Laws and Lawmakers

Science, Metaphysics, and the Laws of Nature

Marc Lange

OXFORD
UNIVERSITY PRESS

2009

Contents

1

Laws Form Counterfactually Stable Sets

1.1. Welcome

I am sitting in the waiting room of a car-repair shop, pounding out these words on my laptop computer while waiting for my car to be fixed. In a host of ways, the laws of nature impinge upon my situation. The auto mechanics are relying on their knowledge of various natural laws in trying to explain why my car is misbehaving and to predict whether it would stop doing so if some of its components were adjusted or replaced. My laptop's cooperative behavior arises from the electrons in its circuits obeying various laws that the laptop's designers were justified in predicting they would obey. Those electrical engineers designed the laptop's insides as they did, rather than according to some other blueprint, because they knew that it would not have worked had they designed it in that other way. Presumably, in using the laws of nature to predict how laptops of various possible designs would behave, the engineers took it for granted that the actual laws would still have been laws even if the laptop's innards had been different. The same applies to the auto mechanics who are relying upon their knowledge of the natural laws to tell them how various possible repairs would affect my car. By the same token, the electrical engineers and auto mechanics also presumably believe that the actual laws would still have been laws even if computers and cars had never been invented—or, for that matter, even if humanity had never evolved or the sun had never formed.

These examples illustrate the most important roles that laws of nature play in science: in connection with scientific explanations, predictions of the future, and "counterfactuals" (that is, "predictions" of what would have happened under certain circumstances that do not actually come about). Whereas scientists aim to discover what the laws of nature are, philosophers aim to identify what it is to *be* a law of nature—in other words, what *makes* a given fact qualify as a law rather than an "accident." For example, philosophers aim to specify what it is about reality in virtue of which it is a law that all gold objects are electrically conductive but an accident (if it is true at all) that all gold cubes are smaller than one cubic mile—to take Hans Reichenbach's famous example.[1] What it is to be a natural law must account for the work that natural laws do in connection with scientific explanations, predictions, and counterfactuals.

In this book, I propose an account of what natural laws are that explains why they do what they do. Admittedly, I lack any sort of knockdown argument for my proposal. Nevertheless, I think that my account has some novel and even (dare I say it?) elegant features that make it attractive.

Although I will not comprehensively survey the other accounts of natural law that are currently on the market, I will occasionally highlight some respect in which my proposal contrasts with its chief rivals. I will also present a few general recipes for generating worries about other proposed accounts of natural law. My own proposal is designed to avoid these worries. Ready? OK, then—let's begin.

1.2. Their Necessity Sets the Laws Apart

Laws of nature have traditionally been contrasted with two other kinds of facts: accidents and "broadly logical" truths. What separates these three kinds of facts?[2]

Let's start with the accidents. The word "accident" here is a bit of potentially misleading philosophical jargon. Please do not confuse it with the ordinary meaning of "accident"—what I mean when I say to you, "Our meeting here was no accident; I was looking for you," or when the owner of a car dealership confidentially informs us, "It is no accident that every car in my showroom smells so fresh; I put the same chemical in each of them, to give them all that fabled 'new car smell.'" To call some events "accidental" in this ordinary sense is roughly to say that they were unintentional ("I accidentally spilled my soup"), that they were just a coincidence ("I encountered him by accident; I wasn't trying to find him"), or that there is no common explanation from which they all spring. On the other hand, in philosophical discussions of natural law, an "accident" can be no coincidence. Indeed, an accident typically has an explanation. For example, that every car in the showroom has the same "new car" smell is no coincidence (the dealer has just told us the explanation), and it certainly was intentional on the dealer's part, but it is still an "accidental" truth in the philosophical sense—the sense relevant to this book. An "accident" in that sense is simply a truth that does not follow from the natural laws (and the "broadly logical" truths) alone. In other words, an accident could have failed to hold without any violation of the natural laws. For example, no natural law has to be violated for the showroom to contain a car without the "new car" smell.

Take another example: suppose that many apples are hanging on the tree in my backyard, and all of them are now ripe. Their ripeness is an accident, since even if some of them were not ripe, the laws of nature could still all have held. If the warm weather had arrived a few weeks later, for instance, then the apples would not yet have been ripe, though the natural laws would have been no different. Nevertheless, it is no coincidence that every single one of those apples is ripe today. Their ripeness resulted from the recent weather conditions, the levels at which various plant hormones have been flowing through the tree, and so forth. Since each of these factors was common to every one of those apples, they all ripened together. Certain laws of nature governing chemical reactions are also responsible for the apples' ripeness. These laws determined how the weather, the plant hormones, and so forth influenced the rate at which the apples ripened. Again, that there are laws and other conditions explaining why all of those apples are now ripe does not keep this fact from qualifying as an accident. Those other conditions are themselves accidental; there is no explanation of the ripeness of any of those apples that appeals to no accidents at all, but exclusively to laws of nature.

Let's now contrast the laws of nature with the "broadly logical" truths. A broadly logical truth possesses a kind of *necessity* that is possessed neither by natural laws nor by accidents. For instance, one kind of broadly logical truth consists of the mathematical truths, such as the fact that there is no way to divide 23 evenly by 3. There does not merely *happen* to be no integer that added to itself, and then added again, equals 23—in the way that there merely happens to be no gold cube larger than a cubic mile (and even in the way that like charges merely happen to repel rather than to attract). Rather, there *couldn't* have been an integer that added to itself, and then added again, equals 23. That it is *impossible* to divide 23 evenly by 3 explains why no one has ever succeeded in figuring out a way to do so, no matter how much mathematics she knows—and why every time someone tries to divide 23 objects evenly into thirds, she fails. None of these efforts could have succeeded. They all fail because they must; their failure was inevitable. Analogous considerations apply to other kinds of broadly logical truths, such as conceptual truths (for example, "All sisters are female"), narrowly logical truths ("Either all emeralds are green or some emerald is not green"), and metaphysical truths ("Red is a color" or perhaps "Water is H_2O").

Just as each broadly logical truth possesses some variety of necessity that accidents and natural laws lack, so likewise there is a species of necessity possessed by every natural law but by none of the accidents. An accident (even one that is not an utter coincidence, such as that all of the apples on my tree are ripe) just happens to obtain. A gold cube larger than a cubic mile could have formed, but (I presume) the requisite conditions happened never to arise. In contrast (again following Reichenbach), it is no accident that a large cube of uranium-235 never formed, since the laws governing critical masses and nuclear chain-reactions prohibit it. (Actually, they merely render it extremely unlikely, but let's ignore this detail for the sake of Reichenbach's nice example.) In short, things *must* conform to the laws of nature. The laws *prohibit* perpetual-motion machines. It is not simply that all material objects accelerated from rest *as a matter of fact* fail to exceed the speed of light. Rather, they *could not* have exceeded light speed; their failure to outpace light is inevitable, unavoidable—necessary. As a popular tee-shirt says, the speed limit of 3×10^8 meters per second isn't just a good idea; it's "THE LAW."

That a characteristic variety of necessity is possessed by the natural laws, setting them apart from ordinary, run-of-the-mill facts, has been recognized for as long as the concept of a natural law has been explicitly invoked. Here, for example, is Richard Hooker way back in 1593, explaining how obedience to the natural laws is compulsory: "[T]hings naturall . . . doe so necessarily observe their certaine Lawes, that as long as they keepe those formes which give them their being they cannot possibly be apt or inclinable to do otherwise then they do."[3] A thing's "forme," Hooker says, enrolls it in a natural "kinde" and explains its "working." According to Hooker, talk of "formes" can be translated into talk of "lawes": a thing's behavior is explained by the laws applicable to the kind of thing it is.

That the laws, by virtue of their necessity, possess an explanatory power that accidents lack was famously codified about 350 years later by Carl Hempel's "covering law" models of scientific explanation. Whatever the shortcomings of Hempel's overall account of explanation, its motivation in the differences between laws and accidents is easy to appreciate.[4] For example, a certain powder burns with yellow flames, rather than flames of any other color, because the powder is a sodium salt and it is a law that all sodium salts, when ignited, burn with yellow flames. (This law, in turn, is explained by more fundamental laws.) The powder *had* to burn with yellow flames considering that it was a sodium salt; this "had-to-ness" expresses the laws' distinctive kind of necessity. In contrast, we cannot explain why my wife and I have two children by citing the regularity that all of the families living on our block have two children—since this regularity is accidental. Were a childless family to try to move onto our block, it would not encounter an irresistible opposing force or automatically acquire two children upon moving in. Like a natural law, a mathematical truth possesses explanatory power by virtue of its necessity. For example, the fact that 23 cannot be divided evenly by 3 explains why it is that every time mother tries to divide 23 strawberries equally among her three children without cutting any (strawberries), she fails.[5]

Although laws possess a variety of necessity, it is widely believed that laws are *not* as necessary as the broadly logical truths. It has generally been thought that the natural laws could have been different (as they are in certain science-fictional universes, where starships travel beyond light speed), whereas the broadly logical truths are a great deal less malleable. (How *could* there have been a round square?) The status of the natural laws thus sounds paradoxical: they are contingent necessities, falling somewhere between the broadly logical truths and the accidents. Faced with the laws' apparently anomalous(!) character, some philosophers (such as Brian Ellis and other "scientific essentialists") have tried to argue that the laws are in fact every bit as necessary as the broadly logical truths. Other philosophers (such as David Lewis and other "Humeans") have taken the opposite approach, arguing that the laws are not separated from the accidents by a profound metaphysical gap. In contrast to both of these approaches, my account will recognize the laws' distinctive, intermediate status.

In what does the laws' necessity consist? To answer this question is one of the main goals of this book, since to do so would be to understand what laws of nature are. Laws are set apart from accidents by the necessity they possess and from broadly logical truths by the necessity they lack. An account of the laws' necessity should reveal not only how laws differ from accidents and from broadly logical truths, but also how the laws' "natural necessity" qualifies as a kind of *necessity*—as a species of the same genus as the variety (or varieties) of necessity possessed by [narrowly] logical, conceptual, mathematical, and metaphysical truths.

1.3. The Laws' Persistence under Counterfactuals

Whereas a large gold cube could have formed, it is inevitable, unavoidable—necessary—that all bodies accelerated from rest fail to exceed light's speed. Had Bill Gates wanted to build a large gold cube, then (I dare say) there might well have been a gold cube exceeding a cubic mile.[6] But even if Bill Gates had wanted to arrange for a body to be accelerated from rest to beyond the speed of light, no body would have been so accelerated. Even if Bill had possessed 23rd-century technology, he would have failed had he tried to accelerate a body from rest to beyond the speed of light.

In elaborating the laws' necessity, I have just made use of "counterfactual conditionals" (or "counterfactuals," for short): statements of the form "Had p been the case, then q would have been the case" (symbolically: $p \mathrel{\Box\!\!\rightarrow} q$, where I will refer to p as the "counterfactual supposition" or "antecedent" and to q as the "consequent").[7] For example: Had we tried to build a perpetual-motion machine, the law of energy conservation would have prevented our success. We cannot get around the natural laws; they are unavoidable. Counterfactuals are used to express the fact that the laws would still have held even if various other things had been different. The laws of nature govern not only what actually happens, but also what would have happened under various circumstances that did not actually come to pass.

Counterfactual conditionals, though asserted by each of us every day, are notorious in philosophy. To begin with, they seem to violate rules of reasoning that look sensible and are obeyed by many other kinds of conditionals. For instance, $p \mathrel{\Box\!\!\rightarrow} q$ and $q \mathrel{\Box\!\!\rightarrow} r$ do not logically entail $p \mathrel{\Box\!\!\rightarrow} r$. As an illustration, let $p \mathrel{\Box\!\!\rightarrow} q$ be "Had Socrates been a woman, then Socrates would not have been a philosopher." That's true, considering the mores of ancient Greek society. (Of course, a woman in ancient Greece could have lived "a life of the mind," and some did. But evidently and regrettably, she could not have been a "professional" philosopher.) Now let $q \mathrel{\Box\!\!\rightarrow} r$ be "Had Socrates not been a philosopher, then Socrates would still have been a man." Of course, that's true, too. But $p \mathrel{\Box\!\!\rightarrow} r$ is "Had Socrates been a woman, then Socrates would still have been a man," which is obviously false.

Another notorious feature of counterfactuals is that it is mysterious what makes (some of) them true. It is quite clear which features of the world are responsible for making it true (if it is true) that all gold cubes are smaller than a cubic mile: namely, the sizes of the various gold cubes throughout the universe's history. But it is much harder to identify the features of the world in virtue of which it is true (if it is true) that had Julius Caesar been in command during the Korean War, then he would have used the atomic bomb.[8] Nevertheless, the standard view is that such a counterfactual is made true by certain features of the actual world, such as Caesar's hotheaded personality. But how, in general, are the responsible features of the world to be picked out? Nelson Goodman famously grappled with this problem.[9] He began with the attractive thought that "Had the match been struck, it would have lit" is made true by various natural laws along with the fact that the match is actually dry, oxygenated, well made, and so forth. However, no matter how many facts about the actual match Goodman included, he found that they did not suffice to make the counterfactual conditional obtain. To them must be added the truth of various other *counterfactual conditionals:* that had the match been struck, it would *still* have been dry, oxygenated, and so forth. It is difficult to see how the original counterfactual's truth can be reduced entirely to noncounterfactual facts (and the laws of nature), though there have been many ingenious attempts to solve this problem.[10]

Another feature of counterfactuals that is widely considered suspect is their context-sensitivity. The counterfactual "Had Caesar been in command during the Korean War, then he would have used the atomic bomb" is true in some contexts, whereas in other conversational environments, the counterfactual "Had Caesar been in command during the Korean War, he would have used catapults" is true instead. Which features of the actual world are preserved under a given counterfactual supposition (that Caesar was in command during the Korean War) and which are allowed to vary (Caesar's personality or his knowledge of armaments) depends to some extent upon our interests in entertaining that supposition. If we are in the midst of illustrating Caesar's gung-ho personality, then the point of the counterfactual is to contribute to that discussion, and so it is true in that context that had Caesar been in command during the Korean War, he would have used the atomic bomb. Of course, not every counterfactual supposition is relevant in every context. If we are trying to describe Caesar's personality, then to ask, "What weapons would Caesar have used in the Korean War, had Caesar been more cautious and less ambitious?" is utterly beside the point.

Although counterfactuals are context-sensitive, violate attractive-looking logical principles, and are not straightforwardly made true by features of the actual world, counterfactuals are not utterly disreputable. There are strict logical principles regulating their use, even if those principles are not quite the most familiar ones. For instance, $p \mathbin{\Box\!\!\rightarrow} q$ and $(p\&q) \mathbin{\Box\!\!\rightarrow} r$ logically entail $p \mathbin{\Box\!\!\rightarrow} r$. (The Socrates example where $p \mathbin{\Box\!\!\rightarrow} q$ and $q \mathbin{\Box\!\!\rightarrow} r$ are true, but $p \mathbin{\Box\!\!\rightarrow} r$ is false, does not violate this principle, since $(p\&q) \mathbin{\Box\!\!\rightarrow} r$ is "Had Socrates been a woman and a nonphilosopher, then Socrates would still have been a man," which is false.) Even young children have little trouble figuring out whether certain counterfactuals are true. Without giving it a second thought, we routinely assert such counterfactuals as "Had I not gotten lost along the way, I would have arrived here sooner." We are sensitive to the context-sensitivity of counterfactuals just as we easily grasp which sorts of remarks are relevant in a given conversation. Science ascertains that various counterfactuals are true, as when Lavoisier discovered that someone who is standing up and moving about would have consumed less oxygen had she instead been sitting quietly.

Indeed, our observations confirm the truth of various counterfactuals just as they confirm various predictions about the actual world.

Our past observations of emeralds confirm not only that all of the actual emeralds lying forever undiscovered in some far-off land are green, but also that had there been emeralds in my pocket right now, then my pocket would have contained something green. (It is not self-evident which of these predictions is more "remote" from our observations.) When we confirm that my pocket would contain something green were there emeralds in it, that confirmation is unaffected by whatever evidence we may have regarding whether there actually are any emeralds in my pocket. So in confirming that my pocket would contain something green were there emeralds in it, we may be confirming both a prediction about the actual world and a counterfactual conditional. Claims about what *would have been* are confirmed right along with claims about what actually *is*.

(A claim like "Were there emeralds in my pocket, then my pocket would contain something green" is a "subjunctive conditional" rather than a counterfactual, since it fails to presuppose that there are no emeralds in my pocket. The corresponding counterfactual conditional is "Had there been emeralds in my pocket, then my pocket would have contained something green." The connection between the antecedent and consequent of a subjunctive conditional that is true is presumably the same as the connection between the antecedent and consequent of a counterfactual conditional that is true. I shall use the symbol "$p \mathrel{\Box\!\!\rightarrow} q$" to represent both subjunctive and counterfactual conditionals, and I shall often use the term "counterfactuals" to encompass both.)

Although a given counterfactual conditional may have different truth-values in different contexts, this phenomenon is hardly confined to counterfactuals. Claims with indexicals ("I," "now") and demonstratives ("this") display the same behavior. Plausibly even some claims without indexicals or demonstratives do, too: how close Jones's height must be to exactly six feet, in order for "Jones is six feet tall" to be true, differs in different contexts. One possible explanation of these phenomena is that the same sentence expresses different propositions on different occasions of use. Claims with indexicals and demonstratives certainly appear to do so without provoking undue suspicion.

In chapter 4, I shall say more about the facts expressed by counterfactuals. Fortunately, we do not need to have a philosophical account of the truth-conditions of counterfactuals in order to be entitled to use counterfactuals as we ordinarily do: to have (in a given context) great confidence in the truth of certain counterfactuals and the falsehood of certain others. Our goal in this chapter is to identify precisely how laws differ from accidents in their relation to counterfactuals. Having done so, we will then be in a good position to understand the laws' characteristic species of necessity. (That will be our aim in chapter 2.)

1.4. Nomic Preservation

Many examples suggest that laws and accidents stand in different relations to counterfactuals.[11] Had Jones missed his bus to work this morning, for instance, then every actual law of nature would still have held, but some of the actual accidents (such as that Jones always arrived at work on time) would not still have held. This example suggests that an accident's range of invariance under counterfactual suppositions is *narrower* than a law's—in other words, that for any law and any accident, the range of counterfactual suppositions under which the law is preserved wholly contains and goes somewhat beyond the range of counterfactual suppositions under which the accident is preserved. Accidents are thus more "delicate" than laws—more easily broken.

However, this thought is incorrect. Suppose a large collection of electrical wires, all of which are made of copper, are lying on a table. For the sake of the wires' utility, it is a good thing that copper is electrically conductive. Had copper been electrically insulating, then the wires on the table would not have been much good for conducting electricity. Now look at what just happened: under the counterfactual supposition that copper is an insulator, the law that all copper objects are electrically conductive obviously fails to be preserved. But the accident that all of the wires on the table are *made* of copper *is* preserved (at least in certain, easily imagined conversational contexts). To repeat: Had copper been electrically insulating, then the wires on the table, being made of copper, would have been useless for conducting electricity. So it is false that an accident's range of invariance under counterfactual suppositions is strictly narrower than a law's. There are counterfactual suppositions under which (in a given conversational context) a given accident is preserved but a given law is not.[12]

This example illustrates another important fact: although accidents may in some respect be more delicate than laws in having less resistance to being overthrown by counterfactual insults, a mere accident may nevertheless possess plenty of resilience. It is no law that all of the families living on our block have two children, yet this accident would still have held even if I had failed to brush my teeth this morning or worn an orange shirt today. Here is another accident possessing considerable persistence under counterfactual suppositions: whenever the gas pedal of my car is depressed by *x* inches and the car is on a dry, flat road, then the car's acceleration is given by the function *a*(*x*). Let's call this accident "*g*" for "*g*as pedal." (Of course, *g* is no coincidence; whenever the gas pedal is depressed, the same facts about the car's internal makeup help to explain its acceleration. But since those facts are accidents, *g* is accidental, too.) My knowledge that *g* would still have held, had the gas pedal on a certain occasion been pressed down a bit farther, has been relevant many times recently to the guidance I have given my daughter, Rebecca, in teaching her how to drive my car.[13]

Let's find a better way to capture the difference between laws and accidents in their persistence under counterfactual suppositions. The copper-wire example required a counterfactual supposition under which the law that all copper objects are electrically conductive fails to be preserved. I resorted to the brute-force approach: "Had copper been electrically insulating." It would have sufficed to use a counterfactual supposition about the electron-band structure of copper atoms or about the behavior of electrons or about the operation of electric fields. Nevertheless, each of these suppositions would have to be like my original brute-force supposition in being logically inconsistent with *some* natural law (even if logically consistent with the law that all copper objects are electrically conductive).[14] A counterfactual supposition must contradict some law or other in order for it to undercut the law that every copper object is electrically conductive. In contrast, for any accident, we can find a counterfactual supposition that is logically *consistent* with all of the laws, but under which that accident fails to be preserved. For instance, no law is violated by Bill Gates wanting to have a large gold cube built, yet under this supposition, the accidental generalization about gold cubes might not still have held. This suggests the following idea, which I call "Nomic Preservation" (NP):

> NP *m* is a law if and only if *m* would still have held under any counterfactual (or subjunctive) supposition *p* that is logically consistent with all of the laws (taken together).

In other words, m is a law exactly when $p \mathbin{\Box\!\!\rightarrow} m$ is true for any p that is logically consistent with the laws (taken all together). Nomic Preservation allows an accident to be invariant under a wide range of counterfactual suppositions—even under a supposition that contradicts laws. Yet NP still manages to distinguish laws from accidents.

However, NP requires several important refinements. They will occupy our attention for the rest of this section.

Let's start with an easy one. In a given conversational context, only certain counterfactual antecedents are relevant, considering our interests there. For example, suppose that several emergency room physicians are discussing whether the victim of a motor vehicle accident, who has just died, might have survived under various counterfactual circumstances: had she suffered only certain injuries without others, or had she been wearing a seat belt, or had the ambulance arrived at the accident scene sooner. In that context, presumably, counterfactual antecedents such as "Had human evolutionary history been different so that our vital organs were arranged differently" or "Had the human aorta been constructed out of steel" are irrelevant. The physicians in that context are concerned with human anatomy as it actually is, not as it might have been.

A counterfactual conditional with an antecedent that is irrelevant in a given context is perhaps neither true nor false in that context. Therefore, even if m is a law and p is logically consistent with all of the laws (taken together), it may be that $p \mathbin{\Box\!\!\rightarrow} m$ is not true in a given context because p is irrelevant there.[15] NP will have to be refined to leave room for this possibility.

Furthermore, even if m is an accident, it may be that *in a given context,* $p \mathbin{\Box\!\!\rightarrow} m$ holds for all counterfactual antecedents p that are relevant in that context and logically consistent with the laws. For example, suppose that I have just driven from Chapel Hill to Myrtle Beach in order to meet you, but I have arrived 30 minutes late. We discuss whether I would (or at least might) have arrived on time had I departed Chapel Hill an hour earlier, or had I taken U.S. Highway 15 instead of Interstate Highway 95, or had there been no accident to slow traffic on I-95, and so forth. You might ask whether I would have arrived any earlier had I turned left at a given intersection, and we might conclude that had I turned left, then I would have arrived even later because I would then have entered the ramp onto I-95 north (away from Myrtle

Beach) rather than south (toward Myrtle Beach). In this familiar sort of conversation, a counterfactual antecedent such as "Had Myrtle Beach been 100 miles nearer to Chapel Hill" is irrelevant. In this context, under every *relevant* counterfactual antecedent, the locations of Chapel Hill and Myrtle Beach and the routes taken by various highways are preserved. Our concern in this context is how I might have arrived on time, given the actual distances to be traveled and the roads actually available. Nevertheless, the locations of Chapel Hill and Myrtle Beach and the routes taken by various roads are accidents, not laws.

Their accidental character is reflected in the fact that there are *other* contexts where these facts are *not* preserved under counterfactual antecedents that are relevant there and logically consistent with the laws. For example, there are contexts where "Had the fall line been 150 miles farther inland in South Carolina, then I-95 would have been constructed farther from the coast there" is true, not to mention contexts where "Had the earth been only 40 million miles from the sun, then human beings would never have evolved, and so I-95 would not have been constructed" is true.

Although NP should allow an accident to behave like a law in *certain* contexts, NP should require that for any accident, there is *some* context where it does *not* behave like a law. Thus:

> *m* is a law if and only if for any conversational context, and for any *p* that is relevant as a counterfactual antecedent in that context and logically consistent with all of the laws (taken together), the proposition expressed in that context by "$p \square\!\!\rightarrow m$" is true.[16]

Let's now look at another way in which NP must be refined. Even if NP succeeds in distinguishing laws from accidents, NP fails to distinguish laws from broadly logical truths. If the laws are preserved under every counterfactual supposition in a given range, then the broadly logical truths (which have an even stronger variety of necessity than the laws) may well be preserved there, too.

However, although the broadly logical truths are not *merely* natural laws, they are *at least* natural laws. They possess whatever variety of necessity the laws possess and more. So they are "by courtesy" counted among the laws. After all, if it is a law that burning hydrogen in oxygen yields only water, and it is a law that burning hydrogen in oxygen yields only H_2O, then (if every logical consequence of laws is a law) it is a law that water is H_2O—although (some philosophers say) this is a metaphysical truth, reflecting water's *essence*.[17] Likewise, if it is a law that the speed of light is 3×10^8 meters per second and a law that the speed of light is half of 6×10^8 meters per second, then (if every logical consequence of laws is a law) it must be a law that $3 \times 10^8 = ½ \times 6 \times 10^8$, although this is a mathematical truth. NP can succeed in distinguishing laws from accidents only if the broadly logical truths qualify as laws.

There are motives for denying that every logical consequence of laws is a law. For instance, although it is a law that all emeralds are green and a law that all rubies are red, is it really a law that all things that are emeralds or rubies are green or red? (Presumably the reason why the stone in the King's ring is green or red is not because it is a ruby or an emerald, but rather because it is a ruby; that is why it is red, and hence green or red.)[18] A logical consequence of laws that is not itself a law will be preserved under every counterfactual supposition under which the laws entailing it are preserved. So no principle like NP can distinguish the laws proper from any of their logical consequences that are not laws. All of the laws' logical consequences hold "as a matter of law" even if not all of them are laws, narrowly speaking.

For the sake of simplicity, let's stipulate that as our default for the rest of this book, we shall interpret "natural law" expansively so that it includes all of the broadly logical truths as well as all of the logical consequences of laws. Over the course of this chapter and the next, we will see how the broadly logical truths differ from the other laws of nature by their greater invariance under counterfactual suppositions. On the other hand, I will have nothing to say here about how laws like "All emeralds are green" and "All rubies are red" differ from laws like "All things that are emeralds or rubies are green or red."

Now let's turn to the counterfactual antecedents under which NP demands that the laws be preserved: every counterfactual supposition p that is logically consistent with all of the laws. The antecedent must be logically consistent with all of the laws *taken together,* not merely with each law individually. In other words, there must be a logically possible world where p and all of the laws hold.

But for the laws to *hold* there, is it enough that every law (that is, every n where it is a law that n) is *true* there, or must they also be *laws* there? Which of these interpretations we place on NP can make a big difference. Suppose *p* is that *it is not a law* that energy is conserved. Obviously *p* is not logically consistent with the conjunction of all truths of the form "It is a law that *n*," since one of these truths is that it is a law that energy is conserved. But *p* is logically consistent with the bare fact that energy is conserved (in other words, with the fact that the total quantity of energy is the same at every moment) together with the truth of every other law. In a possible world where every actual law is true but *p* is also true, the total quantity of energy remains the same at every moment *as a matter of accidental fact.*

However, although this is *a* possible world where *p* is true, this is not the *closest* possible world where *p* is true—by which I mean simply that this is not what would have happened, had *p* been true. Rather, energy would *not* still have been conserved, had *p* been true (that is, had the laws not made energy conservation mandatory). Had there been no law prohibiting perpetual-motion machines, then scientists might well have built one by now. By the same token, had the laws not imposed 3×10^8 meters per second as a cosmic speed limit, there presumably would have been bodies accelerated from rest to beyond that speed—perhaps by the Stanford Linear Accelerator cranked up to full power. So NP is false if it demands that every law would still have been true under any counterfactual supposition that is logically consistent with the *truth* (though perhaps not with the *lawhood*) of all of the laws.

One way to avoid this problem is to refine NP so that it covers only counterfactual suppositions *p* where *p* is "sub-nomic," thereby excluding "Had energy conservation not been *a law*."[19] To explain this approach, let's start with the facts that we are trying to partition into laws and accidents.[20] Put aside all of those facts that could themselves hold at least partly by virtue of which facts are laws and which are not. The survivors include the fact that all emeralds are green (a law) as well as the fact that all gold cubes are smaller than a cubic mile (an accident)—but not the facts that it is *a law* that all emeralds are green, that it is *not a law* that all gold cubes are smaller than a cubic mile, and that no *laws* privilege any particular moment or location. Call the survivors the "sub-nomic facts," and let the "sub-nomic claims" be the claims that in any possible world are true (or false) exclusively by virtue of the sub-nomic facts there—not by virtue of which facts are laws and which are not. In other words, a claim is "sub-nomic" exactly when in any possible world, what makes it true (or false) there does not include which facts there are laws there and which are not. Note that as I shall use the term "sub-nomic," a sub-nomic claim, such as "Like charges repel" (or "All gold cubes are smaller than a cubic mile"), can nevertheless qualify as a law (or an accident).[21] That like charges repel is sub-nomic and a law, though the *fact that it is a law* is not a sub-nomic fact.

Let me also stipulate that the definition of "sub-nomic" treat the other species of modality in the same way as it treats natural modality. For example, just as "It is a law that like charges repel" is not sub-nomic, so likewise "It is a (broadly) logical necessity that all triangles have three sides" is not, but "All triangles have three sides" is. Sub-nomic claims do not contain such modal terms as "actually" or implicit references to other "possible worlds."[22]

The sub-nomic facts are, roughly speaking, the facts that laws might "govern" but that cannot themselves concern the composition of the "government." Accordingly, I take the sub-nomic facts to include facts about single-case objective chances, since science treats these facts as governed by laws just like facts about, say, the distribution of electric charge. For example, it is a sub-nomic fact that every atom of polonium-210, at each moment it exists, has a 50% chance of surviving for the next 138.39 days (the isotope's half-life). But the fact that this is a law is not a sub-nomic fact. It is a sub-nomic fact that every atom now in a given vial (labeled "polonium-210") has a 50% chance of surviving for the next 138.39 days. But the fact that this is an accident is not a sub-nomic fact.[23]

The sub-nomic facts lie at the base of a hierarchy of facts, where the facts on a given rung of the hierarchy are "governed" by some of the facts one rung higher. A given rung of the hierarchy consists of the facts (at least partly) about what's governing the facts on the rung immediately below (see fig. 1.1). Let's start three rungs up in the hierarchy. The facts there include that it is a law that the laws one rung below (the "first-order laws") exhibit various symmetries: roughly speaking, that they privilege no locations or moments, that they are the same in every reference frame in a certain family (as demanded by Einstein's "principle of relativity"), and so forth. The facts on this rung, then, specify what laws govern the first-order laws. (Such "meta-laws" will be discussed in chapter 3.) One rung below, on the second rung of the hierarchy, are the facts that the meta-laws govern, such as the facts specifying the laws that exhibit those symmetries (for example, the fact that it is a law that all electrons have negative electric charge and the fact that it is not a law that all gold cubes are smaller than a cubic mile). These facts are concerned not with what governs other *laws,* but rather with what governs the facts on the rung just below, which is the lowest rung of the hierarchy. The facts on the bottom rung are not about what governs some other facts, since there is nothing lower in the hierarchy to govern. These facts are governed but govern nothing themselves. They are the sub-nomic facts. The laws governing the sub-nomic facts are insufficient to determine all of the sub-nomic facts, since they do not entail the "initial conditions" (and even the laws together with the initial conditions fail to determine all of the sub-nomic facts, if the laws are statistical). Likewise the meta-laws are insufficient to determine the first-order laws; they merely impose some constraints on them.

That all like charges repel is a sub-nomic fact, whereas that it is a law that all like charges repel is not sub-nomic; it belongs on the second rung. Likewise, that no locations or moments are privileged by first-order laws belongs on the second rung, whereas that it is a law

Figure 1.1 The sub-nomic facts (tiny figure at extreme right) being bossed around by the laws governing them (second from right), which in turn are being constrained by the meta-laws (middle). I am not prepared to say for how many levels this hierarchy actually extends; that is for science to discover. The sub-nomic facts have no facts to boss around.

that first-order laws are so symmetric belongs on the third rung. That all first-order laws possess a certain feature may describe the facts governing the sub-nomic facts without itself explaining any sub-nomic fact (unlike the fact that it is a law that all like charges repel, which explains why in fact all like charges repel).

Let's now use the notion of "sub-nomic claims" to refine NP. If we take NP as concerned only with counterfactual suppositions *p* where *p* is sub-nomic, then we exclude from consideration the counterfactual suppositions that were causing trouble for NP, such as "Had energy conservation not been *a law*."

To reduce clutter, let's henceforth reserve lower-case italicized English letters (such as *p* and *m*) for sub-nomic claims and leave mostly tacit the various other qualifications that I have just introduced. So we have arrived at

> NP   *m* is a law if and only if in any context, $p \;\Box\!\!\rightarrow m$ holds for any *p* that is logically consistent with all of the *n*'s (taken together) where it is a law that *n* (that is to say, for any *p* that is logically consistent with the first-order laws).

I endorse this principle.

1.5. Beyond Nomic Preservation

Although NP tells us something important about how laws differ from accidents in their relation to counterfactuals, we will see in the next two sections that we can go much further: by generalizing NP, we will see how the laws can be separated from the accidents *solely* by which counterfactuals are true. Before turning to that challenge, however, I want to identify a few further ideas that are suggested by the same sorts of thoughts that motivated NP. These ideas will prove fruitful later in this chapter. But I shall not add them officially to NP. It will be more convenient to reserve "NP" for the principle that I have just given.

The idea that some truth *m* is "preserved" under a given counterfactual supposition *p* was supposed to be captured by the fact that $p \;\Box\!\!\rightarrow m$—in other words, by the fact that *m* would still have been true, had *p* been the case. But what if not only *m,* but also ~*m* (that is, *m*'s negation) would have held, had *p* been the case? That's not really what we had in mind by *m*'s "preservation"! Of course, not only would energy still have been conserved, had I tried to build a perpetual-motion machine (that is: $p \,\Box\!\!\rightarrow m$), but also it is *not* the case that had I tried to build a perpetual-motion machine, then energy would *not* have been conserved (that is: $\sim (p \,\Box\!\!\rightarrow \sim m)$). However, there may be more exotic counterfactual suppositions under which (at least in some conversational contexts) both *m* and ~*m* obtain, such as "Had there been a round square" or "Had some object been entirely made of rubber and also entirely made of copper." For ~*m* to hold along with *m* under some counterfactual supposition would be a disappointingly cheap way for *m* to be "preserved" there. Here is a means of capturing the idea that under various counterfactual suppositions *p,* the laws are preserved, but *not* in this cheap way:

> *m* is a law if and only if in any context, $p \,\Box\!\!\rightarrow m$ and $\sim (p \,\Box\!\!\rightarrow \sim m)$ hold for any *p* that is logically consistent with all of the *n*'s (taken together) where it is a law that *n*.

We can shorten this principle a bit by replacing the *would*-conditionals in it with a *might*-conditional. I have already introduced some might-conditionals. For instance, I said that had Bill Gates wanted to build a large gold cube (*p*), then there might have been a gold cube exceeding a cubic mile (*m*)—symbolically: $p \,\Diamond\!\!\rightarrow m$. There are two important connections between might-conditionals and would-conditionals. Firstly, if it is not the case that ~*m* might have held, had *p* held, then *m* would have held, had *p* held. In other words, $\sim (p \,\Diamond\!\!\rightarrow \sim m)$ logically entails $(p \,\Box\!\!\rightarrow m)$. Secondly, if it is not the case that ~*m* might have held, had *p* held, then it is not the case that ~*m* would have held, had *p* held. In other words, $\sim (p \,\Diamond\!\!\rightarrow \sim m)$ logically entails $\sim (p \,\Box\!\!\rightarrow \sim m)$.[24] By virtue of these two might-would connections, both $p \,\Box\!\!\rightarrow m$ and $\sim (p \,\Box\!\!\rightarrow \sim m)$ in the indented principle above follow from $\sim (p \,\Diamond\!\!\rightarrow \sim m)$. Perhaps, then, we should consider the principle

> *m* is a law if and only if in any context, $\sim (p \,\Diamond\!\!\rightarrow \sim m)$ holds for any *p* that is logically consistent with all of the *n*'s (taken together) where it is a law that *n*.

This looks like a good way to capture the laws' preservation.

Let me set this principle aside temporarily and turn to another idea suggested by the same sort of examples that motivated NP. As

I mentioned earlier, had the laws not required that energy be conserved, then energy might not have been conserved. This seems closely related to the thought that the reason why energy is in fact conserved is because its conservation is required by law; energy is conserved because it is a law that energy is conserved. (Here we have a covering-law explanation: That *m* is a law explains why *m* is the case by making *m* inevitable, unavoidable—necessary.) Now according to NP, energy would still have been conserved had *p* been the case, as long as *p* is logically consistent with the first-order laws. In the "closest *p*-world," then, why is it the case that energy is conserved? What is the scientific explanation there for the fact that energy is conserved? Presumably, the closest *p*-world is like the actual world in that energy is conserved there not by accident, but because a law requires it: energy conservation is a law in that world, too. If it is not the case that had someone tried to build a perpetual-motion machine, energy conservation would still have been a *law,* then why is it the case that had someone tried to build a perpetual-motion machine, energy would still have been conserved? Without the principle of energy conservation remaining a *law* under this counterfactual supposition, there is no need for the principle to remain *true* under that supposition.

That the laws of nature would have been no different, had Jones missed his bus to work this morning or Bill Gates wanted a large gold cube built, seems intuitively plausible and can be captured by this principle:

> *m* is a law if and only if in any context, "Had *p* been the case, then *m* would have been a law" holds for any *p* that is logically consistent with all of the *n*'s (taken together) where it is a law that *n*.

Let's now explore one step further. If it is actually a law that *m,* then according to the above principle, $p \mathrel{\Box\!\!\rightarrow} (m \text{ is a law})$, and by NP, (*m* is a law) entails that $q \mathrel{\Box\!\!\rightarrow} m$—as long as *p* is logically consistent with the first-order laws, and *q* is likewise. Hence, if *m* is a law, then $p \mathrel{\Box\!\!\rightarrow} (q \mathrel{\Box\!\!\rightarrow} m)$ for any such *p* and *q,* and likewise no matter how many layers of counterfactuals are nested (or "embedded").

Please do not confuse the nested counterfactual $p \mathrel{\Box\!\!\rightarrow} (q \mathrel{\Box\!\!\rightarrow} m)$ with $(p\&q) \mathrel{\Box\!\!\rightarrow} m$. Their difference is especially evident when *p* and *q* are broadly logically inconsistent. For example, consider the nested counterfactual "Had the object been entirely made of rubber, then here's a counterfactual conditional that would have been true: had it been entirely made of copper, it would have been electrically conductive." That is true but plainly not equivalent to "Had the object been entirely made of rubber and been entirely made of copper, then it would have been electrically conductive." The nested counterfactual $p \Box\!\!\rightarrow (q \Box\!\!\rightarrow m)$ is not equivalent to $(p\&q) \Box\!\!\rightarrow m$ even when p is broadly logically consistent with q. For example, suppose that you and I run a race, I win, and I would always win were I to try. Had you won, then had I tried, I would have won. This nested counterfactual is obviously not equivalent to the false conditional "Had you won and I tried, then I would have won."

That the laws would still have been true, even under nested counterfactual antecedents, can be captured by the principle:

> m is a law if and only if in any context, $p \Box\!\!\rightarrow m$ holds, $p \Box\!\!\rightarrow (q \Box\!\!\rightarrow m)$ holds, and so forth, as long as p is logically consistent with all of the n's (taken together) where it is a law that n, q is likewise, and so forth.

For instance, we believe that had Jones missed his bus to work this morning, then the natural laws would still have been exactly the actual laws and so (by NP) if Jones, after missing his bus, had done nothing about getting to work but click his heels and make a wish to materialize instantly at his office, he would not have gotten to work. That was a nested counterfactual that just went by. It concerned whether a given counterfactual conditional ("Had Jones done nothing but click his heels and make a wish to materialize instantly at his office, he would have gotten to work") would have been true under a certain counterfactual supposition ("Had Jones missed his bus to work this morning"). This nested counterfactual is covered by our latest principle.

Likewise, had Bill Gates tried to accelerate a body beyond the speed of light, then he would have failed, and moreover (here comes the nested counterfactual) had he access to 23rd-century technology, then had he tried to accelerate a body beyond the speed of light, he would still have failed. (Even if he had access to 23rd-century technology, it would still have been a law that no body is accelerated to beyond light speed.) On the other hand, it is merely an accident that all hurricanes rotate counterclockwise in the Northern Hemisphere. This regularity's accidental character is reflected in the truth of this nested counterfactual: had the

Earth rotated westward instead of eastward, then had there now existed some hurricanes in the Northern Hemisphere, they would all have been rotating clockwise.[25]

A bit earlier, I suggested that to require that the laws be preserved, but not in a cheap way, we can demand not only that $p \mathbin{\Box\!\!\rightarrow} m$, but also that $\sim (p \mathbin{\Box\!\!\rightarrow} {\sim}m)$, where both of these counterfactuals follow from $\sim (p \mathbin{\Diamond\!\!\rightarrow} {\sim}m)$. The same argument applies to preservation principles that include nested counterfactuals. If the preservation principle requires that $q \mathbin{\Box\!\!\rightarrow} (p \mathbin{\Box\!\!\rightarrow} m)$, then in order to rule out the possibility that $(p \mathbin{\Box\!\!\rightarrow} m)$'s preservation under q is of the cheap kind, the principle should also require that it *not* be the case that $(p \mathbin{\Box\!\!\rightarrow} m)$ might have been lost had q obtained. In other words, it should require that $\sim (q \mathbin{\Diamond\!\!\rightarrow} \sim (p \mathbin{\Box\!\!\rightarrow} m))$ hold. Now one of the connections we saw between might-conditionals and would-conditionals was that $\sim (p \mathbin{\Diamond\!\!\rightarrow} {\sim}m)$ logically entails $(p \mathbin{\Box\!\!\rightarrow} m)$. In other words, $\sim (p \mathbin{\Box\!\!\rightarrow} m)$ logically entails $(p \mathbin{\Diamond\!\!\rightarrow} {\sim}m)$. Therefore, if it is not the case that $(p \mathbin{\Diamond\!\!\rightarrow} {\sim}m)$ might have held, then it is not the case that $\sim (p \mathbin{\Box\!\!\rightarrow} m)$ might have held. Hence, our new requirement's $\sim (q \mathbin{\Diamond\!\!\rightarrow} \sim (p \mathbin{\Box\!\!\rightarrow} m))$— in other words, that it is not the case that $\sim (p \mathbin{\Box\!\!\rightarrow} m)$ might have held, had q held—follows from $\sim (q \mathbin{\Diamond\!\!\rightarrow} (p \mathbin{\Diamond\!\!\rightarrow} {\sim}m))$.

This last expression has other convenient implications as well. By another application of the same might-would connection, $\sim (q \mathbin{\Diamond\!\!\rightarrow} (p \mathbin{\Diamond\!\!\rightarrow} {\sim}m))$ entails $(q \mathbin{\Box\!\!\rightarrow} \sim (p \mathbin{\Diamond\!\!\rightarrow} {\sim}m))$, which, by one final application of the same might-would connection, entails $(q \mathbin{\Box\!\!\rightarrow} (p \mathbin{\Box\!\!\rightarrow} m))$. That was the first nested counterfactual that we incorporated into a preservation principle.

What we have just seen, then, is that $\sim (q \mathbin{\Diamond\!\!\rightarrow} (p \mathbin{\Diamond\!\!\rightarrow} {\sim}m))$ entails every other result involving nested counterfactuals that we wanted to include in a preservation principle. Expressions like $\sim (q \mathbin{\Diamond\!\!\rightarrow} (p \mathbin{\Diamond\!\!\rightarrow} {\sim}m))$ are all we need to use in order to construct a preservation principle that is powerful enough to encompass all of the various conditionals we want:

> *m* is a law if and only if in any context,
> $\sim (p \mathbin{\Diamond\!\!\rightarrow} {\sim}m)$,
> $\sim (q \mathbin{\Diamond\!\!\rightarrow} (p \mathbin{\Diamond\!\!\rightarrow} {\sim}m))$,
> $\sim (r \mathbin{\Diamond\!\!\rightarrow} (q \mathbin{\Diamond\!\!\rightarrow} (p \mathbin{\Diamond\!\!\rightarrow} {\sim}m))$,. . .
> all hold, as long as *p* is logically consistent with all of the *n*'s (taken together) where it is a law that *n*, *q* is likewise, *r* is likewise, and so forth.

To keep things simpler, I have refrained from building these additional details officially into NP. But the principle at which we have just arrived succeeds in cashing out some further ideas that are suggested by the same thoughts that motivated NP. These ideas will prove valuable shortly.

1.6. A Host of Related Problems: Triviality, Circularity, Arbitrariness

Let's return to NP:

> NP *m* is a law if and only if in any context, $p \; \Box\!\!\rightarrow m$ holds for any *p* that is logically consistent with all of the *n*'s (taken together) where it is a law that *n*.

NP and the other principles I have just mentioned accord well with our routine scientific practice of using the natural laws to ascertain what would have happened under various hypothetical circumstances. Intuitively, the laws supply a control panel of knobs for setting the universe's initial conditions (or any system's boundary conditions), and these knobs can be twisted (hypothetically!) in any fashion whatsoever that is logically consistent with the laws. No matter the setting to which the knobs are turned (counterfactually) within these generous limits, the actual laws would still have held.[26]

One entertaining example of knob-twisting takes place in Arthur Upgren's *Many Skies: Alternative Histories of the Sun, Moon, Planets, and Stars* (2005). An astronomy professor at Wesleyan and Yale, Upgren explains that in thinking about what the world would have been like under various counterfactual circumstances (such as had the solar system been located in a star cluster), "I have . . . not changed the laws of physics." Upgren takes what we believe the laws to be and extrapolates from them to the conditions that would have prevailed under various counterfactual circumstances.[27]

Principles like NP (though without some of the qualifications and elaborations that I have introduced) have been advanced by a host of philosophers.[28] They have also been contested by some philosophers. I relegate further discussion of these objections to an endnote[29] because we have an even bigger problem to worry about. Despite all of our work refining NP, we must face the fact that even if NP is true, it cannot reveal how the laws are set apart from the accidents by their relation to counterfactual conditionals.

One source of concern is that in one direction, NP is trivial. It is obvious that no accident would still have been true under every *p* that is logically consistent with the first-order laws, since if *m* is an accident, one such *p* is ~*m,* and ~*m* □→ *m* is plainly false (at least when ~*m* is broadly logically possible). The trouble with NP's truth being trivial in this direction is that we might have expected "All gold cubes are smaller than one cubic mile" to reveal its accidental character not by failing to be preserved under "Had there been a gold cube larger than a cubic mile" (its failure to be preserved under *that* antecedent is surely not its fault!) but rather by failing to be preserved under a counterfactual supposition with which it is logically consistent, such as "Had Bill Gates wanted a large gold cube to be constructed."

NP is trivial in one direction because the range of counterfactual suppositions falling under NP encompasses exactly the *p*'s that are logically consistent with the first-order laws. So for each accident *m,* the range includes ~*m,* whereas for each law *m,* the range does not include ~*m.* This bias toward the laws takes us to the heart of the problem with NP: it gives preferential treatment to the laws, allowing them to determine which counterfactual suppositions get to be considered. This amounts to the laws stacking the deck in favor of themselves.

Let me explain NP's "circularity" a little more carefully. NP uses the laws to pick out the range of counterfactual suppositions that, in turn, it uses to pick out the laws. This means that NP cannot distinguish the laws from the accidents solely on the basis of the truth-values (in all conversational contexts) of all of the counterfactual conditionals *p* □→ *m.* Rather, for NP together with the truth-values of those counterfactuals to reveal which sub-nomic facts are laws, the laws must first be distinguished on some *other* grounds, thereby picking out the range of counterfactual suppositions *p* that are logically consistent with the laws. Only then can NP pick out the laws among the sub-nomic facts as exactly the sub-nomic facts that are preserved under all of those suppositions.

But that's not all. Even if NP is true, NP cannot explain why the laws' distinctive sort of persistence under counterfactual suppositions makes the laws *special* or *important;* it cannot reveal the sense in which the laws bear an especially intimate relation to counterfactuals. Once again, the source of the trouble is that the concept of natural law appears in NP on both sides of the "if and only if"; NP allows the laws themselves to delimit the range of counterfactual perturbations under which a fact must be invariant in order for it to qualify as a law. Therefore, NP can portray the laws as special, in virtue of their invariance under this range of counterfactual suppositions, only if this particular range of counterfactual suppositions is itself special already. But this range is distinguished only by its ranging over exactly those suppositions that are logically consistent with all of the laws. Hence, NP can reveal what makes the laws special, as far as their invariance under counterfactual suppositions is concerned, only if there is some independent reason why the laws are special. Unless there is already some reason why this particular range of counterfactual suppositions is special, the laws' invariance under this range fails to reveal anything special about the laws.

NP's circularity is closely related to NP's triviality as far as accidents are concerned. NP permits the gold cubes accident to be invariant under a wide range of counterfactual suppositions. NP insists only that there be some counterfactual supposition *p,* logically consistent with every law, under which the gold cubes regularity would not still have held. Plainly, there is: Had there been a gold cube larger than a cubic mile! NP regards the failure of the gold cubes regularity to be preserved under this *p* as showing that the regularity is not a law, but only because this *p* is logically consistent with the laws—that is, only because the gold cubes regularity is not a law! There's the circle again.

It seems arbitrary to privilege the counterfactual suppositions that are logically consistent with the laws. We could, it seems, just as well have privileged the counterfactual suppositions that are logically consistent with, say, the fact that George Washington was the first president of the United States. But the facts that would still have held, under every one of *those* counterfactual suppositions, should not merit our attention in the way that the laws of nature should. (They have no special explanatory power, for instance.) What is so noteworthy, then, about preservation under one range of counterfactual suppositions as opposed to some other range? Once again, although NP may be true, it leaves unexplained why the laws are especially significant. For NP to tell us why the laws are special, we would have to know already what is special about being invariant under every counterfactual supposition that is logically consistent with (wait for it!) the laws, and so we would have to know already why the laws are special.

A few philosophers have mentioned a problem along these lines (though without necessarily elaborating it in terms of triviality, circularity, or arbitrariness) and despaired of ever resolving it.[30] I shall resolve it now. I shall spend the rest of this book trying to squeeze every last drop of philosophical juice out of this single move. So it had better be good!

1.7. Sub-nomic Stability

NP's problems arose from its giving special treatment to a range of counterfactual suppositions designed expressly to suit the laws. What if the same courtesy that NP gives the laws were extended to every set of sub-nomic truths? NP says that the laws would still have held under every counterfactual supposition that is logically consistent with the laws. So let's consider whether some set of sub-nomic truths containing accidents would still have held under every counterfactual supposition that is logically consistent with that set. We thus avoid arbitrarily privileging the range of counterfactual suppositions that is logically consistent with the laws. Rather than allowing the laws to dictate to any set of truths the range of counterfactual suppositions under which that set's invariance is to be assessed, let's allow each set to pick out for itself a range of counterfactual suppositions that it finds comfortable.

For example, take the set containing exactly the sub-nomic claims that are logical consequences of "All gold cubes are smaller than one cubic mile." Are this set's members all preserved under every sub-nomic counterfactual supposition that is logically consistent with all of them (taken together)? Of course, the set's members are *not* all preserved under the counterfactual supposition "Had there existed a gold cube larger than a cubic mile." But that supposition is not logically consistent with the set. To see that the set's members are not all preserved under every counterfactual supposition that is logically consistent with the set, notice that they are not all preserved under the supposition that Bill Gates wants a large gold cube to be constructed. Although that supposition *is* logically consistent with the set (and so its members *could* all still have held under it), they *wouldn't* all still have held under it.

This approach, I shall argue, allows us to distinguish the set of first-order laws (that is, the set containing exactly the sub-nomic truths *m* where it is a law that *m*) from any set of sub-nomic truths that contains some (but not all) of the accidents—while avoiding the problems of triviality, circularity, and arbitrariness afflicting NP. I will call the key property "sub-nomic stability." Roughly speaking, a set of sub-nomic truths is "sub-nomically stable" if and only if whatever the conversational context, the set's members would all still have held under every sub-nomic counterfactual (or subjunctive) supposition that is logically consistent with the set—even under however many such suppositions are nested. In other words, for any member *m* of the set, and for any sub-nomic suppositions *p, q, r*, . . ., each of which is logically consistent with the set, all of the counterfactuals $p \mathbin{\Box\!\!\rightarrow} m$, $q \mathbin{\Box\!\!\rightarrow} (p \mathbin{\Box\!\!\rightarrow} m)$, $r \mathbin{\Box\!\!\rightarrow} (q \mathbin{\Box\!\!\rightarrow} (p \mathbin{\Box\!\!\rightarrow} m))$ . . . hold in every context. Moreover, as we saw in section 1.5, we should preclude *cheap* preservation of the set's members, which we can do by also requiring that it not be the case that their negations might have held under these suppositions—for instance, by requiring not only that $p \mathbin{\Box\!\!\rightarrow} m$, but also that $\sim (p \mathbin{\Diamond\!\!\rightarrow} \sim m)$. We found that nested might-conditionals entail all of the conditionals we need. I will now put those nested might-conditionals to use in defining "sub-nomic stability":

> Consider a nonempty set Γ of sub-nomic truths containing every sub-nomic logical consequence of its members. Γ possesses *sub-nomic stability* if and only if for each member *m* of Γ (and in every conversational context),
>
> $\sim (p \mathbin{\Diamond\!\!\rightarrow} \sim m)$,
> $\sim (q \mathbin{\Diamond\!\!\rightarrow} (p \mathbin{\Diamond\!\!\rightarrow} \sim m))$,
> $\sim (r \mathbin{\Diamond\!\!\rightarrow} (q \mathbin{\Diamond\!\!\rightarrow} (p \mathbin{\Diamond\!\!\rightarrow} \sim m)))$,. . .
>
> for any sub-nomic claims *p, q, r*, . . . where $\Gamma \cup \{p\}$ is logically consistent, $\Gamma \cup \{q\}$ is logically consistent, $\Gamma \cup \{r\}$ is logically consistent,. . . .

The motivations behind NP suggest that the set of first-order laws is sub-nomically stable. I shall call this set "Λ"—that is, "lambda" (for "law"). In contrast, the set containing exactly the sub-nomic logical consequences of the gold cubes generalization fails to be sub-nomically stable (since some members of the set might have been false, had Bill Gates wanted a large gold cube to be constructed).[31]

I shall argue that there is no sub-nomically stable set containing an accident—except perhaps the set of *all* sub-nomic truths. According to many proposed logics of counterfactuals, $p \,\square\!\!\rightarrow q$ is true trivially whenever $p\&q$ is true (a principle known as "Centering"), and likewise for nested counterfactuals. If Centering is correct, then each member of the set of all sub-nomic truths is trivially preserved under every sub-nomic supposition p that is true. There are no sub-nomic suppositions p that are false and logically consistent with the set. (If p is a false sub-nomic supposition, then $\sim p$ is a member of the set.) In that case, the set of all sub-nomic truths trivially possesses sub-nomic stability. Accordingly, I will argue that Λ is the largest *nonmaximal* set that is sub-nomically stable. (On the other hand, if Centering is false, then even the set of all sub-nomic truths may lack sub-nomic stability. In fact, I believe that Centering fails in a universe where there are objective chances (other than 0 and 1), but perhaps it holds in a universe where there aren't.[32] But the truth of Centering need not concern us now; our focus is on the proposal that laws differ from accidents in belonging to a sub-nomically stable set that does not contain every sub-nomic truth.)

This proposal for distinguishing laws from accidents avoids the circularity afflicting the idea that the laws are the truths that would still have held under every counterfactual supposition that's logically consistent with the laws. Sub-nomic stability does not start by giving special privileges to the laws. It is very egalitarian; it does not grant the laws the right to dictate to every set the range of counterfactual suppositions under which that set's invariance is to be tested. Stability thus has the potential to be a genuinely special feature of the laws. Whether sub-nomic stability can realize this potential is the subject of the rest of this chapter (and indeed this book).

Let me allay one concern that you may have at this point. I have asked you to think about whether various counterfactuals are true. In trying to evaluate those counterfactuals, you may have found yourself thinking about what the laws of nature are. For instance, in considering whether there might have been a large gold cube had Bill Gates wanted one to be constructed, you may have said to yourself, "Well, I guess there might then have been a large gold cube, since it is just an accident that all gold cubes are smaller than a cubic mile." Likewise, in thinking about whether there might have been a perpetual-motion machine had Bill Gates wanted one to be constructed, you might have said to yourself, "No, there wouldn't have been, even if Bill had wanted one, since the laws of nature prohibit perpetual-motion machines." Accordingly, you may be worried that insofar as we are using our beliefs about the laws to ascertain which counterfactual conditionals are true, it is problematic for us to turn around and appeal to the truth or falsehood of various counterfactuals in ascertaining how laws differ from accidents in their relation to counterfactuals.

However, I am not trying to discover whether some fact is a law by consulting various counterfactuals that I know to be true only by already knowing whether that fact is a law. Rather, various truths are laws (we believe), and various counterfactual conditionals hold (we believe), and I am trying to figure out the relation between these two sets of facts. Since these facts are closely related, it is entirely to be expected that we will sometimes consult our beliefs about which truths are laws in order to arrive at our beliefs about which counterfactuals are true, just as we may sometimes use our beliefs about various counterfactuals to arrive at beliefs about the laws. ("It can't be a law that every family on my block has exactly two children, because the Jones family could have moved onto our block without violating any law of nature, and they would then still have had three children.") We have been looking at various proposals regarding the laws' special relation to counterfactuals, and we have been testing those proposals partly by seeing whether they fit our beliefs about which truths are laws and which counterfactuals are true. These tests remain severe even if we sometimes draw upon our beliefs about the laws in order to arrive at or to justify our beliefs about counterfactuals (or vice versa).

Admittedly, it would be hugely problematic if we took the laws as helping to *make* certain counterfactuals true while also taking the truth of those counterfactuals as helping to *make* certain facts qualify as laws. But I have not done that. In this chapter, I am concerned only with *identifying* the special relation between laws and counterfactuals, not with figuring out *why* this relation holds: whether laws are responsible for counterfactuals, or counterfactuals are responsible for laws, or neither is responsible for the other. We will grapple with those questions in chapters 2 and 4.

1.8. No Nonmaximal Set Containing Accidents Possesses Sub-nomic Stability

Let's see an argument for that bold assertion!

Take a set of sub-nomic truths containing every sub-nomic logical consequence of its members and including the fact that all gold cubes are smaller than a cubic mile. What would it take for this set to be sub-nomically stable (or just "stable," for short)? As we have seen, it is not the case (in every conversational context) that all of the set's members would still have been true had Bill Gates wanted to have a large (exceeding one cubic mile) gold cube built. How can the set be stable despite failing to be preserved under this counterfactual supposition? Stability requires the set's members all to be invariant under every sub-nomic counterfactual supposition that is logically consistent with them all (taken together). The only way for this set to be stable, despite failing to be preserved under the supposition "Bill Gates wants to have a large gold cube built," is for that supposition to be logically inconsistent with the set. To be stable, then, the set's members must logically entail "Bill Gates never wants to have a large gold cube built," so that the supposition that Bill Gates wants to have a large gold cube built is logically inconsistent with the set's members. Since the set contains every sub-nomic logical consequence of its members, the set must contain the fact that Bill Gates never wants to have a large gold cube built.

However, the set's stability is not yet assured. Presumably, had Melinda Gates wanted to own a large gold cube, then Bill (Melinda's husband) would have wanted one built. Hence, a member of the set ("Bill Gates never wants to have a large gold cube built") is not preserved under the supposition that Melinda Gates wants to own a large gold cube. How can the set be stable despite failing to be preserved under this supposition? The argument we gave a moment ago applies here as well. If the set includes the fact that Melinda Gates never wants to own a large gold cube, then the supposition that she wants to own one is logically inconsistent with the set, and so the set's failure to be preserved under this supposition does not preclude its stability. Therefore, if a stable set includes the fact that Bill Gates never wants to have a large gold cube built, then it must also include the fact that Melinda Gates never wants to own one.

With each additional claim p that is admitted into the set in order to keep the set's behavior under some supposition $\sim p$ from rendering it unstable, we must worry about the suppositions $\sim q$ under which p is not preserved. To prevent the set's behavior under those suppositions from rendering it unstable, further claims q must be admitted into the set. The process snowballs until the set contains every sub-nomic truth. Thus, no nonmaximal set containing accidents possesses stability.

For example, suppose the set omits the accident that all of the apples on my tree are ripe. Then the following counterfactual supposition is logically consistent with the set: had either some gold cube exceeded one cubic mile or some apple on my tree not been ripe. Under this supposition, there is no reason why the generalization about gold cubes (which is in the set) should take priority in every conversational context over the apple generalization (which we have supposed not to be in the set). So it is not the case that in every context, the gold cubes generalization is preserved under this counterfactual supposition. Hence, to be stable, the set must include the apple generalization (thereby rendering the supposition logically inconsistent with the set). Therefore, if the set is stable and includes one accident, then it must include every accident.

Let's confirm this conclusion by looking at another example. Return to the accidental truth g: whenever the gas pedal of a certain car is depressed by x inches and the car is on a dry, flat road, then the car's acceleration is given by $a(x)$. Had the pedal on a certain occasion been depressed a bit farther, then g would still have held. However, a set containing g is unstable unless it also includes a description of the car's engine, since g might not still have held had the engine contained six cylinders instead of four. With a description of the car's engine in the set, the set's failure to be preserved under "Had the engine contained six cylinders instead of four" does not compromise the set's stability. But now to be stable, the set must also include a description of the engine factory, since had the factory been different, the engine might have been different. By packing more and more into the set, will we ever achieve stability before the set contains every sub-nomic fact?

I do not think so. Take a logically closed set containing g but omitting the fact that Jones is not wearing an orange shirt. Now consider what the world would have been like, had either $\sim g$ held or Jones been wearing an orange shirt. Would g still have held? In every conversational context? Certainly not! In at least some contexts, *g* might still have held, but Jones might just as well still not have been wearing an orange shirt.[33] In those contexts, it is *not* the case that a given truth (*g*) *within* the set would still have held, had either it or an arbitrary truth *outside* of the set been false. That is enough to make the set unstable. To disarm this threat to the set's stability, we must ensure that the set contains the fact that Jones is not wearing an orange shirt; the threatening supposition (~*g* or Jones is wearing an orange shirt) is then logically inconsistent with the set, and so to be stable, the set has no need to be preserved under that supposition. But if a set containing *g* is rendered unstable by omitting even an arbitrary sub-nomic fact, then the set is unstable if it fails to include *every* sub-nomic fact.

The same sort of argument could be made regarding any set Γ of sub-nomic truths containing every sub-nomic logical consequence of its members and *some* accidents but not *all* of them. There exist two intuitively unrelated accidental truths, *p* and *q,* where Γ includes *p* but omits *q*. For Γ possibly to be stable, each member of Γ, such as *p,* must be invariant (in every conversational context) under the counterfactual supposition that either ~*p* or ~*q,* since this supposition is logically consistent with Γ. (If (~*p* or ~*q*) were inconsistent with Γ, then since Γ contains every sub-nomic logical consequence of its members, Γ would have to contain ~(~*p* or ~*q*), that is, (*p*&*q*), and since (once again) Γ contains every sub-nomic logical consequence of its members, Γ would have to contain *q,* which Γ was stipulated as omitting.) But it is not the case that (~*p* or ~*q*) □→ *p* holds in every context. In picturesque terms, the supposition (~*p* or ~*q*) pits *p* against *q,* as far as remaining true is concerned. They cannot both be preserved under this supposition; at least one must go. With *p* and *q* utterly unrelated and neither a law, it is not the case that *p* takes priority over *q* in every context, no matter which facts are salient there. Hence, Γ is unstable. Although the set has picked out a comfortable range of counterfactual suppositions, the set is not invariant under all of these suppositions.

The above argument made no appeal to nested counterfactuals, although the definition of "sub-nomic stability" requires that each member *m* of a stable set be preserved even under arbitrarily many nested suppositions, so that *p* □→ *m, q* □→ (*p* □→ *m*), *r* □→ (*q* □→ (*p* □→ *m*)),. . . all hold. These nested counterfactuals may seem remote from actual scientific practice. But in fact, scientists routinely employ nested counterfactuals ("Had the chamber been completely evacuated, then had a few CO_2 molecules been present, they would have had a long mean free path"; "Had gravity declined with the cube of the distance, then a solar system, had it begun with many planets, would not long have so remained; planets would have soon escaped or spiraled into the sun"). Although the argument I just gave shows how difficult it would be, I suppose it is barely possible for there to be a nonmaximal set of sub-nomic truths containing an accident where (in every context) each member *m* would still have held under any sub-nomic counterfactual supposition *p* that is logically consistent with the set. But if there is such a set, its invariance is just a fluke. That is, its invariance is not likewise invariant. (Or if it is, then *that* is just a fluke: *its* invariance is not likewise invariant. (Or. . . .)) In other words, although every counterfactual p □→ m requisite for stability may hold, there is some *q* that is logically consistent with the set where q □→ (p □→ m) fails (or one of the further nested counterfactuals fails). To ensure that no nonmaximal set containing accidents manages to possess sub-nomic stability, we need nested counterfactuals in the definition of "sub-nomic stability."

The argument that I have just given against the stability of any nonmaximal set Γ containing accidents cannot be used against Λ's stability. Although context wields great influence over counterfactuals, there is a limit to its influence: in no context does an accident *q* take priority over a law *p* under the counterfactual supposition (~*p* or ~*q*). In note 29, I look at several sorts of cases that initially might appear to violate Λ's sub-nomic stability, and I argue that all of them are best understood without denying Λ's stability. In addition to those arguments, here is a general way of thinking about why context cannot enable an accident *q* to take priority over a law *p* under (~*p* or ~*q*).

Consider any ordinary counterfactual conditional in a context where it is true. For example,

> Had my family gone out to dinner last night, we would have gone to an ethnic restaurant in North Carolina.

(We live in Chapel Hill, North Carolina; we ate at home last night; when we eat out, we tend to visit local ethnic restaurants.) Let us gradually turn the counterfactual supposition into a disjunction (~*p* or ~*q*). In the same context, the following conditional is true:

> Had my family gone out to dinner last night, we might have gone to Chinese Noodle Restaurant in Chapel Hill, but we would not have gone to McDonald's in Istanbul.

Hence (in the same context),

> Had my family eaten dinner last night either at Chinese Noodle Restaurant in Chapel Hill or at McDonald's in Istanbul, we would have gone to Chinese Noodle and not to Istanbul.

In this context, the fact that we did not go to Istanbul last night for dinner takes priority over the fact that we did not go to Chinese Noodle for dinner. Now consider a scenario even more remote than our going to McDonald's in Istanbul last night for dinner: our breaking some law of nature last night. If our going to Chinese Noodle is "closer" in the given context than our dining at McDonald's in Istanbul, then it is "closer" than something even more outlandish than our dining at McDonald's in Istanbul: our violating the laws. How does that qualify as more outlandish?[34] There is a variety of possibility (namely, natural possibility) such that our dining last night at Chinese Noodle was possible but our violating the laws was impossible—and anything possible is "closer" than everything impossible. In other words, it is possible (naturally) for my family to eat dinner at Chinese Noodle or to break the law of gravity—and whatever would have happened, under some possible circumstance, must qualify as possible. So in the given context, this counterfactual is true:

> Had my family either eaten dinner last night at Chinese Noodle Restaurant in Chapel Hill or broken the law of gravity, we would have eaten dinner last night at Chinese Noodle and not broken the law of gravity.

This argument was given for an arbitrary context where all of the above suppositions can be entertained; the same kind of argument could be given for any context. So under a supposition ($\sim p$ or $\sim q$) that pits a law's preservation against an accident's, the law takes priority in any context—since context is powerless to override the principle that anything possible is nearer than everything impossible, and there is a species of necessity associated with the natural laws.

These ideas about necessity and possibility will play prominent parts in the next chapter, where I will defend them further.

1.9. How Two Sub-nomically Stable Sets Must Be Related: Multiple Strata of Natural Laws

I have suggested that Λ is a sub-nomically stable set and that no non-maximal set of sub-nomic facts containing an accident is sub-nomically stable. Are there any other sub-nomically stable sets? There is at least one: the set containing exactly the broadly logical truths that are sub-nomic. For example, 3 would still have failed to divide 23 evenly even if I had been wearing an orange shirt, or even if gravity had declined with the cube of the distance—indeed, under any counterfactual supposition that is logically consistent with the broadly logical truths.

Furthermore, for any two sub-nomically stable sets, one must be a proper subset of the other, so the sub-nomically stable sets must fall into a natural hierarchy. Here is the proof (see fig. 1.2).

> Suppose (for the sake of *reductio*) that Γ (gamma) and Σ (sigma) are both sub-nomically stable sets, *t* is a member of Γ but not of Σ, and *s* is a member of Σ but not of Γ.
>
> Let's start with Γ. The claim (~*s* or ~*t*) is logically consistent with Γ. (Since Γ is stable, Γ contains every sub-nomic logical consequence of its members, so since Γ does not contain *s,* it follows that Γ does not entail *s,* and so ~*s* is logically consistent with Γ, and hence (~*s* or ~*t*) is, too.)
>
> Since Γ is sub-nomically stable, every member of Γ would still have been true, had (~*s* or ~*t*) been the case.
>
> In particular, *t* would still have been true.

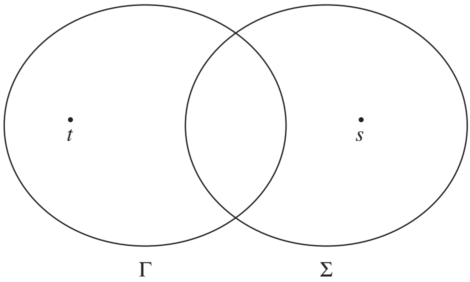

Figure 1.2 Sets Γ and Σ with their members *t* and *s*, respectively.

Thus t & ($\sim s$ or $\sim t$) would have held, had ($\sim s$ or $\sim t$).

Hence, ($\sim s$ or $\sim t$) $\Box\!\!\rightarrow \sim s$.[35]

Now let's work from the Σ side. Since ($\sim s$ or $\sim t$) is logically consistent with Σ, and Σ is sub-nomically stable, all of Σ's members are preserved under the supposition that ($\sim s$ or $\sim t$). It is not the case, for any of Σ's members, that its negation would (or even might) have held, had ($\sim s$ or $\sim t$).

Take s in particular: $\sim$ (($\sim s$ or $\sim t$) $\Box\!\!\rightarrow \sim s$).

But this result contradicts our earlier conclusion that ($\sim s$ or $\sim t$) $\Box\!\!\rightarrow \sim s$.

The promised *reductio* has been achieved: we have shown that for any two sub-nomically stable sets, one must be a proper subset of the other.[36] This result can guide our search for other sub-nomically stable sets besides Λ and the set of sub-nomic, broadly logical truths. For instance, this result (together with the stability of the set of sub-nomic, broadly logical truths) entails that any stable set containing some sub-nomic, broadly logical truths, along with some sub-nomic truths that are not broadly logical truths, must contain every sub-nomic, broadly logical truth. But any sub-nomically stable set must contain every sub-nomic, narrowly logical truth (since those truths are logical consequences of its members), and any narrowly logical truth is a broadly logical truth. Hence, any stable set containing some sub-nomic truths that are not broadly logical truths must contain every sub-nomic, broadly logical truth. Furthermore, as we saw in the previous section, no nonmaximal set containing an accident is stable. Thus, no nonmaximal superset of Λ is stable. Therefore, to find further promising candidates for stability, let's look among proper subsets of Λ containing all of the sub-nomic, broadly logical truths.

Many of these sets clearly lack sub-nomic stability. Consider Coulomb's law: between any two point bodies that have for a long while been at rest at any positions $\boldsymbol{r}_1$ and $\boldsymbol{r}_2$ ($\boldsymbol{r}_1 \neq \boldsymbol{r}_2$) while carrying any electric charges q_1 and q_2, there is a mutual electrostatic repulsion $\boldsymbol{F} = q_1 q_2 / |\boldsymbol{r}_1 - \boldsymbol{r}_2|^2$. Now consider a logical consequence of Coulomb's law, namely, that the above Coulombic regularity holds at all times *after today*. Take the proper subset of Λ containing exactly this restricted Coulombic regularity, the broadly logical sub-nomic truths, and their sub-nomic logical consequences. Is this set stable? Consider this counterfactual supposition: Had Coulomb's law been violated sometime *before* today. This supposition is logically consistent with the restricted Coulombic regularity, since that regularity concerns events only after today. Therefore, the chosen set is sub-nomically stable only if its members are all preserved under this supposition. But they are not. Had Coulomb's law been violated sometime before today, then Coulomb's "law" would not have been a law and so would not have been around to mandate the restricted Coulombic regularity. With Coulomb's law out of the way, there would have been nothing to keep the course of events after today from violating Coulomb's law. Therefore, it might well have been violated after today—just as energy might well have failed to be conserved, had there been no law making energy conservation compulsory.

However, *some* proper subsets of Λ containing all of the broadly logical sub-nomic truths *are* plausibly sub-nomically stable. Take the fundamental law of dynamics, which governs the relation between the forces on a body and the body's motion. At one time, this law was believed to be Newton's second law of motion $\boldsymbol{F} = m\boldsymbol{a}$, relating the net force $\boldsymbol{F}$ on a body to its mass m and acceleration $\boldsymbol{a}$. In 1830, George Biddell Airy used Newton's second law of motion to figure out how bodies would have behaved had they been subjected to various weird hypothetical kinds of forces.[37] His investigation presupposes that Newton's second law of motion would still have held, even if the force laws had been different. Similarly, Paul Ehrenfest in 1917 famously showed that had gravity been an inverse-cube force or fallen off with distance at any greater rate, then planets would eventually have collided with the sun or escaped from the sun's gravity.[38] Ehrenfest's argument also presumes that Newton's second law of motion would still have held, had gravity obeyed a different force law.

Plausibly, the fundamental dynamical law would still have held, had the world been populated by different kinds of forces or different kinds of fundamental particles, or had the strengths of those forces or the characteristic properties of those particles been different. Had the electrostatic force been half as strong, or light speed been half as fast, or the electron's charge been half as great, then the fundamental dynamical law would have been no different. Any additional kinds of forces and particles, had they existed, would have obeyed the same fundamental dynamical law as the actual kinds do. The fundamental dynamical law would still have held, had there been charged leptons other than muons, electrons, and taus—the actual species of charged leptons. (This counterfactual supposition violates the "closure law" belonging to Λ that all charged leptons are muons or electrons or taus. Closure laws will arise again in chapter 3.)

Likewise, consider the law of the composition of forces. It specifies how various component forces, of whatever kind, add (via the "parallelogram of forces") to yield a total resultant force—the "net force" appearing in Newton's second law of motion. The law of the composition of forces would still have held, even if the world had been populated by different kinds of forces.[39] By the same token, it is commonly suggested that the conservation laws transcend the particularities of the force laws. Even if there had been different kinds of forces (or, say, the electrostatic force had been twice as strong), momentum and energy would still have been conserved. Regarding energy conservation, momentum conservation, and so forth, Eugene Wigner writes: "[I]t is clear that their validity transcends that of any special theory—gravitational, electromagnetic, etc.—which are only loosely connected . . . ."[40]

Likewise, consider the Lorentz transformations, which are central to Einstein's special theory of relativity, entailing such famous relativistic results as time dilation and length contraction. They specify how an event's space-time coordinates in one reference frame in a certain family relate to its coordinates in that family's other frames. In his first relativity paper in 1905, Einstein derived these transformations by using the principle that light travels at the same rate in one of these frames whatever the motion of its source. However, as Einstein later wrote,[41] this derivation is misleading because features of light, a particular inhabitant of space-time, are not responsible for the coordinate transformations. The widespread recognition that the transformations lie deeper than the particular kinds of fundamental forces or particles there happen to be has led to a long tradition, beginning as early as 1909, of deriving the Lorentz transformations without appealing to any details of electromagnetism or any other force.[42] It is only because scientists believe that the Lorentz transformations transcend the fundamental force laws that they believe that were there additional fundamental force laws or had the fundamental force laws been different, the Lorentz transformations would still have held. As Roger Penrose says, if it is just "a 'fluke'" that certain dynamical laws exhibit Lorentz covariance, then "[t]here is no need to believe that this fluke should continue to hold when additional ingredients of physics are" discovered. But physicists generally do regard special relativity as prior to the force laws.[43] For instance, physicists commonly assert that had the force laws been different so that photons, gravitons, and other kinds of particles that actually possess zero mass instead possessed nonzero mass, the Lorentz transformations would still have held (though these particles would not have moved with the speed *c* that famously figures in the transformations).[44]

Thus, by using the concept of sub-nomic stability, we can cash out what it would be for the conservation laws, the parallelogram of forces, the fundamental dynamical law, and the coordinate transformations to rise above the specific kinds of forces there happen to be. It would be for a set containing laws such as these (and the sub-nomic, broadly logical truths), but omitting the force laws, to possess sub-nomic stability. Here we have a plausible candidate for a stable proper subset of Λ. In short, there appears to be a hierarchy of sub-nomically stable sets that includes at least these members (see fig. 1.3).

The pyramid in figure 1.3 suggests that there are at least two "strata" or "levels" of natural law. Any metaphysical account of what natural laws are should leave room for laws to come in multiple strata.

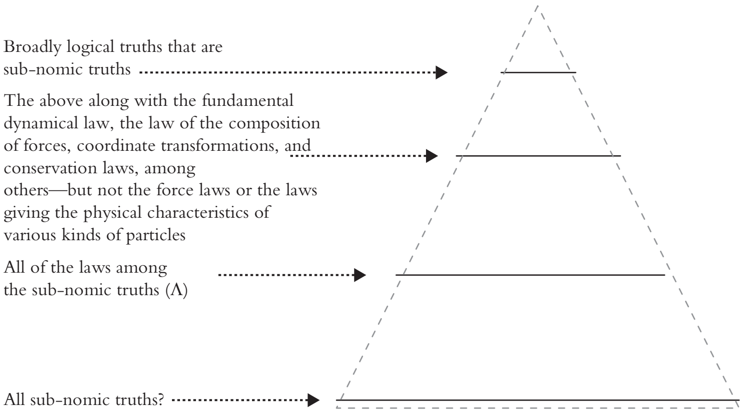

Figure 1.3 Some (though perhaps not all) good candidates for sub-nomically stable sets. This pyramid is *not* the hierarchy depicted in figure 1.1, which placed sub-nomic facts at the bottom, the laws governing *them* (first-order laws) one rung higher, the laws governing them (meta-laws) one rung higher, and so forth. Unlike that earlier hierarchy, every rung of this pyramid contains exclusively sub-nomic facts.

In later chapters, I shall impose this criterion of adequacy on familiar philosophical accounts of natural law and say more about the significance of the laws' coming in multiple strata.

1.10. Why the Laws Would Still Have Been Laws

I have argued that it is a law that *m* if and only if *m* belongs to at least one nonmaximal sub-nomically stable set (or equivalently—since the sub-nomically stable sets form a pyramid—if and only if *m* belongs to the largest nonmaximal sub-nomically stable set). I have also suggested that this biconditional captures the special relation in which laws stand to counterfactuals. It is a consequence of what lawhood is, not a peculiarity of what the actual laws are in fact like. If these suggestions are correct, then a philosophical analysis of what a natural law is should account for this relation between lawhood and stability. I will take up that challenge later in this book. For now, I shall note only that any analysis that succeeds in explaining why stability and lawhood are so related will automatically have a further payoff: it will thereby explain why the (first-order) laws would still have been *laws* under any sub-nomic counterfactual supposition *p* that is logically consistent with the laws. Here the nested counterfactuals in the definition of sub-nomic stability come into play.

Suppose that *m* is a member of Γ, a sub-nomically stable set, and that each of *q, r, s*... is individually logically consistent with Γ. Then Γ's stability ensures that $\sim (q \Diamond\!\!\rightarrow \sim m)$, $\sim (q \Diamond\!\!\rightarrow (r \Diamond\!\!\rightarrow \sim m))$, $\sim (q \Diamond\!\!\rightarrow (r \Diamond\!\!\rightarrow (s \Diamond\!\!\rightarrow \sim m))$, and so forth. One of the connections we saw (in section 1.5) between might-conditionals and would-conditionals was that $\sim (q \Diamond\!\!\rightarrow \sim m)$ logically entails $(q \Box\!\!\rightarrow m)$. So the counterfactuals in the above sequence respectively entail $(q \Box\!\!\rightarrow m)$, $(q \Box\!\!\rightarrow \sim (r \Diamond\!\!\rightarrow \sim m))$, $(q \Box\!\!\rightarrow \sim (r \Diamond\!\!\rightarrow (s \Diamond\!\!\rightarrow \sim m))$, and so forth. So had *q* been the case (that is, in the "closest *q*-world"), the following all hold: *m*, $\sim (r \Diamond\!\!\rightarrow \sim m)$, $\sim (r \Diamond\!\!\rightarrow (s \Diamond\!\!\rightarrow \sim m)$, and so forth—simply each of the earlier counterfactuals with their opening $q \Box\!\!\rightarrow$'s lopped off, since we are talking about what's true in the closest *q*-world. But this sequence supplies exactly what is needed for Γ to be sub-nomically stable in that world: *m* and its colleagues in Γ are all true and preserved under every counterfactual supposition that is logically consistent with Γ. Hence, if Γ is in fact sub-nomically stable, then Γ would still have been sub-nomically stable had *q* been the case, for any *q* that is logically consistent with Γ.

Therefore, for any such *q*, the actual laws would still have been *laws* had *q* been the case—if the laws under *q* are exactly the members of at least one set that would under *q* have been nonmaximal and sub-nomically stable. (Moreover, any *stratum* of laws would still have constituted a stratum of laws, had *q* been the case.) We thereby save the intuition that had Earth's axis of rotation been nearly aligned with its orbital plane (so that Earth was "lying on its side," as Uranus actually is), then although terrestrial seasons would have been quite different, the actual laws of nature would still have been laws—which is *why* terrestrial seasons would have been so different. The laws' collective invariance under counterfactual suppositions is no "accident"; the laws' invariance (and hence their lawhood) is itself invariant. This fact will resurface in the chapters that follow.

1.11. Conclusion: Laws Form Stable Sets

This chapter sets up everything else in the book. So let's review the main point: Λ is the largest nonmaximal sub-nomically stable set. Whereas NP uses the *laws* to pick out a range of counterfactual suppositions under which a set's invariance is to be tested, stability allows any set to pick out a comfortable range for itself. Thus, in using stability to explain how laws differ from accidents in having a special relation to counterfactuals, we avoid arbitrarily privileging the laws from the outset; we avoid specifying the laws as the truths that would still have held under every counterfactual supposition that is logically consistent with the laws. Although we could talk about the set of truths that are invariant under every supposition that is logically consistent with (say) "George Washington is the first president of the United States," these truths fail to form a stable set. By identifying the laws as the members of at least one nonmaximal stable set, we discover how a sub-nomic fact's lawhood is fixed by the sub-nomic facts *and the subjunctive facts about them*. The laws' stability turns out to account not only for the sharp distinction between laws and accidents, but also for the fact that the laws would still have been *laws* had *q* been the case, for any *q* that is logically consistent with the laws.

I began this chapter by suggesting that laws are set apart from accidents by their *necessity* (of a certain kind). Let's now see how the notion of stability helps us to understand the laws' characteristic necessity ("natural necessity"). On to chapter 2!