does. This looks slightly circular, and when we consider everyone asking why their toaster does this strange thing, we can see that the explanations we are all being given are completely circular.

So why does being told "They all do that" look like an explanation? The answer is that, although it is not itself an explanation, it *points* to an explanation. "They all do that" rules out as unlikely the possibilities that it is the way I have mistreated the toaster, or that it is a one-off fault. In the context of all the other toasters behaving this way, the best explanation of why my toaster does so is that some feature or by-product of the design or manufacturing process causes the toaster to do this. This is a genuine explanation, and it is because this is clearly suggested by the shopkeeper's remark that the remark appears to be explanatory. What this serves to show is precisely that regularities are not explanations of their instances. What explains the instances is something that explains the regularity, although the fact of the regularity may provide evidence that suggests what that explanation is.

Laws, regularities, and induction

If regularities do not explain their instances, then a question is raised about inductive arguments from observed instances to generalizations. The critic of minimalism, whom I shall call the fullblooded theorist says that laws explain their instances and that inferring a law from the observation of its instances is a case of *inference to the best explanation* -- e.g. we infer that there is a law of gravitation, because such a law is the best explanation of the observed behaviour of bodies (such as the orbits of the planets and the acceleration of falling objects). (Inference to the best explanation is discussed in Chapters 2 and 4). Because, as we have seen, the minimalist is unable to make sense of the idea of a law explaining its instances, the minimalist is also unable to employ this inference-to-the-best-explanation view of inductive inference.[14] For the minimalist, induction will in essence be a matter of finding observed regularities and extending them into the unobserved. So, while the minimalist's induction is of the form *all observed Fs are Gs therefore all Fs are Gs*, the full-blooded theorist's induction has an intermediate step: *all observed Fs are Gs, the best explanation of which is that there is a law that Fs are Gs, and therefore all Fs are Gs.*

Now recall the problem of spurious (accidental, contrived, single case, and trivial) regularities that faced the simple regularity theory. The systematic regularity theory solves this problem by showing that these do not play a part in our optimal gystematization. Note that this solution does not depend on saying that these regularities are in themselves different from genuine laws -- e.g. it does not say that the relationship between an accidental regularity and its instances is any different from the

[14] This view is promoted by David Armstrong in his *What is a law of nature?*, pp. 52-9.

relationship between a law and its instances. What, according to the minimalist, distinguishes a spurious regularity from a law is only its relations with other regularities. What this means is that a minimalist's law possesses no more intrinsic unity than does a spurious regularity.

Recall the definition of "grue" (see p. 18). Now define "emerire" thus:

X is an emerire = *either* X is an emerald and observed before midnight on 31 December 2000
or X is a sapphire and not observed before midnight on 31 December 2000.

On the assumption that, due to the laws of nature, all emeralds are green and all sapphires are blue, it follows that all emerires are grue. The idea here is not of a false generalization (such as emeralds are grue), but of a mongrel true generalization formed by splicing two halves of distinct true generalizations together. Now consider someone who has observed many emeralds up until midday on 31 December 2000. If their past experience makes them think that all emeralds are green, then they will induce that an emerald first observed tomorrow will be green; but if they hit upon the contrived (but real) regularity that all emerires are grue, then they will believe that a sapphire first observed tomorrow will be blue. As both generalizations are true, neither will lead this person astray. The issue here is not like Hume's or Goodman's problems -- how we know which of many possible generalizations is true. Instead we have two true generalizations, one of which we think is appropriate for induction and another which is not -- even though it is true that the sapphire is blue, we cannot know this just by looking at emeralds. What makes the difference? Whatever it is, it will have to be something that forges a link between past emeralds and tomorrow's emerald, a link that is lacking between the past emeralds and tomorrow's sapphire.

I argued, a couple of paragraphs back, that in the minimalist's view there is no intrinsic difference between a spurious regularity and a law in terms of their relations with their instances. And so the minimalist is unable to give a satisfactory answer to the question: What makes the difference between inducing with the emerald law and inducing with the emerire regularity? Being intrinsically only a regularity the law does not supply the required link between past emeralds and future emeralds; any link it does provide is just the same as the one provided by the emerire regularity that links emeralds and sapphires. (Incidentally, the case may be even worse for the systematic regularity theorist, as it does seem as if the emerire regularity should turn out to be a derived law because it is derivable from two other laws.)

The full-blooded view appears to have the advantage here. For, if laws are distinct from the regularities that they explain, then we can say what makes the difference between the emerald law and the emerire regularity. In the former case we have something that explains why we have seen only green emeralds and so is relevant to the next emerald, while the emerire regularity has no explanatory power. It is the explanatory role of laws that provides, unity to its instances -- they are,all explained by the same thing. The contrast between the minimalist and full-blooded views might be illustrated by this analogy: siblings may look very similar (the regularity), but the tie that bindi them is not this, but rather their being born of the same parents, which explains the regularity of their similar appearance.

3. A full-blooded view- nomic necessitation

The conclusion reached is this. A regularity cannot explain its instances in the way a law of nature ought to. This rules out regularity theories of lawhood. The same view is achieved from the reverse perspective. We cannot infer a regularity from its instances unless there is something stronger than the regularity itself binding the instances together.

The task now is to spell out what has hitherto been little more than a metaphor, i.e. there is something that binds the instances of a law together which is more than their being instances of a regularity, and a law provides a unity not provided by a regularity. The suggestion briefly canvassed above is that we must consider the law that Fs are Gs not as a regularity but as some sort of relation between the properties or *universals* Fness and Gness.

The term *universal* refers to properties and relations that, unlike particular things, can apply to more than one object. A typical universal may belong to more than one thing, at different places but at the same time, thus greenness is a property that many different things may have, possessing it simultaneously in many places. A *first-order universal* is a property of or relation among particular things; so greenness is a first-order universal. Other first-order universals would be, for example: being (made of magnesium, being a member of the species *Homo sapiens*, being combustible in air, and having a mass of 10 kg. A *second-order universal* is a property of or relation among first-order universals. The property of being a property of emeralds is a second-order universal since it is a property of a property of things. The first order universal greenness has thus the second-order property being a colour. *Being generous* is a first-order property that people can have. *Being a desirable trait* is a property that properties can have - for instance, being generous has the property of being desirable; so the latter is a second-order universal.

Consider the law that magnesium is combustible in air. According to the full-blooded view this law is a relation between the properties of being magnesium and being combustible in air. This is a relation of *natural necessity.* It is not just that whenever the one property is instantiated the other is instantiated, which would be the regularity view. Rather, necessitation is supposed to be stronger. The presence of the one property brings about the presence of the other. Necessitation is therefore a property (more accurately a relation) of properties. This is illustrated in Figure 1.3, where the three levels of particular things, properties of those things (first order universals) and relations among properties (second order universals) are shown. The arrow, which represents the law that emeralds are green, is to be interpreted as the relation of necessitation holding between the property of being an emerald and the property of being green.

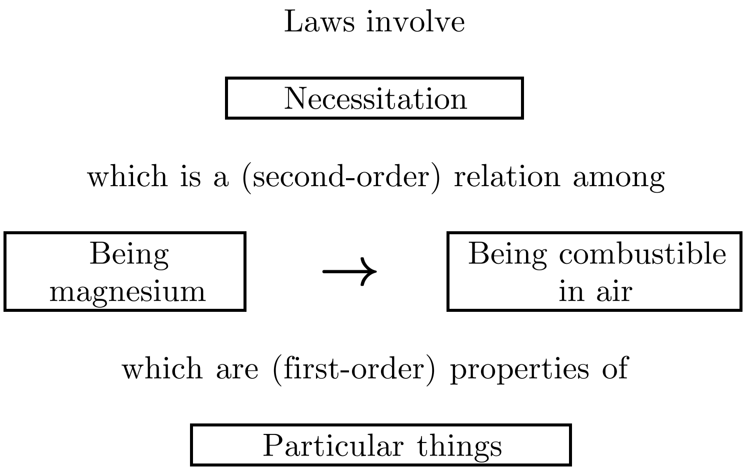

Figure 1.3

The advantages of this view are that it bypasses many of the problems facing the minimalist. Many of those problems involved accidental regularities or spurious "cooked-up" regularities. What were problems for the minimalist now simply disappear - necessitation among universals is quite a different matter from a regularity among things. The existence of a regularity is a much weaker fact than necessitation between universals. Two universals can coexist in precisely the same objects without one necessitating the other. We would have something like the diagram in Figure 1.3, but without the top layer and so without the arrow. This is the case in a purely accidental regularity. In an accidental regularity every particular has both of two universals, but the universals are not themselves related. This explains why not every regularity is a law. It also explains why every deterministic law entails a regularity. If there is a law that Fness necessitates Gness then every F will be G. This is because if x is F then the presence of Fness in x will bring about Gness in x, i.e. x will be G. So the existence of a deterministic law will bring about the corresponding regularity without the reverse holding.

This deals with accidental regularities. Some of the more spurious regularities, for instance those employing grue-type predicates, are even more easily dealt with, because the cooked-up predicates do not correspond to universals. There need be no genuine property of being grue, and hence none to be related by nomic necessitation.

Nomic necessitation among universals also makes more sense of explanation and induction. Regularities, we saw, cannot explain their instances in the way that laws are supposed to, as to do so would be circular. No such problem arises for the full-blooded view. The fact of a's being both F and G and the general fact of all Fs being Gs are both quite distinct from the fact of Fness necessitating Gness. The former are facts about individuals, while the latter is a fact about universals. And, as explained, while the latter fact entails the regularity that all Fs are Gs, it also goes beyond it, as it is not entailed by it. So reference to necessitation among universals as an explanation of particular facts or regularities is genuinely informative.

It also provides the unity among instances of a law that is necessary for induction to be possible. The thought is that if we see various different pieces of magnesium burn in air, we can surmise that the one property, being magnesium, necessitates the other, combustibility. If this conclusion is correct, then this relation will show itself in other pieces of magnesium. We induce not merely a resemblance of unobserved cases to those we have seen, but rather we induce a single fact, the relation of these two properties, which brings about the resemblance of unobserved to observed cases. In this way we should think of induction to facts about the future or other unobserved facts not as a one-stage process:

All observed Fs are Gs
∴all unobserved Fs are Gs

but instead as a two-stage process

All observed Fs are Gs
∴Fness necessitates Gness
∴all unobserved Fs are Gs

What is necessitation?

So far so good. The idea of necessitation between universals seems to do the job required of it better than the minimalist's regularities, be they systematic or otherwise. The objection the minimalist has the right to raise is this. At least we know what we mean by *regularity*. But what do we mean by *necessitation*? It is not something we can see. Even if in a God-like manner we could see all occurrences of everything throughout time and space, we would still not see any necessitation. The necessitation of Gness by Fness looks just like the regularity of Fs being Gs. This may not be enough to show that there is nothing more to a law than being a regularity. But it does put the onus on the full-blooded theorist. For without a satisfactory answer the suspicion may remain that perhaps "Fness necessitates Gness" just means "there is a law that Fs are Gs", in which case the full-blooded theorist will really be saying nothing at all.

What then can we say about this notion of "necessitation"? A leading full-blooded theorist, David Armstrong, lists the following properties of necessitation, which you will recognize from our discussion:

(1) If Fness necessitates Gness then this entails that everything which is F is also G.
(2) The reverse entailment does not follow. Instances of Fness may only coincidentally also be instances of Gness, without the being any necessitation.
(3) Since necessitation is a relation, it is a universal. Furthermore, since necessitation is a relation among universals, it is a second-order universal.
(4) Since necessitation is a universal it has instances. Its instances are cases of, for example, *a*'s being G because *a* is F (*a*'s being F necessitates *a*'s being G).

Is this enough to tell us what necessitation is? It might be thought that if we listed enough of the properties of necessitation, then that would be sufficient to isolate the relation we are talking about. If I told you that I was thinking about a relation R, and that R is a relation among people that is used to explain why people have the genetic properties they do, then you would be able to work out that what I had in mind was something like the relation "being an ancestor of". So perhaps the properties (l)-(4) are sufficient to isolate the relation of necessitation.

In fact, it turns out that these conditions do not succeed in distinguishing the nomic necessitation view from the regularity view of laws.[15] To see this let the relation RL be taken to hold between universals F and G precisely when it is a law that Fs are Gs, according to the Ramsey-Lewis systematic account. The requirements on RL are as follows:

The properties F and G are RL related precisely when:

(a) all Fs are Gs;
(b) the above is an axiom or theorem of that axiomatic system which captures the complete history of the universe and is the maximal combination of strength and simplicity.

The point is that the RL relation does everything that necessitation is supposed to do according to Armstrong's properties (1)-(4). Let us take (l)-(4) in turn: (1) If F and G are RL related, then, by (a), all Fs are also Gs. (2) Because of (b) the reverse entailment does not follow. (3) The RL relation is a relation among properties. Hence it is a second-order relation. (4) We can regard "*a*'s being G because *a* is F" as an instance of the RL relation -- when F and G are RL related, *a* is both F and G, and the capturing of the fact that *a* is G by the regularity that all Fs are Gs contributes to the systematization mentioned in (b). (I discuss this point at greater length in the next chapter.)

As the requirements (1)-(4) can be satisfied by taking laws to be a certain species of regularity, these requirements cannot gives us any insight into necessitation that accounts for the most important metaphysical features of laws, i.e. those which we discussed above:

(i) a law explains its instances;
(ii) particular facts can count as evidence for there being a law;
(iii) it is possible for systematic regularities to diverge from the laws that there are (i.e. there can be a lot of systematic regularity for which there are no corresponding laws).

We could add (i)-(iii) to (1)-(4), which would then exclude RL relations. But to do so would be to give up on trying to give an illuminating explication of the concept of necessitation. There would be no reason for someone who is doubtful about the idea of

[15] Points to this effect are made by Jeremy Butterfield in his review of Armstrong *What is a law of Nature?* in *Mind* **94**, 1985.

necessitation to think that such a relation really exists. Thus an alternative account of necessitation that satisfies (1)-(4) and (i)-(iii) is still required. This is what I will try to provide next.