1 Introduction

1.1 Scientific vs. Metaphysical Questions about Laws

Nature is full of regularities. Some are obvious: that bread nourishes, that heavy objects are attracted to the earth, and so on. Others, such as the patterns described by Schrödinger's equation or Einstein's field equations, are much harder to discover but more resilient. Indeed, we take them to be necessary: it's no accident that oppositely charged particles attract; under normal conditions, they *must* attract. Such regularities are associated with *laws of nature*.

Laws of nature are a subject of interest for both science and philosophy. However, scientists and philosophers focus on different aspects of laws. Scientists are primarily interested in *which* laws of nature there are. For example, physicists ask questions such as: What are the values of various physical constants? Is gravity Newtonian or relativistic? Should we accept classical or quantum mechanics? To clarify, let's take a closer look at the third question.

Classical Mechanics: The behavior of fundamental things is described by Newton's laws of motion.

Quantum Mechanics: The behavior of fundamental things is described by the laws of quantum mechanics, such as the Schrödinger equation.

These theories agree that laws describe the motions of objects. However, their laws differ, so they posit different regularities at the more fundamental levels of nature.

In contrast, philosophers are interested in metaphysical questions about the *nature* of laws. They ask questions such as: What kind of thing is a law? What makes some regularities lawlike and others accidental? Why does nature contain regularities in the first place?[1]

A thought experiment will help to motivate some different answers to these metaphysical questions.

Virtual Physicists: The departments of Physics and Computer Science have collaborated to produce a video game, *Virtual Physicists*. The objective is to explore a simulated environment, make observations, collaborate with other players, and together formulate a scientific theory that explains and predicts events in the virtual world. The game ends when the players discover the true final theory: the set of lawlike generalizations and boundary conditions in the program running the simulation.

[1] We'll see later (especially in Section 9) that philosophers are interested in other sorts of questions about laws, too.

Your task in the game is to answer *scientific* questions about laws. (It's not called "Virtual Philosophers"!) But suppose we ask why the virtual world contains regularities as opposed to irregularities. The answer is obvious: The regularities – namely, patterns on your screen – are *not* simply a rock-bottom stopping point for explanation, nor do they occur purely by chance; rather, they are *imposed* on the virtual environment by the program and ultimately by the programmers/developers.

The nature of laws in our world might be analogous or it might not be. Everyone agrees that laws and regularities are intimately related. However, the central metaphysical dispute about laws concerns the nature of this relationship. Compare:

Humeanism: The world is like a grand mosaic whose tiles could have been arranged any which way. Nothing is ultimately responsible for its regularities. They are just a basic, fundamental feature of our world. Statements of laws, then, are merely descriptions of the most significant regularities that happen to occur. Thus, our world is *not* analogous to the worlds in *Virtual Physicists*.[2]

Non-Humeanism: The world is like a grand mosaic, but it's *not* the case that its tiles could have been arranged any which way. Something imposes structure on it, analogously to the way in which the computer program *Virtual Physicists* imposes structure on its virtual environments. For example, some Non-Humeans posit a god as the enforcer of natural laws, some treat laws as primitives, and there are other options besides. But whatever the details, all Non-Humeans accept some sort of basic, fundamental necessity: something that governs, produces, or somehow constrains patterns of events in nature.[3]

This philosophical dispute concerns the nature of laws (or natural necessities) and the order of metaphysical explanation: Humeans hold that regularities are prior to laws (or any other sort of natural necessity), whereas Non-Humeans hold that laws (or some natural necessities) are prior to regularities. Generally speaking, Humeans and Non-Humeans can agree about *which* sentences are statements of laws – a matter left to scientists – while disagreeing about what makes them lawlike.

Of course, there are also disputes internal to Humeanism and Non-Humeanism. For example, I said in this section that Humeanism takes the most

[2] This view is so called because David Hume (1748/2000) famously argued that the concept of natural necessity had to be analyzed in terms of observed patterns. However, it's controversial whether Hume endorsed Humeanism as I've defined it (Strawson, 2015).

[3] The terms 'Non-Humean' and 'Anti-Humean' are used interchangeably in the literature. I prefer 'Non-Humean'. As I've said elsewhere (Hildebrand, 2020b), although I'm not a Humean, I'm not *anti*-Humean!

significant regularities to be laws, but the notion of significance is open to interpretation. This leads to different versions of Humeanism. Similarly, there are many versions of Non-Humeanism, differing with respect to the primitive necessities they posit. Some appeal to God, some treat laws themselves as primitive, some invoke special sorts of properties, and so on. We'll explore narrower versions of both theories in later sections.

1.2 Philosophical Method

As noted, Humeanism and Non-Humeanism are typically formulated in an effort to accommodate our best scientific theories of laws. That seems to suggest that our metaphysical theories will be empirically equivalent. How, then, are we to choose one?

In comparison, it might seem easy to settle scientific disputes about laws. Precisely because scientific disputes about laws are disputes about *which* laws (and thus *which* regularities) there are, we can settle them by careful observation. However, even the question "How do we choose a scientific theory?" is difficult to answer. Consider the apparent position of stars during the famous Eddington experiment in 1919. Newtonian theories of gravity (of the day) and Einstein's theory of gravity made different predictions about where stars would appear during a solar eclipse. Eddington's team looked at the stars and found the predictions of Einstein's theory to be more accurate. Did everyone immediately accept Einstein's theory over Newton's? No. Newtonian theories could be modified to make predictions matching the observations of the experiment. Thus, Newtonian theories could be made empirically equivalent to Einstein's theory, at least for known observations of the day. How, then, can scientists justify a choice among empirically equivalent theories?

The answer, in short, is that scientists can and do appeal to *nonempirical theoretical virtues*. For example, Einstein's theory seems simpler and more elegant than its empirically equivalent Newtonian competitors. Generally speaking, *nonempirical criteria for theory choice* are required to solve problems of empirical underdetermination.

Analogously, philosophers appeal to similar criteria for theory choice when doing metaphysics.[4] I'll provide a brief sketch of some important criteria, but it includes a strong disclaimer. In this section, I won't try to justify them or say how they fit together to form a big-picture method for metaphysics. I'll say a bit more about this in Sections 8 and 9, after we've seen how the criteria are employed in philosophical practice.

[4] For defenses of metaphysics that proceed along these sorts of lines, see Paul (2012), Sider (2011, 11–15), and Tahko (2015).

For starters, a theory of laws should be intelligible; it shouldn't involve primitive concepts that we cannot understand. For example, Non-Humeans need to explain how we come to possess the concept of primitive necessity that features in their theories.

Theories should be simple and/or parsimonious. They shouldn't posit too many (kinds of) entities when fewer (kinds of) entities will do.

Theories of laws should cohere with scientific practice concerning laws and align with the ordinary concept of laws. Every philosophical analysis must start somewhere. If our account of laws is to be an account of *laws* – the things scientists talk about when they use the term – the account should classify the things we call "laws" as laws, and it should accommodate at least some of our normal intuitions about the work laws are supposed to do. For example, a metaphysical account of laws should make sense of the fact that scientists appeal to laws to explain observations and make predictions.

Relatedly, it would be nice if a theory of laws could explain *why* there are regularities in nature. But providing such explanations can be difficult, in part because some explanations merely push the problem back a level. If we posit a new entity to explain a regularity, we can ask why that new entity is the way that it is. We don't want to posit "turtles all the way down."

Finally, an account of lawhood should fit within our broader metaphysics. A metaphysical naturalist – one who thinks that nothing exists beyond the world of spacetime – would find it costly to endorse a theory of laws that posits a god. But if your metaphysics of laws fits seamlessly within your broader metaphysical commitments, that would be a mark in its favor.

In sum, there are many criteria relevant to the selection of a philosophical theory of laws. We'll examine them more carefully in due course, but this overview should provide a sufficient foundation to get started.

1.3 Why Care?

To conclude this section, I'll briefly mention some ways in which philosophical reflection on the nature of laws might be valuable.

First, science is one of the best things humans have created. That makes it worthy of philosophical reflection, if only for our own curiosity and enjoyment. More importantly, because science is valuable to us, it's worth trying to understand how scientific theories are to be interpreted; it's worth asking why science has the virtues it has and why it's possible for beings like us. Answers to such questions require us to say at least something about the nature of laws.

Second, the metaphysics of laws intersects with other philosophical issues. Generally speaking, an account of the nature of the world is relevant to accounts

of how we ought to form our beliefs and live our lives. Only one connection of this sort is discussed in this Element (Sections 3.4, 8, and 9), but it's important: Laws are essential for *induction*, the form of inference – crucial to both science and ordinary life – in which we make predictions about the unobserved on the basis of the observed.

Third, studying philosophy can teach us humility. A better understanding of what laws might be helps us to understand our limitations and the strengths and weaknesses of scientific methods. Occasionally, a prominent scientist will make disparaging remarks about philosophy (I won't name names). I think this reflects a failure to understand the nature of scientific or philosophical projects, as well as the relations between them. Historically, science and philosophy have been viewed as complementary by leading scientists (e.g., Newton and Einstein).

Unfortunately, I don't have much space to discuss questions of value in this Element. However, my hope is that exposure to philosophical questions themselves will help readers to see the value for themselves. So, let's dive in. We'll begin with a careful examination of Humeanism.

2 Humeanism

At the fundamental level, Humeanism doesn't posit any primitive necessity or other modally robust primitives; it just posits non-modal events in spacetime – i.e., *the Humean mosaic*. That's all. Its ontology (the entities that exist according to the theory) is economical, and its conceptual primitives (the primitive predicates required to express the theory) are easily grasped. In contrast, Non-Humeans posit primitive necessities in addition to events in spacetime. But why worry about exotic metaphysical necessities if we don't have to? In this section, we'll examine some different versions of Humeanism.

To begin, here was our initial statement of Humeanism:

> **Humeanism:** The world is like a grand mosaic whose tiles could have been arranged any which way. Nothing is ultimately responsible for its regularities. They are just a basic, fundamental feature of our world. Statements of laws, then, are merely descriptions of the most significant regularities that happen to occur…

There are two distinct claims here: first, that nature consists of a "Humean mosaic"; second, that laws are, or reduce to, regularities in the Humean mosaic. We'll spend the bulk of our efforts examining the second claim, but I'll say a little more about the Humean mosaic before we begin.

2.1 The Humean Mosaic

I have provided a *negative* characterization of the Humean mosaic, in terms of its lack of primitive necessities. However, many Humeans would prefer a *positive* characterization – one that tells us what the mosaic is like without invoking any modal language at all. As we'll see in Section 5, some Humeans maintain that we don't understand the modal language invoked in Non-Humean theories. It would be a problem if their own theory required the same language.

Here is an example of a positive characterization due to David Lewis:

> We have geometry: a system of external relations of spatiotemporal distance between points. … And at those points we have local qualities: perfectly natural intrinsic properties which need nothing bigger than a point at which to be instantiated. For short: we have an arrangement of qualities. And that is all. (Lewis, 1986b, ix)

Lewis doesn't invoke any modal language, but he does tell us more about the fundamental properties instantiated by points in the mosaic. They are *perfectly natural*, *intrinsic*, *point-sized*, and so on. Don't worry if you're not familiar with these concepts. What matters for present purposes is just that these concepts aren't supposed to have primitive modal character. The upshot is that it does seem possible to provide a positive characterization of the Humean mosaic.

However, the details are controversial. For example, Lewis's characterization seems unable to accommodate entangled states of quantum mechanics, since such states seem to involve connections between distinct points of the Humean mosaic (Maudlin, 2007). If that's right, Lewis's account of the Humean mosaic is incompatible with one of our best scientific theories. There are various proposals for characterizing the Humean mosaic that seek to avoid this problem, but this is largely an internal dispute among Humeans and unfortunately there isn't space to consider them here.[5]

What matters for our purposes is that Humeans are united in their rejection of primitive modality. Thus, we'll stick with a negative characterization of the Humean mosaic. This gives Humeans freedom to adjust their positive characterization of the mosaic to match our best scientific theories, and it will suit us just fine in our attempt to understand the major differences between Humeanism and Non-Humeanism. Namely, it suffices to make sense of one of the primary motivations for seeking a Humean theory of laws: Its ontology is economical and its conceptual resources (whatever they turn out to be) will be easier to understand insofar as they do not invoke primitive modality. Let's shift our attention to the claim that laws are analyzed in terms of regularities.

[5] See, for example, Loewer (1996), Earman and Roberts (2005a), and Bhogal and Perry (2017).

2.2 The Naïve Regularity Theory

Suppose the world is a Humean mosaic. According to Humeanism, laws are, or reduce to, regularities. But not just any regularities will do. They have to be "significant." What does that mean, and why should we restrict lawhood to significant regularities in the first place?

Well, here are some statements that do *not* seem lawlike:

(1) Most elementary particles have mass.
(2) All electrons have positive charge.
(3) The result of adding 2 and 2 is always 4.
(4) On January 1, 2022, all the books on my desk were written by David Armstrong.

Each statement describes a regularity. Why, then, aren't they statements of laws? Statement (1) isn't universal, but we typically think of (fundamental) laws as holding at all times and places. Statement (2) is false, which seems to disqualify it. Statement (3) is necessarily true. However, mathematical necessity seems stronger than natural necessity. We can't imagine worlds in which $2+2 \neq 4$, but we *can* imagine worlds with different laws. Finally, statement (4) is tied to specific people, places, and times. The laws of respectable scientific theories aren't – at least not at more fundamental levels of science.

Each statement suggests a condition on lawhood, giving rise to the following Humean theory:

The Naïve Regularity Theory: *L* is a *statement of a law of nature* if and only if *L* is (i) universally quantified, (ii) true, (iii) contingent, and (iv) contains only nonlocal empirical predicates apart from logical connectives and quantifiers.[6]

Notice that we have replaced the vague term 'significant' with a more precise set of conditions.

The Naïve Regularity Theory has some attractive features. Because its language is precise and it avoids reference to mysterious primitives, it fares well with respect to the virtues of conceptual clarity and ontological economy. Also, its classification of laws seems accurate. According to this theory, none of (1)–(4) are statements of laws, whereas paradigm statements of law such as the Schrödinger equation *are*.

[6] This statement is borrowed with slight modifications from Armstrong (1983, 12), who borrowed it from Molnar (1969, 79).

2.3 Extensional Problems for the Naïve Regularity Theory

Unfortunately, the Naïve Regularity Theory has a serious problem: It is *extensionally inadequate*, which is to say that it fails to correctly classify statements of laws. In other words, the expressions on either side of the 'if and only if' pick out different classes. We'll illustrate with two counterexamples.

The following two statements suggest that conditions (i)–(iv) are *insufficient* for lawhood.

(1) All solid spheres of Uranium-235 have radius < 1 km.
(2) All solid spheres of gold have radius < 1 km.

Both statements meet all four of the Naïve Regularity Theory's conditions on lawhood: they are universally quantified, true, contingent, and avoid local predicates. Thus, the theory classifies both as statements of laws. But, intuitively, only (1) is a statement of law. There *couldn't* be a large sphere of Uranium-235, because it would exceed the critical mass. In contrast, (2) seems *accidental*. There are no large spheres of gold, but there could have been. Nothing about the nature of gold would prevent the formation of a large gold sphere. Thus, the theory implies that (2) is a statement of law when in fact it isn't.

This is an example of *the problem of accidental regularities*, so called because the theory seems to classify accidents such as (2) as laws. It is easy to generate other examples of this sort. Have a try yourself.[7]

Let's turn to a famous case that suggests that conditions (i)–(iv) are *unnecessary* for lawhood.

> All the fruit in Smith's garden at any time are apples. When one attempts to take an orange into the garden, it turns into an elephant. Bananas so treated become apples as they cross the boundary, while pears are resisted by a force that cannot be overcome. Cherry trees planted in the garden bear apples, or they bear nothing at all. If all these things were true, there would be a very strong case for its being a law that all the fruit in Smith's garden are apples. And this case would be in no way undermined if it were found that no other gardens, however similar to Smith's in all other respects, exhibited behavior of the sort just described. (Tooley, 1977, 686)

The expression "All the fruit in Smith's garden are apples" makes reference to Smith, so it violates condition (iv), according to which statements of laws must involve only nonlocal predicates. Nevertheless, it seems lawlike in Smith's world.[8]

[7] For more careful discussion of such objections, see Armstrong (1983, chap. 2).
[8] See Armstrong (1983, 26) for a less fanciful case and Lange (1995) and Dosanjh (2021) for further discussion of laws involving individuals.

The point of these examples is that our concept of *law* and our concept of *a regularity meeting conditions (i)–(iv)* seem to be different. As a result, the Naïve Regularity Theory does not capture our ordinary concept of law.

A comment on philosophical method: In describing these objections, I have appealed to intuitions, or the way that things seem, or to our ordinary way of thinking about the role of laws in science. This is standard practice in philosophy, though it remains controversial. I'll say more about these methodological issues in Sections 3 and 8, but for now I'll just mention two possible responses to these sorts of objections. One option is bite the bullet. A proponent of the Naïve Regularity Theory could simply accept that it *is* a law that all solid spheres of gold have radius < 1 km, and that there *could not be* any law involving Smith's garden. However, these are but two of many putative counterexamples (Armstrong, 1983, chap. 2), so applying this strategy across the board would require a radical revision to our concept of law. The more popular strategy is to revise the analysis of laws. Over the next few subsections, we'll examine a version of Humeanism that does a better job of classifying laws and accidents.

2.4 The Humean Best Systems Account

Here is a more popular account of how to analyze laws in terms of the Humean mosaic.

Humean Best Systems Account (BSA): Statements of laws are contingent generalizations in the best systematization of the Humean mosaic.

The best systematization of a mosaic is, roughly speaking, an efficient summary of its events. To make better sense of this, we need to clarify the concept of a *systematization* and we need to say something about which is *best*.

A *systematization* of a mosaic is just a set of sentences about that mosaic.[9] Consider a deterministic world in which Newton's laws hold. Suppose we provide a complete description of the world, point by point, analogous to a bitmap image file in which each pixel is assigned a color. We'll call this set of true singular sentences *HM*, short for 'Humean mosaic'. Now consider the set of all sentences entailed by *HM*.[10] This new set is the *deductive closure* of *HM*, so we'll label it *DC*. Now consider a set whose members are solely (*i*) statements describing the initial conditions and (*ii*) Newton's laws. We'll call this set *PL* after Pierre-Simon Laplace (1814/1999), who suggested that we imagine a god

[9] They are usually stipulated to be true sentences but not always (Braddon-Mitchell, 2001).
[10] Talk of one set entailing another is shorthand for saying that the members of one set entail the members of the other.