

Laws of Nature and Their Modal Surface Structure

In this and the following chapters, I advocate a practice-oriented approach to questions in the metaphysics of science. I take metaphysics to study – inter alia – the most general features of reality, among them the issues covered here: laws, causation, reduction and foundationalism. My approach starts with the *role* played by laws, causation and reduction in scientific practice. The best explanation of the success of the scientific practice we have, I argue, requires making a number of metaphysical assumptions about the structure of reality. Thus, the purpose of this and the following chapters is to examine which metaphysical assumptions we need to make in order to understand the role that laws of nature, causation and reduction play in scientific practice. In this context, our practices of explanation, confirmation, manipulation and prediction play the role of the explananda in an inference to the best explanation.

I will assume that for an explanation to qualify as the best explanation it should be *minimal*: it should contain no assumption that does not do any work in explaining scientific practice. One may worry that the conclusions we can draw from a minimality constraint thus defined depend on where we start our investigation. If we start by looking at practice P_1 and move to an investigation of P_2 and P_3 , it may turn out that, given assumptions A_1 and A_2 , postulating A_3 does not do any additional explanatory work. However, had we started with an analysis of P_3 , A_1 might have turned out to be explanatorily irrelevant. As a rejoinder I would like to point out that our starting point is non-arbitrary. Our practice of confirmation, explanation, etc. in terms of laws of nature – as I will argue in later chapters – is presupposed in causal reasoning as well as in our reductive practices but not vice versa. Thus, we cannot avoid looking at the role laws play in confirmation, explanation, etc. outside causal or reductive contexts – and how to account for it. Analysing the role laws of nature play in scientific practice is thus the natural starting point of a minimal metaphysics of scientific practice.

We will see how far one can get with scientific practice as the main epistemic source for metaphysical arguments. Further sources that traditionally play a significant role in metaphysical theorising, such as appeal to intuitions or a preference for desert landscapes, will only be admitted if there is an argument as to why such intuitions or preferences should be considered to be truth conducive.¹

Chapters 1 and 2 are both devoted to a characterisation of laws of nature. While the first chapter focuses on how to best reconstruct law statements and on modal aspects of laws, Chapter 2 will be concerned with the practice of hedging laws with *ceteris paribus* clauses. A full account of what is the best account of laws will thus have to wait until the end of Chapter 2.

I will start by arguing that the practices of explanation, confirmation, manipulation and prediction require a particular reading of the law statements² involved as invoking two different kinds of generalisations – internal generalisations and external generalisations (Section 1.1). Having reconstructed what law statements say in the light of how they are used in confirmation, explanation, etc., I will then explain why law statements thus reconstructed can successfully play the role they play. I argue that law statements make claims about *systems* – more precisely, attributing multi-track properties to systems (Section 1.2). Furthermore, I will analyse the modal surface structure of law statements. It is part of my approach to eschew questions as to the origin of the modal features that are delineated by laws. The question whether or not the modal surface structure is reducible to non-modal facts may be an interesting question on its own, but answers to this question typically do not do any work in explaining the success of the scientific practice we have. Still, my account will comprise three claims about the modality of laws. First, law statements attribute a space of possible states to systems. Second, laws constrain the temporal development of systems by virtue of what I will call law equations. Third, the laws' inviolability or natural necessity can be explicated in terms of the fact that they are invariant with respect to a number of different kinds of circumstances (Section 1.3).

¹ In Chapter 3, I argue that we do have an argument to consider causal intuitions – as opposed to intuitions of, say, law-governing or metaphysical fundamentality – to be by and large truth conducive.

² Although a number of people have reservations about using the expression 'law' or 'law statement' (e.g., Woodward and Hitchcock 2003, *iff*), I will use 'law statement' in order to be able to distinguish law statements from other kinds of generalisations that are involved in scientific practice (see, e.g., Section 1.1.3).

1.1 Law Statements and the Role of Different Kinds of Generalisations

1.1.1 Law Statements

Let me start with Galileo's law. It may be thought that Galileo's law is simply identical to the following equation:

$$s = \frac{1}{2} g t^2 \quad (1.1)$$

(where s is distance covered, t is time and g is a constant). That seems wrong to me. It is fairly uncontroversial to take laws or law statements to be those (maybe complex) generalisations that play a role in extrapolation, confirmation, explanation and other aspects of scientific practice. With this characterisation of a law statement as a starting point, we can immediately infer the following consequence: if a law statement is what is confirmed or disconfirmed in trials (or used in the contexts of explanation, prediction or manipulation), an equation on its own cannot be an example of a law (or a law statement – in what follows I will use these two terms synonymously). As a matter of fact, nobody takes Galileo's law to be disconfirmed by balls uniformly rolling on a horizontal plane or by stones lying on the ground, both of which fail to satisfy Eq. (1.1). What is missing is a claim about *the kinds of systems* that are meant to be represented by the equation. Galileo's law is not simply a mathematical equation. Nor does it suffice to add that t represents time and s the path taken by an arbitrary object. Galileo's law is the claim that the behaviour of a particular class of systems can be represented by this equation. A full statement of Galileo's law might thus be something like the following:

Free-falling bodies behave according to the equation $s = \frac{1}{2} g t^2$.

Similarly, $F = ma$ is merely a mathematical equation. It becomes a law statement once it is asserted that this equation is meant to represent the behaviour of physical systems; indeed, of all physical systems whatsoever. And again, the Schrödinger equation with the Coulomb potential on its own does not qualify as a law statement; that is, it is not what we confirm or disconfirm. By contrast, the claim '*Hydrogen atoms* behave according to the Schrödinger equation with the Coulomb potential' is a law statement.

The fact that equations such as $s = \frac{1}{2} g t^2$ come with a domain of systems³ for which they are meant to be relevant has been noted by others, e.g.,

³ Nothing hinges on the term 'system' – 'object' or 'thing' would be fine too. I say a little bit more about systems in Section 1.2.

within the semantic account of theories. Thus, Bas van Fraassen, referring to Ronald Giere, defines a *theory* (not a law⁴) as consisting of:

- (a) the *theoretical definition*, which defines a certain class of systems; and
- (b) a *theoretical hypothesis*, which asserts that certain (sorts of) real systems are among (or related in some way to) members of that class (van Fraassen 1989, 222).

A preliminary general characterisation of law statements might thus be the following:

(A) All systems of a certain kind K behave according to Σ .

Here ' Σ ' – the law predicate – typically stands for an equation or a set of equations. The expression 'of a certain kind K' may refer to all physical systems whatsoever, as it does in the case of Newton's second law or in the case of the bare Schrödinger equation. Or it might refer to a more circumscribed class of systems such as free-falling bodies or hydrogen atoms, thus giving rise to so-called *system laws*.⁵ It is important to note that the behaviour attributed to the systems in question is in general complex and relational. In the case of free-falling bodies, the length of the path and the time taken are related, not only for actual values of the variables but for all possible (or some restricted domain of) values. Taking 'All ravens are black' as a paradigm for law statements ignores the complex structure usually attributed to systems.

Another example that illustrates the structure of law statements is Euclidean geometry. Euclidean geometry on its own is a mathematical theory without any empirical import. We get an empirically testable claim (a law) if we take a certain class of systems (space-times) to be adequately characterised in terms of Euclidean geometry.⁶

⁴ In fact, Giere and van Fraassen deny that there are laws of nature (van Fraassen 1989, 183ff). According to my reconstruction, what Giere and van Fraassen call a 'theory' should be taken to be a law statement.

⁵ One might worry about the exact characterisation of the system to which Σ is attributed. The worry is that one needs Σ to individuate the systems in question. That, of course, would make the law statement an analytical truth and thus devoid of empirical content. It has to be assumed that the relevant class has been individuated antecedently, for example in terms of experimental procedures ('free falling bodies') or by other means that do not depend on Σ . This is a thorny issue that I will not go into in this book.

⁶ This point was famously observed by Einstein: 'As far as the laws of mathematics refer to reality, they are not certain; and as far as they are certain, they do not refer to reality.' In fact, Einstein – in his paper 'Geometry and Experience' (Einstein 1921) – suggests a view of laws or theories pretty close to the one suggested here.

Law statements as just characterised play a prominent role not only in physics but also in other disciplines. Thus, the Lotka–Volterra equations describe the temporal development of a biological system consisting of two populations of different species, one being a predator and the other prey. The relevant equations for prey and predator populations are (1) $dx/dt = x(a - by)$ and (2) $dy/dt = -y(c - gx)$, where x represents the number of prey and y the number of predators and a , b , c and g are constants. Again, we can distinguish between the system to which the equations apply, on the one hand, and the equations or the description of the behaviour, on the other.

Even in cases in which the behaviour in question is not represented mathematically, it is possible to distinguish the behaviour from the systems to which it is attributed. Thus, according to Schmalhausen's law, a population at the extreme limit of its tolerance in any one aspect is more vulnerable to small differences in any other aspect. On the one hand, we have populations (the systems); on the other, we have a qualitative description of what may happen to the population (the behaviour).

1.1.2 Internal and External Generalisations

Characterising law statements in terms of (A) allows me to draw attention to an important distinction between different kinds of generalisations. Take the example of Galileo's law,

Free-falling bodies behave according to the equation $s = \frac{1}{2} gt^2$.

Even though there are often no explicit quantifiers, law statements usually involve at least two different kinds of generalisations (as will be illustrated by an example in Section 1.1.3). In Galileo's law we can distinguish one form of generalisation that quantifies over systems (for all x that are falling bodies). This quantification specifies the *objects* (or systems) to which a certain kind of behaviour is attributed. Besides generalisations that pertain to objects or systems, there are *system-internal generalisations*. These generalisations concern the values of the variables that appear in the equation. When we claim that a system behaves according to the equation $s = \frac{1}{2} gt^2$, what is implied is that for every value of t the path s that the body has fallen is determined by $s = \frac{1}{2} gt^2$.

We can therefore distinguish two kinds of generalisations (see Scheibe 1991a)⁷:

⁷ Hitchcock and Woodward (2003, 189) draw attention to this distinction, albeit in different terms, when they remark with respect to explanation that 'the nomothetic approach has focused on

- (1) *System-internal generalisations*: Generalisations concerning the values of variables. For instance, in the case of Galileo's law, the system-internal generalisation is that the equation holds for all values of the variable t (or at least for all values within a certain range).
- (2) *System-external generalisations*: Generalisations concerning different systems such that the equation pertains to all systems of a certain kind (e.g., free-falling bodies).

In the case of the Lotka–Volterra equations, the internal generalisations concern the variables x (number of prey) and y (number of predators), while the external generalisations concern ecological systems consisting of prey and predator populations.

In our preliminary law characterisation (A), the system-external generalisation ('All systems of a certain kind') is explicitly mentioned while the system-internal generalisations are implicit in Σ . One reason why internal generalisations are not made explicit may be the fact that usually more than one internal generalisation is allowed by the law equation and it is the context that determines which of those are relevant for the characterisation of a particular phenomenon. To be a bit more specific, law equations (in contrast to the structural equations discussed in the causation literature) are not in general asymmetric. As a consequence, a law equation, such as $pV = \nu RT$, allows us to infer not only that once the values for p and V are given those of T are determined but also that the values for p and T determine those for V , etc. The law equation thus implies at least three different internal generalisations. Which of those is relevant for a particular situation may depend on the quantities on which we want to intervene or on other features determined by the context. Similarly, when we claim that a system behaves according to the equation $s = \frac{1}{2}gt^2$, what is implied is not only that for every value of t the path s is determined by $s = \frac{1}{2}gt^2$ but also that for every path s , the time t that the body has taken to fall is determined by $t = \sqrt{(2s/g)}$. The law statement allows us to assert both generalisations.

The distinction between internal and external generalisations goes hand in hand with a distinction between different kinds of counterfactuals noted by Woodward and Hitchcock. First, we have what Woodward and Hitchcock (2003, 20) call 'other object counterfactuals'. Examples of other object counterfactuals include 'If b had been a raven it would have been black' or 'If b had been an ideal gas, it would have behaved in accordance

a particular kind of generality: generality with respect to objects or systems other than the one whose properties are being explained'. By contrast, their own account of explanation relies on generalisations that pertain to the values of variables.

with the equation $pV = \nu RT$.' These counterfactuals presuppose external generalisations; traditional accounts of laws as external generalisations stress the fact that laws support other object counterfactuals.⁸

Woodward and Hitchcock contrast other object counterfactuals with *same object counterfactuals*. These pertain to particular systems, such as in the statement 'If the ideal gas in question had had a volume $V = V_o$ and a pressure $p = p_o$ its temperature would have been $T = T_o$.' Same object counterfactuals presuppose internal generalisations. Woodward and Hitchcock argue that it is same object counterfactuals that are relevant for scientific explanation. I largely agree, though I do not accept the interventionist account of the truth conditions for these counterfactuals. In Chapter 2 (Section 2.4.3), we will encounter a third kind of counterfactual connected with law statements.

The fact that law statements come with internal generalisations is a first hint at the complexity of what law statements assert. Take the ideal gas law as an example. The law statement is highly complex because it is a functional law. It implies an infinite number of statements of the form 'If the value of p and the value of V of a particular gas had been such and such then the temperature T would have been thus and so.' Thus, the behaviour or properties that law statements attribute to systems are typically infinitely *multi-track* (I will examine this issue in more detail in Section 2.4.3). Note that the infinity of implied statements does not preclude that the ideal gas law can be stated in finite terms.

The Schrödinger equation provides another illustration of the complexity of law statements. When we claim that hydrogen atoms can be characterised in terms of the Schrödinger equation with the Coulomb potential, the Σ in our canonical statement 'All systems of a certain kind K behave according to Σ ' comprises the conceptual apparatus of quantum mechanics. So, when we say that hydrogen atoms behave according to the Schrödinger equation with the Coulomb potential, we are saying that they behave according to quantum mechanics in which the Schrödinger equation is concretised via the Coulomb potential. The essential point is that law statements attribute a behaviour to systems already identified as being of a certain kind by invoking law predicates that typically involve a highly complex mathematical apparatus. This complexity becomes invisible if we are operating with examples of the 'All ravens are black' sort.

⁸ The role of other object counterfactuals is somewhat controversial because they may be thought to involve metaphysical impossibilities in the antecedent (see Tan 2019 for discussion). That debate, however, is not relevant for the purposes of this chapter.

Examples of the latter kind are misleading because they suggest that law statements can be analysed solely on the basis of external generalisations. This assumption shaped much of the debate about laws of nature in the twentieth century.

1.1.3 Excursus: The Role of Internal and External Generalisations in Standard Explanation

The distinction between internal and external generalisations will prove fruitful in later sections of this chapter. It also helps to understand how certain simple cases of explanations work. Consider an example of what may be called a ‘standard explanation’ (for a discussion of this example see Skow 2016, 75ff.).

An explanation or answer to the question ‘Why did that rock hit the ground at a speed of 4.4 m/s?’ might consist in the statement ‘It hit the ground at that speed because it was dropped from a height of one metre.’ If the further question arises as to why the one explains the other, an answer will refer to the equation $s = \frac{1}{2} gt^2$.

How exactly does this explanation work? The explanandum in this case is the velocity of a particular rock (more generally, the state of a system). The explanans mentions the height and the equation $s = \frac{1}{2} gt^2$. So, we have the general structure that in a standard explanation we explain the state of a system in terms of initial conditions and a (dynamic) law equation.

The explanation points out how the speed of the rock – in virtue of the equation $s = \frac{1}{2} gt^2$ – depends on the height from which the stone was dropped; in doing so, the explanation appeals to an internal generalisation. By contrast, the fact that Galileo’s law *covers* the rock, as implied by the external generalisation, is a *presupposition* for the explanation to work. Without knowing that the law covers the rock it would not make sense to explain the velocity in terms of this equation.

More generally, in simple standard explanations like the one just described, the external generalisation claims that the law predicate Σ is relevant for the explanandum. This appeal to an external generalisation is not a part of the explanans but rather, a presupposition of the explanation. By contrast, internal generalisations typically figure explicitly in the explanans since they provide information about how one quantity depends on other quantities.⁹

⁹ How does this account relate to Skow’s distinction between levels of reasons why? Skow (2016, chapter 4) distinguishes a first-level question ‘Why did the rock hit the ground at 4.4 m/s?’ (Answer/

In the 1960s, Scriven and Hempel debated the role of laws in explanation. One of the contentious issues was whether laws explicitly figure in the explanans or whether they provide a ‘role justifying ground’ ‘roughly showing that the explanans is relevant for the explanandum’ (Salmon 1984, 17 fn6). By relying on the distinction between external and internal generalisations we see that law statements play both roles. The external generalisation does not figure explicitly in our simple standard explanation; it provides a ‘role justifying ground’. By contrast, the internal generalisation is explicitly appealed to in the explanation.

To sum up: what I do hope to have shown is that not only external but also internal generalisations can play a role in explanation and that the distinction is thus significant for understanding scientific practice.

1.2 Systems

Up to this point I have reconstructed what we should assume law statements to say given their role in confirmation, explanation, etc. I will now turn to an explanation of why law statements thus reconstructed can successfully play the role they play.

In Section 1.1.1, I argued that law statements attribute a certain kind of behaviour to systems: All physical systems of a certain kind K behave according to Σ ? A pretty straightforward explanation of why we can successfully work with such law statements is the assumption that systems that display the relevant behaviour do in fact exist.

To talk of ‘systems’ and ‘behaviours’ is very common in the sciences. Systems come in all kinds of sizes, including space-times, economies, interacting predator–prey populations, cells, gases and hydrogen atoms. In contrast to the notion of substance (on some interpretations), systems do not carry with them notions of fundamentality or indivisibility; systems

reason: it was dropped from 1 m) from a second-level question ‘Why is being dropped from 1 m a reason for it having the speed of 4.4 m/s?’ (Answer/reason: $s = \frac{1}{2}gt^2$, etc.) I have not made this distinction between these different levels of why questions, but it would do no harm to my argument if I did. What is important for me is to distinguish a third-level question that Skow does not consider: ‘Why is it that the equation $s = \frac{1}{2}gt^2$ is relevant for the second-level answer?’ An answer to the third-level question has to appeal to the external generalisation, while an answer to the second-level question appeals to the internal generalisation. Thus, for my purposes the issue of whether first and second level need to be distinguished is not relevant; what is relevant, however, is that the third level can be distinguished.

may be fairly complex, they might be constituted out of subsystems, and so on.¹⁰

In the case of physical systems, we can distinguish different aspects of their behaviour. Some quantities of a physical system are constant; others vary with time. For instance, for a single classical particle we can distinguish position and momentum as changing quantities, whereas mass remains constant. The values of the varying quantities at a particular time are called the ‘state’ of the physical system at this time. However, the constants and the state of a system do not determine the complete behaviour of the system. We also have equations that describe the connections between the various quantities involved and in particular how the state of the system develops over time: its *temporal development* or *dynamics*.

Talk of ‘behaviour’ indicates that often what is attributed to systems is not just a set of static properties but rather a certain temporal development. The Lotka–Volterra equations, for example, tell us how the predator–prey system will develop over time. Similarly, the Schrödinger equation applied to certain kinds of systems describes their dynamics.

I have argued that, in order to make sense of the fact that law statements are taken to be confirmable or disconfirmable, we need to understand them as attributing behaviour to systems. The assumption that we can identify systems that behave in certain ways is thus essential for understanding scientific practice, at least to the extent that this practice involves law statements.

Thus, whatever the content of a scientific theory or law, that is, whatever structure is attributed to reality by the law predicate Σ , the analysis of this structure cannot show that there are no systems or that there are no things. For instance, it is sometimes suggested that contemporary physics, particularly the phenomenon of indistinguishable particles, shows that there is no place for objects (or systems) with intrinsic natures in metaphysics (Ladyman and Ross 2007, 131). Such a claim is a non-sequitur, as I will now argue.

For the purpose of illustration, consider a two-electron system. (Normalised) vectors in two-dimensional Hilbert spaces represent the spin states of the separate particles. The possible spin states of the

¹⁰ As I will argue in more detail in Chapter 3, nature may suggest but does not dictate how to draw the boundaries of the systems in which we are interested; that is, it does not dictate how to individuate systems. This is particularly obvious when we are dealing with macroscopic systems. A commitment to systems is thus not a commitment to nature having joints that are completely independent of pragmatic considerations (see Section 3.5.2).

compound system are all those states that can be represented as (normalised) vectors in the tensor product of the Hilbert spaces associated with the separate particles $H_s = H_1 \otimes H_2$. If we take as a basis for H_1 the eigenvectors in the spin z -direction $|\psi^{z\text{-up}_1}\rangle$ and $|\psi^{z\text{-down}_1}\rangle$ and as a basis for H_2 $|\psi^{z\text{-up}_2}\rangle$ and $|\psi^{z\text{-down}_2}\rangle$, the following superposition state will be among the possible states of the (normalised) compound system:

$$\Phi = 1/\sqrt{2} |\psi^{z\text{-up}_1}\rangle \otimes |\psi^{z\text{-down}_2}\rangle - 1/\sqrt{2} |\psi^{z\text{-down}_1}\rangle \otimes |\psi^{z\text{-up}_2}\rangle.$$

If the electrons are in such an entangled state, it is true that the electrons cannot be described as two individual particles with intrinsic properties, where intrinsic properties are conceived of as properties that the systems have independently of the properties of other systems (see, e.g., Ladyman and Ross 2007, 135ff). However, even if the indistinguishability and non-individuality of the electrons are granted, it does not follow that there are no things or systems. The system to which the entangled state is attributed (the ‘two-electron system’) may still be a system with intrinsic properties or states.

The problem with Ladyman and Ross’s claim is not only that the argument from the indistinguishability of particles to the claim that there are no things or systems is a non-sequitur. In fact, when we try to confirm claims about the indistinguishability of particles, we have to presuppose the existence of things or systems.

This fact is nicely illustrated by the research that was awarded the Nobel Prize in Physics in 2001. As one of the recipients of the prize notes, ‘The phenomenon of Bose-Einstein condensation (BEC) is the most dramatic consequence of the quantum statistics that arise from the indistinguishability of particles’ (Ketterle 2007, 159). The prize was awarded for the empirical confirmation of the consequences of the indistinguishability of particles, specifically ‘for the achievement of Bose-Einstein condensation in dilute gases of alkali atoms, and for early fundamental studies of the properties of the condensates’ (Nobel Prize press release 2001). What this reinforces is that confirming the empirical consequences of the indistinguishability of particles requires systems or objects, such as ‘dilute gases’ or ‘condensates’.

Furthermore, the Nobel Prize-winning research concerning the indistinguishability of the parts of a system does not tell us that the properties of the condensate fail to be intrinsic. Bose–Einstein condensates may very well be systems with intrinsic properties or behaviour, which can be studied independently of relations to other systems. It tends to be overlooked that law statements are *statements about systems* if theories are simply

taken to be sets of equations and if they are analysed without taking into account that their empirical import is generated by the assumption that these are true of something in the real world.

It may be objected that considerations of quantum entanglement lead us to the view that there is only one system with intrinsic properties – the universe as a whole. That may very well be true, but it would still be in conflict with the claim that there are *no* systems. Furthermore, what we need to understand is why we can successfully treat subsystems of the universe as if they exhibited their behaviour independently of other systems in the world (more on this in Chapter 2).

According to another objection, there is less of a conflict than it first appears between Ladyman and Ross's denial of systems or things and my insistence that we have to assume their existence. While it is true that structural realists sometimes claim that there are *no things* (Ladyman and Ross 2007, p. 130: 'a first approximation to our metaphysics is: "There are no things. Structure is all there is"'); see also French 2014, chapter 7: 'The Elimination of Objects'), neither Ladyman and Ross nor French deny that we can meaningfully talk about (everyday) objects. Both have their ways of accommodating true assertions about medium-sized objects in a world that (according to their view) strictly speaking does not contain any. Thus, according to French, 'we can reject tables, people, everyday objects in general as elements of our fundamental ontology, whilst continuing to assert truths about them' (2014, 167). However, in contrast to Ladyman, Ross and French, for the reasons outlined in this section I don't think we can reject systems or objects as elements of our fundamental ontology. They cannot be analysed away if we want to understand the success of our scientific practice.

Let us now have a closer look at how Ladyman and Ross argue against intrinsic natures:

[...] talk of unknowable intrinsic natures and individuals is idle and has no justified place in metaphysics. This is the sense in which our view is eliminative; there are objects in our metaphysics but they have been purged of their intrinsic natures, identity and individuality, and they are not metaphysically fundamental. (Ladyman and Ross 2007, 131)

It seems, however, that scientific practice in at least some cases gives us very good reasons to suppose that there are systems with *intrinsic* properties, provided an intrinsic property can be taken to be a property of a system that the system has independently of the properties of other systems and its interactions with such systems.

In experimentation we typically try to shield from interfering factors; we try to causally isolate the system under investigation – that is, we try to figure out how the system would behave if it were on its own – not interacting with the world. Shielding and isolation point to the fact that in experimentation we try to determine properties of systems that these systems have independently of the properties of other systems and their interactions with such systems. Presumably, even in the case of the experimental determinations of the behaviour of dilute gases of alkali atoms, as well as studies of the properties of the condensates, the experimenters determined intrinsic properties of the systems in question.

If we take into consideration scientific practice as a source of metaphysics of science, we have very good reasons to stick with at least some of the features Ladyman and Ross classify as ‘standard metaphysics’ (Ladyman and Ross 2007, 151) – namely, that science deals with systems and their intrinsic properties.

1.3 Modal Surface Structure

What I have argued so far is that law statements attribute complex (multi-track) behaviour to systems via law predicates. I will now examine the modal aspects of this behaviour in more detail. Traditionally, when the nomological or modal character of laws is discussed, the focus is on external generalisations of the kind ‘All Fs are Gs.’ Armstrong, for instance, argues that the properties of, say, being an electron and having a certain charge are related by a *sui generis* relation of nomological necessitation that explains why all electrons have charges. Bird holds that the fact that all negative charges repel each other obtains in virtue of the negative charges’ essence. In this section I will attempt to show that examining internal generalisations and their role in scientific practice will unfold a rich modal structure underlying the characterisation of the behaviour of systems. More particularly, I will advocate the following claims: Law statements attribute a space of possible states to systems (Section 1.3.1). Laws constrain the temporal development of systems by virtue of law equations (Section 1.3.2). The laws’ ability to constrain, their natural necessity, can be explicated in terms of the fact that they are invariant with respect to a number of different kinds of circumstances (Section 1.3.3).

I use the term ‘modal *surface* structure’ because even though I will argue that nomological or natural necessity, as well as all the natural dependence relations we will encounter in later chapters (dispositional modality, causal dependence, part-whole dependence), can be explicated in terms of

invariance, invariance is a modal notion itself, and I will make no attempt to reduce modal facts to non-modal facts.

1.3.1 *A Space of Possibilities*

In order to understand the role of internal generalisations in scientific practice we need to distinguish two features. First, internal generalisations quantify over *a domain of values for variables*, which serve as possible initial or boundary conditions. Second, internal generalisations typically include a *law equation*. The law equation restricts or determines the values for the variables of the system or the values for variables that characterise the temporal development of the states of the system. I will deal with the law equation in the next section.

By virtue of internal generalisations, laws attribute a space of possible states to systems. In the case of dynamical laws, it is assumed that the systems have a set of possible initial states. With respect to these states we can distinguish two cases. Either the domain of quantification comprises all possible states (e.g., in the case of Newton's second law or the Schrödinger-equation) or, as is the case in more specific laws, the domain of quantification comprises only a restricted range of states. Hooke's law, for example, holds only for a limited range of elongations.

What is essential for our investigation is the fact that in both cases we are dealing with a *modal* presupposition because it is not only actual states or actual behaviour with which the internal generalisations are concerned. In fact, the internal generalisation on its own does not even tell us which state of the system is the actual state. The internal generalisation's concern is possible behaviour only (whether actual or non-actual). Thus, the fact that internal generalisations come with a domain of values for variables requires the assumption that law statements attribute a *space of possible (and mutually exclusive) states* to systems.

That laws attribute to systems a space of possible behaviour has recently been discussed as a threat to Humean accounts of laws of nature, because (a) it seems to be *prima facie* problematic to square this feature with the requirement of informational strength and (b) it may indicate that laws give not only information about patterns in the Humean mosaic but also genuinely modal information (see Hall 2015; Hicks 2018; Jaag and Loew 2020). For instance, Ned Hall observes that

it is worth noting that breadth or permissiveness of the [range of initial conditions] makes for a certain kind of explanatory strength. For it is, other

things equal, a point in favor of a physical theory that it recognizes a wide range of nomologically possible initial conditions. Compare, for example, Keplerian and Newtonian accounts of the solar system. Granted that the Newtonian account is much more empirically accurate; it is also, from the standpoint of scientific investigation, better in a distinct sense: for it allows us to answer questions not merely about how the elements of the solar system did, do, and will behave, but also about how they would have behaved under alternative physical conditions. (Hall 2015, 263)

Many laws of nature – in particular the dynamical laws of fundamental physical theories – allow for a wide range of initial conditions. The fact that laws come with a range of possible states is essential for the role laws of nature play in scientific practice, as the following examples illustrate.

One case is the application of law statements in engineering contexts. Suppose an engineer considers different ways to build a bridge. Specifically, she will be considering different (e.g., Newtonian) models for the bridge: she will consider a scenario S_1 in which the bridge is built with materials M_1 and a scenario S_2 in which the bridge is built with materials M_2 . First, the engineer will determine how Newtonian mechanics describes what would be the case if the bridges were built. She will furthermore ascertain what would happen in these models if certain parameters were varied: whether the hypothetical bridge would remain stable if the traffic were of a certain kind, if the weather conditions changed, and so on. Hence, a scientist or engineer will be interested not only in what is actually the case, but also in what is non-actual but (nomologically) possible. Such information about what is possible and what isn't is needed in order to know how to manipulate a system such that it reaches a designated state, for instance that the bridge doesn't collapse given the expected traffic.

More generally, it might be argued that laws play a role in decision making. It is constitutive for decision making that various possible outcomes are examined – in order to explore how different situations would develop – due to the laws of nature. Laws can only play this role because of the fact that they come with a range of (possible) initial conditions.

Furthermore, in the explanation of events – if we follow Woodward and Hitchcock's account of explanation – we appeal to laws or generalisations not because they tell us what is actually the case but because they provide modal information. According to their account, we can explain why a gas G has a certain temperature T_o by showing how the temperature

T depends on the pressure p of the gas and its volume V (Hitchcock and Woodward 2003, 4). Hitchcock and Woodward contrast their account with the deductive nomological account of explanation:

the generalization [...] not only shows that the explanandum was to be expected, given the initial conditions that actually obtained, but it can also be used to show how this explanandum would change if these initial and boundary conditions were to change in various ways. (Hitchcock and Woodward 2003, 4)

According to Woodward and Hitchcock, the counterfactuals that are explanatory do not only appeal to information about what is actually the case (or what was the case) but also to nomologically possible but non-actual behaviour of systems. This possible but non-actual behaviour is characterised in terms of the law equation and the domain of quantification of the internal generalisation. The counterfactuals rely on non-actual but nomologically possible states of the gas and thus on the modal structure that the law statements attribute to systems.

To conclude: Internal generalisations on their own do not tell us which state a system is actually in. To determine the actual state of the system we need additional information – information about the actual values of the variables that characterise the system. With respect to the states of a system, the law statement (by virtue of the domain of internal generalisations) gives us *purely modal* information: information about (nomologically) possible and mutually exclusive states in which the system might be.

1.3.2 Constraints

In the previous section I argued that law statements attribute a space of possibilities to systems due to the fact that internal generalisations quantify over a domain of values for variables that represent mutually exclusive possible states of a system. Let me now turn to the second aspect of internal generalisations that is relevant for the examination of modal structure, the law equation. I will argue that we need to assume a further modal feature to make sense of our scientific practice concerning law statements: law statements do not simply register the past, present and future behaviour of systems; they describe how this behaviour is constrained. While in this section I will introduce the claim, in the next section I will argue that if we understand this claim in terms of invariance relations, this best explains a certain feature of scientific practice, namely, why we can rely on laws.

Internal generalisations put *restrictions* on the space of possible behaviour of systems by establishing relations between variables (i.e., law equations). These restrictions can either concern the synchronic co-possibility of values of variables that characterise the state of a system – as in the case of the ideal gas law – or the temporal evolution of the states of a system – as in the case of the Schrödinger equation.

In the case of a synchronic law (law of coexistence), such as, e.g., the ideal gas law, the set of possible values for the variables p , V and T is restricted to those that satisfy the equation $pV = \nu RT$. Thus, the possible states of the gas are constrained to a two-dimensional hypersurface of the three-dimensional space that is generated by the variables p , V and T . The internal generalisation does not only provide information about how the *actual state* of a system (if known) is constrained. In addition, it tells us how all possible states of the system are constrained, whether or not they are actual. That the systems are constrained means that those states not on the hypersurface are not accessible to the system. They are classified as states the system cannot possibly occupy, given the law equation, i.e., as nomologically impossible states.

The fact that the gas satisfies the equation of the gas law allows a scientist or an engineer who is able to manipulate pressure and volume to ensure that the gas will have a certain temperature. Similarly, the engineer might want to prevent certain situations, such as preventing a gas from having a certain temperature. In such cases she will rely on the fact that the law tells us that certain combinations of pressure, volume and temperature will not occur; by setting pressure and volume appropriately we can make sure that a certain temperature value will not obtain.

The same holds for internal generalisations that describe the temporal evolution of a state of a system. Provided we prepare the system under consideration in a certain state, and provided the equation in question is deterministic, we can ensure that at a later time the system is in a certain state, and we can also prevent the system from being in certain other states.

In the case of prevention, it is not only that given certain combinations of, say, pressure and volume, certain values for T simply do not occur. There is a sense in which these values *cannot* occur.¹¹ The use scientists and engineers make of internal generalisations in scientific practice is best

¹¹ For this reason, Popper conceived of laws as 'prohibitions' (Popper 1959, §15).

understood by assuming that internal generalisations represent modal, that is, nomologically necessary, relations.¹²

That internal generalisations ought to be understood in this way requires taking a certain perspective, the perspective of a scientist or engineer operating *within* the universe, in contrast to the omniscient outside observer whose sole job is to document the world. In discussions about laws of nature the dominant perspective is that of ‘a scientist operating outside the universe and looking in. This ideal scientist starts with the knowledge of all the facts of the world, so the only task left to her is to organize them’ (Hicks 2018). This ideal, inward-looking scientist has two pertinent features: she is omniscient and she is interested in the description or the organisation of facts only. This, however, is not the perspective taken in scientific practice and it leaves out why laws are best understood as – at least *prima facie* – representing nomologically necessary relations. The essential difference is that a real engineer or scientist – in contrast to the ideal scientist – will *rely* on the internal generalization in the sense that the generalization tells her that any other value than the one that is determined by the equation $pV = \nu RT$ *cannot* occur.¹³ The notion of invariance that I introduce in the next section will clarify what it means that a scientist relies on laws.

1.3.3 Invariance

It may seem that in the previous section I have illegitimately smuggled in modal terminology. Instead of saying that the law equation of an internal generalisation states a relation between the values of different variables, I have said that the values of the variables are *constrained* or *restricted*, that they obtain with *nomological necessity*. Why this modal terminology? Why claim that a certain value of T *cannot* (fail to) occur? Why argue that, provided certain values for p and V , the occurrence of a certain value of T is *nomologically necessary* and that we can rely on this?

The essential point is that laws – as opposed to accidental generalisations – do not simply state that a relation between variables obtains. Looking at

¹² As a reminder, I am not arguing against the Humean at this point; I am interested in the modal *surface structure* that may or may not be reducible to non-modal facts.

¹³ A focus on laws as instruments or tools for limited beings may have always been part of Lewis’s best system analysis (Jaag and Loew 2020). Recently, there have been various attempts to argue along these lines, that is, to claim that laws should be conceived as instruments for cognitively limited beings that operate inside the universe (Hicks 2018; Jaag and Loew 2020; Ismael 2015). However, in the past, in the literature on laws of nature there has always been a focus on explanation and description rather than on manipulation.

scientific practice, i.e., not simply looking at the law statement or the law equation itself but examining their role, i.e., how they are used, reveals that law statements should be understood as implying independence or invariance claims. For example, the ideal gas law is understood as implying that in a gas the value for T is fixed *under a wide variety of circumstances*: if certain values for p and V are fixed, then *whatever other features the gas may have and whatever else is going on in the universe*, a certain value for T is determined. It is part of how the content of the internal generalisation is understood that it is only the values of p and V that determine the value for T . An essential aspect of what laws tell us is that those variables that *do not* occur in the law equation are irrelevant for the determination of the values of certain other variables whatever the circumstances may be.

The ideal gas law, if true, is not a mere truth. Its being a law means that nobody could bring about a situation such that it is false. No person or government could bring it about that in an ideal gas the values of pressure, temperature and volume fail to be on the two-dimensional hypersurface that is determined by the ideal gas law.¹⁴ In this sense, the behaviour of the gas is constrained or nomologically necessary. By contrast, an accidental generalisation may state that a certain equation or correlation between, say, variables representing the bread prices in London and the water levels in Venice may actually obtain. However, the equation representing the accidental generalisation is not taken to be invariant – it is not taken to continue to hold if, for instance, the British government fixes the price of bread.

The fact that the law and the accidental generalisation are used differently, that they play a different role in scientific practice, is best explained by assuming that laws state relations that are invariant. An illustration of this claim is that we can rely on laws but not on accidental generalisations, precisely because the former imply, as I argued before, that many features of the universe are irrelevant for the determination of the behaviour of the system we are interested in. The engineer or scientist *in our world* – in contrast to the ideal, inward-looking scientist – does not know all the facts of the world. She is confronted with epistemic risk when she is predicting, manipulating or constructing systems. What we need to explain is why this scientist or engineer can rely on what she takes to be laws. What is it about laws that accounts for the possibility of relying on them? With the notion of invariance, we can make sense of

¹⁴ The same holds for real gases and more realistic gas law equations such as the Peng–Robinson equation.

why the scientist and the engineer rely on (what they take to be) laws rather than on accidental regularities.

Suppose a scientist *S* at a certain time has encountered the same number of positive instances for two claims, $P_1: \forall x (Fx \rightarrow Gx)$ and $P_2: \forall x (Mx \rightarrow Nx)$. Suppose that in the case of P_2 but not in the case of P_1 there is furthermore evidence for invariance – evidence, say, that P_2 holds under all kinds of changes of the behaviour of other (e.g., neighbouring) systems. For *S* it is now much more reasonable to be confident that P_2 will continue to hold, compared with P_1 , because there is evidence for the fact that P_2 is stable and will not break down if changes in the environment occur. *S* can and will rely on P_2 much more so than on P_1 . Thus, if having evidence for the nomic character of a generalisation means – as I have argued – having evidence for invariance relations, we can understand why scientists and engineers rely on laws much more than on other generalisations.

To sum up, the role that internal generalisations and in particular law equations play in scientific practice is best understood by assuming not only that they describe relations between variables but also that they imply that these relations hold with nomological necessity, which is best understood in terms of the fact that they are invariant in a number of respects. In the remainder of this section I will examine these invariances in some detail.

The idea of spelling out the modal aspects of laws of nature in terms of invariance is not new; Mitchell (2003, 140), Lange (2009) and Woodward, to name a few, have done so before. I will not discuss any of these approaches in any detail. Let me, however, mention Woodward (1992; 2018), whose view is probably closest to the one presented here. He explicitly endorses the idea that nomological necessity should be understood as an invariance claim: ‘We may say that a law, in contrast to an accidentally true generalization, expresses a relationship which not only holds in the actual circumstances but which will remain stable or invariant under some fairly wide range of changes or interventions’ (Woodward 1992, 202).

Invariance is clearly a modal notion. It concerns not only actual but also counterfactual changes. While it might be argued that moving from one modal notion (nomological necessity) to another (invariance) is not much progress, there is certainly at least one advantage. The notion of invariance naturally leads to a closer examination of the modal structure delineated by the internal generalisations. Invariance is a *relative* notion, and we thus have to ask, ‘Invariance with respect to what?’ Furthermore, as we will see later, the concept of invariance helps to understand how modal notions can be empirically accessible.

While Woodward has briefly alluded to some distinctions among invariance relations (2003, chapter 6; 2018), there has been no systematic examination.

We can specify the following invariance relations, which will shed light on the inviolability of laws or their nomological necessity:

- 1) *Invariance of the law equation with respect to initial conditions.* The law equation holds irrespective of which values from a certain range of initial conditions, or boundary conditions or other sets of variables, obtain or would obtain. Newton's second law, a dynamic law, was supposed to hold for any place and any initial velocity characterising a system. By contrast, the ideal gas law holds only for a restricted range of values of p and V (when read as determining T). Both law equations are invariant with respect to at least some initial conditions; note that this is an invariance with respect to the *values* of variables that explicitly figure in the law equation.

The other two kinds of invariances that I will discuss concern invariances with respect to features of the world that are not represented as variables in the law equation.

- 2) *Invariance of the law equation with respect to other features of the systems.* When we determine the speed of a free-falling body there are a number of properties that can be ignored, such as the shape and the colour of the falling body. That is what Galileo's law implies by not mentioning them. The fact that certain variables do not figure explicitly in the law equation (and cannot be determined by those explicitly mentioned) implies that the law equation is invariant with respect to these variables. The law equation in the ideal gas law is in practice understood as saying not only that the temperature of actual ideal gases is determined by the pressure and the volume but also that there is no other feature that might be relevant: neither smell, nor the shape of the molecules, and so on.

In the case of macroscopic laws, two kinds of properties with respect to which invariance may occur ought to be distinguished:

- a. Same-level properties; for example, colour, shape, mass in the case of free-falling bodies.
- b. Lower-level or constitutional properties; for example, the molecular structure of the gases in the case of the ideal gas law. This kind of invariance plays a major role in discussions about universality in the context of phase transitions. Two different kinds of

micro-physical invariances are relevant. First, macroscopic physical properties may be invariant with respect to changes of the system's dynamical state on the micro-level. Second, the macro-behaviour might even be invariant with respect to non-actual counterfactual changes in a system's composition at the micro-level. For instance, in a ferromagnetic system one might add next-nearest neighbour interactions to a system originally having only nearest neighbour interactions and scale down the strength of the original interaction in a manner that would leave the macroscopic magnetization of a system invariant (see Hüttemann, Kühn and Terzidis 2015).

Whether or not a law equation is invariant with respect to features of the system not represented in the law equation is a matter that needs to be established empirically. If such invariances obtain, we are justified in abstracting away from the relevant features of the system in the law equation we use to describe the system's behaviour. It may, of course, happen that as a result of empirical investigations, certain invariance claims with respect to other features of the system under consideration have to be given up. That will lead to law equations that contain more variables. Thus, for instance, the ideal gas equation was at some point replaced by the van der Waals equation, the Peng–Robinson equation, and so on, equations that take into account some features of the molecules that constitute the gas.

- 3) *Invariance of the law equation with respect to the behaviour of other systems in the universe.* Newton's second law adequately describes the behaviour of any (physical) body irrespective of the behaviour of other systems in the universe. This is not to say that other systems cannot have an influence on the system we attempt to characterise in terms of Newton's second law but rather that the law equation continues to hold for the body despite external forces being impressed on it. Whatever the forces are that affect a certain system, Newton's second law will continue to hold for the system under investigation.

The status or role of the three kinds of invariance relations is somewhat different. The first kind of invariance (invariance with respect to initial conditions) merely expresses the fact that the law equation holds for more than one set of values of the relevant variables. This is part of the content of what the law explicitly says about the behaviour of systems. The second kind of invariance (invariance with respect to characteristics of the system

that are not represented in the law equation) provides a rationale for abstracting in the law equation from the features in question. The third kind of invariance (invariance with respect to the behaviour of other systems in the universe) seems to be most relevant for accounting for the inviolability or natural necessity of laws. If this kind of invariance holds, then whatever changes there are, the law equations will remain the same.

There is an interesting and important further contrast between the invariance of the law equation with respect to other features of *the same system* and the invariance of the law equation with respect to the behaviour of *other systems in the universe* when it comes to evidence for the failure of invariance. Whereas in the former case, as already indicated, the law equation will be revised (in the light of sufficient evidence), in the latter case, one strategy is to hedge the law by a *ceteris paribus* clause. Instead of revising the law equation in Galileo's law when it comes to falling objects in water or other media, it may be argued that the law holds *ceteris paribus*. I will deal with this issue in Chapter 2.

Let me add a few remarks about the notion of invariance.

First, I agree with Woodward that invariance-based accounts of nomological necessity 'provide a naturalistic, scientifically respectable and non-mysterious treatment of what non-violability and physical necessity amount to' (Woodward 2018, 160). Invariance claims can be scientifically investigated. For example, if it is argued that a certain law equation is invariant with respect to the colour of the system, it is reasonably clear what is claimed and we know how to check the claim. Furthermore, macroscopic invariance claims (e.g., the invariance of the behaviour of gases with respect to their constitution) can at least in principle be explained in terms of lower-level laws (as well as experimentally investigated).

Let me add that I see no reason to believe that there is a special problem with knowing invariance facts simply because they are modal facts. Here is, for instance, how Galileo's spokesman Salviati, on the basis of experiments, argued for the claim that bodies with different densities would fall with equal speed in a vacuum, that is, that their speed is invariant with respect to the density of the bodies:

We have already seen that the difference of speed between bodies of different specific gravities is most marked in those media which are the most resistant: thus, in a medium of quicksilver, gold not merely sinks more rapidly than lead but it is the only substance that will descend at all; all other metals and stones rise to the surface and float. On the other hand, the variation of speed in air between balls of gold, lead, copper, porphyry, and other heavy materials is so slight that in a fall of 100 cubits a ball of gold

would surely not outstrip one of copper by as much as four fingers. Having observed this I came to the conclusion that in a medium totally devoid of resistance all bodies would fall with the same speed. (Galileo 1954, 71–2)

Salviati's inference becomes problematic only if one already starts out with the idea that our epistemic access is limited to the actual. Laws of nature make modal claims that are empirically accessible.

Second, this latter claim is further illustrated by the fact that additionally, some more specific notions of invariance play an explicit role in physics. Some laws or theories are characterised as Galilei invariant, Lorentz invariant or Gauge invariant. These concepts point to symmetries in the systems under investigation, symmetries with respect to certain classes of transformation. Not every law or theory is Lorentz invariant, and whether or not the behaviour of systems can be characterised as Lorentz invariant is an empirical matter. The notion of invariance that is relevant in these physical discussions is the same as I used previously; that is, the law equations continue to hold under certain kinds of actual and counterfactual changes. More specifically, the homogeneity of space, for instance, is a symmetry in the sense that the laws of classical physics remain the same under translations in space; all space points are equivalent when it comes to Newton's second law or the Schrödinger equation. Absolute space points turn out to be irrelevant for the dynamics of systems (Castellani 2003, 429). The same holds for Lorentz transformation, and so on. Laws are invariant if they stay the same under actual or counterfactual changes. Some of these invariances, namely those listed here, are constitutive of what is usually considered to be nomological necessity. Others, such as Lorentz invariance, are additional invariances that laws may or may not comply with.

A third remark addresses an objection against analysing lawhood or nomological necessity in terms of invariance relations. Psillos criticises Woodward's account for being circular because the notion of invariance presupposes that of a law. Lawhood, Psillos argues, should thus not be explicated in terms of invariance. Rather, 'some laws must be in place before, based on considerations of invariance, it is established that some generalization is invariant under some intervention' (Psillos 2002, 185). The disagreement concerns the question of whether lawhood or invariance should be taken to be the fundamental notion. I have argued that we have good reasons why we should explicate lawhood or nomological necessity in terms of invariance (rather than the other way around), because we thus understand the role nomological necessity plays in scientific practice.

In dealing with the circularity objection, it is important to distinguish an ontological issue and an epistemological issue. The ontological issue concerns the question of what the invariance of a law equation with respect to, say, the behaviour of other systems *consists in*. The simple answer is that the equation continues to hold when there are actual or counterfactual changes in the behaviour of other systems. There is no need to refer to laws of nature when it comes to explicating what invariance consists in.

The epistemological issue is a different issue. It concerns the question of how we might know that a certain invariance claim is true: how we come to know that a certain law equation would continue to hold given changes in the behaviour of other systems, or – to use Psillos’s phrase – how the invariance claim is ‘established’. In order to decide whether or not a certain generalisation is invariant with respect to certain changes we may indeed rely on laws. But that is no problem for the invariance account as long as the ontological and epistemological issues are kept apart. Thus, laws may play a role in establishing invariance claims, but that does not imply that the ontological characterisation of invariance presupposes the notion of a law.¹⁵

1.3.4 External Generalisations and Modal Surface Structure

Let me now turn briefly to external generalisations and the question of whether their role in scientific practice requires additional modal assumptions. In Section 1.1.3 I argued that the explanation of why the stone has a certain velocity presupposes that the system in question, the free-falling stone, falls under Galileo’s law. The truth of the external generalisation, I argued, is a presupposition for the internal generalisations doing their explanatory work. Furthermore, in the case of our example of standard explanation, all that is required is that the external generalisation is true. Thus, standard explanation, which is of course only one aspect of scientific practice, does not commit us to postulating modal structure, let alone *additional* modal structure.

The case is different when it comes to prediction or manipulation. Let us go back to our preliminary law statement:

(A) ‘All systems of a certain kind K behave according to Σ .’

¹⁵ Woodward mixes these two issues up by defining invariance in terms of interventions (see, e.g., Woodward 2003, chapter 6).

As already indicated, in predicting, manipulating or constructing systems we rely on laws, and this reliance can be explicated in terms of the invariance of the law equation with respect to actual or counterfactual changes in the behaviour of other systems in the universe. However, the fact that the law equation continues to characterise the behaviour of the system under consideration presupposes that the external generalisation (A) is invariant with respect to the same changes. No matter what other systems are doing, all systems of kind K will continue to behave according to Σ . So, when it comes to manipulation or prediction, external generalisations commit us not only to their truth but also to their invariance with respect to the behaviour of other systems in the universe. There is, however, no evidence that we are committed to any new kind of invariance relations that we have not encountered in our discussion of internal generalisations.

Let me close this section on external generalisations by giving a diagnosis as to why the modal surface structure that is delineated by law statements has often been ignored. Two facts seem to me to be relevant here. First, the reconstruction of law statements along the lines of the 'All ravens are black' paradigm has encouraged us to overlook the role of internal generalisations. Second, when scientific practice was analysed at all, it has often been assumed that description and explanation is all there is. As I have argued, we can understand the role of external generalisations with respect to description and explanation without making modal assumptions, since the truth of the external generalisation suffices in these cases. It is only when we consider other scientific practices, such as predicting or manipulating, or the role internal generalisations play in these practices, that we are forced to consider the modal surface structure of law statements.

1.4 Conclusion

The purpose of Section 1.3 was to examine the modal surface structure of laws, that is, those modal assumptions that best account for how we make use of external and internal generalisations in scientific practice. I argued that in order to understand not just our explanatory practice but also other practices that involve a reliance on laws, such as predicting, manipulating or constructing systems, we have to make two assumptions. First, we have to assume that law statements attribute a space of possible states to systems. Second, we must assume that both the law equations and the external generalisations are invariant with respect to certain actual and counterfactual changes. The law equation is invariant with respect to (i) a range of

initial conditions and (ii) features of the system that do not figure in the law equation. In addition, both the law equation and the external generalisation are invariant with respect to changes of the behaviour of other systems in the universe. These empirically accessible invariance claims account for what is usually taken to be the law's nomological necessity.

As already mentioned, it is important that invariance claims are modal claims. It is claimed not only that law equations are invariant with respect to actual changes but also that they are invariant with respect to counterfactual changes. Some may feel that for this reason we have made little progress. Nomological necessity, they will argue, is mysterious precisely because it is a modal notion; invariance is a modal notion too, so we are left with a mystery. I think this is wrong. Being able to characterise the modal structure that comes with law statements in more detail and explaining how invariance claims are empirically accessible does seem to me to constitute progress.

Finally, as mentioned before, I am exclusively concerned with the modal surface structure of laws, which the notion of invariance helps to characterise. One may then still ask 'What underpins these invariances?' (Bird 2007, 5). That may be an interesting question, but it is one that transcends the approach taken in this book. As I said at the beginning, I want to confine myself to metaphysical claims that can be established via an inference to the best explanation of why we have the scientific practice we have. Assuming that laws are invariant in various respects does exactly that. Invariance relations, as I will show in the following chapters, can account for (almost) all the natural modalities we encounter in scientific practice. A further analysis of invariance in terms of subjunctive facts (Lange), essences (Bird), dispositional modality (Mumford) or the Humean mosaic does not do any extra work in explaining scientific practice. A minimal metaphysics of scientific practice should thus abstain from such hypotheses of how natural modality is to be further explained or reduced. Those who do conduct such investigations have to appeal to intuitions about the governing of laws, intuitions about the supervenience or non-supervenience of laws on the underlying Humean mosaic, or intuitions about quiddities, or to a preference for desert landscapes. It is hard to see an argument as to why such intuitions or preferences should be considered to be truth conducive. Thus, I am quite happy with the modal surface structure for which I have argued. 'Those who go beneath the surface do so at their peril' (Oscar Wilde).