CHAPTER TWO

Laws, Mechanisms, and Models

LOOKING AT BIOLOGY from a philosophical point of view, one of the first things people notice is that there is apparently not much role for scientific *laws*. The image of science as a search for the laws governing the natural world is an old and influential one, and many philosophers have held that the investigation of laws is central to any genuine scientific field (Carnap 1966, Hempel 1966). The laws of physics may be basic, but each science tries to find its own laws—laws present in the systems it studies. Perhaps biology is just a cataloguing of the world's contents, and not a theoretical science that gives us real understanding?[1] The progress in biology over the past century has made this seem more and more unlikely. Instead, it appears that good science can be organized differently. Or perhaps laws are present in biology but we are not seeing them clearly and calling them by that name?

This chapter is about the organization of hypotheses and explanations in biology. I start with laws, and then look at two other sets of issues.

2.1. LAWS

What exactly is a law of nature? There is much disagreement, and I will focus on a few features that are widely accepted. First, a statement of a law is a true generalization that is *spatiotemporally unrestricted*; it applies to all of space and time. Second, a law does not describe how things merely *happen* to be, but (in some sense) how they *have* to be. An example of a law that seems to

[1] Ernest Rutherford, who split the atom, allegedly said, "All science is either physics or stamp collecting." See also Smart (1959).

meet these criteria is Einstein's principle that no signal can travel faster than light. The idea that laws describe how things must, or have to, be is sometimes expressed by saying that laws have a kind of *necessity*. That may seem an overly strong word, and some philosophers would avoid it. Laws are not supposed to have the same kind of necessity seen in mathematical or logical truths (such as "*p&q* implies *p*"). But even if the term "necessity" is not used, there is supposed to be a distinction between a natural law and an "accidental" regularity, a pattern that merely happens to hold. A standard example of an accidental regularity that might be true for all space and time is: *all spheres of gold have a diameter of less than one mile*. Contrast: *all spheres of uranium-235 have a diameter of less than one mile*. A sphere of uranium that big would explode, so this second regularity is one that *has* to hold.

Laws might be "strict"—admitting no exceptions—or they might involve probabilities. That divide will not matter much here. There is also a verbal ambiguity: sometimes the term "law" is used for a *statement* of one of these patterns in nature, and sometimes for the pattern itself. I will use the term for the patterns themselves.

A biological example that has been much discussed is "Mendel's First Law." This principle has been revised since the days of early genetics, and it has exceptions. But it is a good illustration of several aspects of the situation. In modern language, the principle says that in the formation of sex cells (eggs and sperm), a diploid organism (one with two sets of chromosomes, like us) puts one gene into each sex cell of the two genes that it received at that place in its genome from its own parents, and each of these two genes has a 50 percent probability of being found in any given sex cell. Exceptions include cases of Down syndrome in humans, and cases where particular genes have evolved the capacity to make their way into more than their fair share of sex cells (§6.3). But let's set those aside for now and treat the generalization as near enough to true in sexually reproducing species like ours. Is this a "law," or an "accidental" regularity?

The best initial answer seems to be "a bit of *both*." There is no reason to think that any sexually reproducing animal *must* do things this way. The genetic system we find in organisms like us evolved from something different, and might evolve into something else in the future (Beatty 1995). You might say that the generalization holds uniformly within organisms on earth over some relevant period, but laws are not supposed to be restricted to some places and times.

On the other hand, if we look at sexual organisms that are around now, it is *no accident* that the regularity holds. Once certain machinery is in place, this machinery has consequences, and these include the patterns described in Mendelian genetics. Mendel's First Law, to the extent that it holds, is a predictable result of the operation of mechanisms that are contingent historical products.

I'll put another couple of examples on the table. The "Central Dogma of Molecular Biology" describes the construction of new protein and nucleic acid molecules. It holds that the specification of the order of the building blocks in these molecules always goes *from* nucleic acid *to* protein, never vice versa and never from protein to protein (Crick 1958, 1970). (The Central Dogma is sometimes described as saying more than that, but I will stick with Crick's version.) "Kleiber's Law" describes the rates at which animals of different sizes use energy. The metabolic rate of an animal (R) depends on its body mass (M) and a constant (c) according to this formula: $R = cM^{3/4}$. Discovered in the 1930s, this relationship holds across a wide range of cases. For different groups of animals there is a slightly different c (so c is a "constant" only within each group) but the ¾ is always the same. For many years this was seen as a striking and mysterious relation, and then it turned out to be possible to derive Kleiber's Law from general features of the transport networks that move substances around the body, such as blood vessels, along with an assumption that efficiency in these networks is maximized. Given those assumptions, Kleiber's Law must hold (West et al. 1997).

Kleiber's Law initially seems independent of history, a manifestation of general facts about transport networks. But what about the assumption that these networks will be efficiently organized? Biologists differ on how unusual it is for evolution to produce inefficient or poorly adapted outcomes, but it is certainly possible.

In both the Mendel case and the Kleiber case, "law or accident?" seems to be the wrong question to ask. So let's start again. Rather than a two-way distinction between laws and accidental regularities, biological patterns show different amounts of what can be called *resilience* or *stability.* (Other terms used in this area are *robustness* and *invariance*.)[2] A resilient pattern is one that holds across many actual cases, and does so in a way that gives us reason to believe it would also hold in some relevant situations that are not actual but merely possible. A resilient pattern need not hold in *any* possible situation, and it might have some actual-world exceptions. Resilience is not a yes-or-no matter; regularities have more or less of it, roughly speaking, though there isn't a single scale on which all can be compared.

Mendel's First Law and Kleiber's Law have some degree of resilience, though in both cases it is clear how exceptions could arise. Let's look again at the Central Dogma. When we look at what the Dogma rules *out*, it looks quite resilient. An exception would be some process in which the order of the amino acids (the building blocks) in a protein molecule was used to determine the order of nucleotides in a molecule of DNA or RNA, or the amino acids in another protein. This is thought to be chemically difficult. We might be wrong in thinking this, and perhaps an actual-world exception will be found. But so far the Central Dogma looks pretty resilient. On the other hand, the thing that made the Dogma important—the thing that made it reasonable to use the term "Central"—was the idea that proteins are made by simply reading off the sequence of nucleotides in DNA. Complications to that picture, including the discovery of widespread "editing" of the RNA intermediate stages, have steadily grown (§6.1), and as they have grown the centrality of the Dogma has shrunk.

So far I have been discussing broad and well-known principles. Biology also has many narrower generalizations with some degree of resilience. In mammals, the sex of an individual is determined by its male parent (except perhaps in one enigmatic vole). Spiders

[2] I borrow this term from Skyrms (1980), one of the first to introduce an idea of this kind, but I use the term differently. See Woodward (2001) for differences within this family.

are carnivorous. (For years I used this as an example of a pattern without any exceptions, but now a vegetarian has been found [Meehan et al. 2009].) Some generalizations in biology describe the ways that biological properties are distributed among actual organisms. Others describe the causal consequences of a setup or interaction of factors, without saying where or how often this setup is found: a species that has lost almost all of its genetic diversity is likely to go extinct. That principle describes the consequences of low diversity without saying which species fall into that category.

Once these facts have all been laid out, you might decide to use the term "law" for all the patterns that have *some* resilience, you might reserve it for cases that have a great deal, or you might think the term should be dropped from biology. That is mostly a verbal choice.

Does this analysis apply to *all* of science, or just to biology and similar fields? Sandra Mitchell (2000) applies a view of this kind to all of science, including physics.[3] Another possibility is that physics is a special case; physics describes laws that govern the fundamental working of the world, and the working of these laws in organisms gives rise to further patterns that are not much like physical laws but have various degrees of resilience.

2.2. Mechanisms

I'll say more about laws later, but first I will look at some newer accounts of how theories work in biology. One family of views hold that large parts of biology are engaged in the *analysis of mechanisms*.[4] A mechanism is an arrangement of parts that produces a more complex set of effects in a whole system in a regular way. Biology describes how DNA replication works, how photosynthesis works, how the firing of one neuron makes another fire. In cases like these, the activities of the parts of a system are described, and

[3] As Mitchell notes about the standard example on the second page of this chapter, given the way gold comes to exist in a universe like ours it is *not* so "accidental" that huge amounts of it have never come together in one place.

[4] Central works here are Bechtel and Richardson (1993), Glennan (1996), and Machamer et al. (2000).

these activities and the relations between them explain how the more complicated capacities of the whole system arise.

To say this is not yet to break from a law-based view. Perhaps the way mechanisms are analyzed is by showing how the parts are governed by laws? But this approach is often seen a replacement for a law-based view. In the analysis of mechanisms, a different kind of causal understanding seems to be sought, or at least available. This is visible in the language used to describe causal relationships: one molecule will *bind* to another, altering how it interacts with other molecules. Or it might *cleave* or *oxidize* another. A stretch of DNA will be *transcribed*, or *silenced* by the *methylation* of some of its sequence. This seems to be a form of causal description oriented around the idea that some events *produce* others, in virtue of how things are physically connected. (I will discuss causal relations again, and modify this picture, in §6.2.) Generalizations expressed in these terms might still be seen as describing laws, but laws don't play an overt role in this sort of analysis. And although the parts of a neuron firing or embryo developing may well be following *physical* laws, that is in the background, and there seems to be no need to find laws of biology if you can describe all the mechanisms in this way.

This kind of work is "reductionist," in a low-key sense of that term: the properties of whole systems are explained in terms of the properties of their parts, and how those parts are put together. Reductionism is sometimes associated with the idea that a whole system is "nothing but" its lower-level parts, but this "nothing but" talk is usually quite misleading. A living system may be entirely composed out of a collection of parts, but the system will have features that none of the parts have. Rather than showing that the higher-level activities do not exist, the point of mechanistic explanation is usually showing *how* the higher-level features arise from the parts.

This view gives a good account of at least part of biology. How far does it extend? One option is to extend it very broadly. Perhaps natural selection, for example, is a mechanism in this sense, and evolutionary biology is about the analysis of mechanisms?[5]

[5] See Skipper and Millstein (2005). The next few pages have been influenced by Levy (2013) and Matthewson and Calcott (2011).

Perhaps the exclusion of one species by another in an ecological system is also a mechanism? I think saying those things requires diluting the sense of "mechanism" that has been useful in the analyses sketched above. Instead, there is a side of biology that analyzes mechanisms and a side that does not.

The philosophers who argue for the importance of mechanistic analysis do not tie their view to a 17th-century sense of the term "mechanism," in which the universe is treated as if it were clockwork. Mechanistic views of that kind, which see the world as governed only by pushes, pulls, and collisions, have been rejected in basic physics. But the biological systems to which mechanistic analysis most directly applies do have a machine-like quality in another sense. They are not only physical systems, but *organized* ones. This is another vague term, but a way to make sense of it is to think about how sensitive a system is to small changes in its parts, especially substitutions of one part for another. If we look at a neuron firing or a protein being made, the process we are interested in occurs as a result of the interactions of parts whose exact relations to each other matter. If you swapped a chromosome for a ribosome, the consequences would usually be large, just as they would be if you swapped a car's fuel pump for its gear box. One part of biology is concerned with systems like this. Other areas are concerned with systems that are "looser." When a population of organisms is evolving, there are parts (the organisms) and a whole (the population), but many of the parts are similar to each other, and if you swapped one for another it often would not make that much difference. Their exact relationships—who is next to whom—do not matter so much. These relationships *might* matter greatly in a particular case, but often they do not. A gas, as studied in physics, is a more extreme case of the same thing. A gas contains many molecules moving about in specific ways, but the details of those ways do not matter to various important properties of the gas, like its temperature and pressure. Those properties depend on broad and general features of the collection, such as the molecules' average speed. If you swap one molecule for another, it usually won't matter. When analyzing a system of that kind, a statistical approach is often taken.

Fields like evolutionary biology, ecology, and epidemiology are concerned with systems of this second kind—or more exactly, not with *un*organized systems but with *less* organized systems. The systems they study are more organized than a gas, but less organized than a cell.

Adapting some terminology used by Richard Levins (1970) and William Wimsatt (2006, 2007), we can distinguish more *organized* systems from more *aggregative* ones. More organized systems include cells and organisms, and more aggregative ones include populations of those organisms. There are intermediate cases, like honey bee colonies. Sometimes if you "zoom out" from an aggregative system, organization will reappear. If you imagine watching gas molecules interacting with blood cells in the lungs of a large animal, what you see will be an aggregative system. If you swap one oxygen molecule for another, it does not make much difference. But if you zoom out so that the lungs become one organ in a whole body, *that* system is a highly organized one. Mechanistic analysis is most appropriate when dealing with organized systems. Aggregative ones are better described in terms of tendencies that arise from the combined action of parts which each have some degree of independence. The two kinds of systems occur at different scales; there are size constraints, it seems, on highly organized systems, due to the difficulty of keeping the parts working together, and organized systems often have distinctive kinds of histories. Organized systems often make use of aggregative activities in their small parts (consider molecules diffusing across a membrane). It is interesting to think about human societies, which can be very large objects, in terms of this distinction.

This distinction can also be used to clear up, or perhaps replace, another. Earlier I mentioned the idea of "reduction." A term often used to express a contrast with features that can be reductively explained is "emergent." Emergent properties are sometimes said to be those that *can't be explained* at a lower level; they are "irreducible." In the philosophy of mind, consciousness is sometimes said to be an emergent property in this sense. The claim is that although consciousness has some material basis in the brain, it can't be explained at the neural level. You could know exactly what all the neurons were doing, and it would still be mysterious why those brain processes gave rise to consciousness.

The term "emergent" is also used in much weaker (more inclusive) ways. Biologists sometimes use it to refer to properties of a whole system that the system's individual parts do not have. The high-level properties might be *explained* in terms of the parts, but are not *present* at the lower level. An example is the "surface tension" phenomenon in water. Surface tension is a consequence of the tendency of water molecules at an air/water boundary to form lots of weak chemical bonds with each other rather than with the air. An individual water molecule does not have surface tension; the phenomenon exists only when many molecules are brought together. This is a sense of "emergent" in which most features of any complex system will qualify.

The underlying phenomenon here is, once again, something like a gradient: higher-level activities in a system can be more or less dependent on the exact relations between the parts. If you want to draw a line between the "emergent" and the "reducible" properties, you could draw it at the divide between cases where higher-level properties are also *present* in the parts and cases where they are not, but then emergent properties are often clearly explainable in a bottom-up way. The distinctions beyond that one are distinctions of degree. How sensitive is a high-level behavior—the music coming from an orchestra, the economic patterns coming from the choices of individuals, the behavior coming from a collection of human cells—to the arrangement of the system's parts, in addition to the parts' individual properties? The idea of a special category of emergent properties that cannot be explained at all in lower-level terms has been influenced by the special perplexities of the mind/body problem. It probably does not help there, and there is no support for such a picture in other parts of biology.

2.3. Models

Of the other styles of work in biology, one is especially relevant here. This is modeling, or model building.

"Model" has many meanings in science and philosophy. Sometimes the term is used to describe any theory or hypothesis, or to describe a theory that is acknowledged to be rough or simplified.

However, the word can also be used to indicate a particular strategy in scientific work, a strategy in which one system is used as some kind of surrogate for another. The usual reason for doing this is that the "target" system we want to understand is too complicated to investigate directly. So it makes sense to choose some of the most important factors operating in the target system, and work out how they interact in a situation in which other factors are absent. Alternatively, one system might be used as a model for another because the best available methods can be more easily applied to the model than to the target, even though the model system is no simpler.[6]

In some cases a "model system" will be a physically built object. Engineers still build scale models of river systems and bays. This is related to the use of "model organisms" in biology. Model organisms, such as fruit flies and *E. coli* bacteria, were initially naturally occurring organisms that were easy to work with in the lab. Now they are often partly artificial, with features that would never occur in the wild. Much modeling work is different from this, however, as there is no model system that is physically present. Instead the model system is imagined or hypothetical. A researcher will write down a set of assumptions that are relevantly similar to those that hold in some real system, and will use mathematical analysis, computer simulation, or some other method to work out the consequences of those assumptions.

Evolutionary game theory is an example of a field where this method is widespread. Game theory uses mathematics to study how rational agents should behave in relation to each other. In the 1970s George Price and John Maynard Smith pioneered the use of this method to deal with animal behavior.[7] Rather than assuming that animals are rational, they assumed that natural selection will lead to the proliferation of behaviors that promote survival and reproduction, eliminating behaviors that do not. The first application of these methods was to fighting; the aim was to work out why bluffing and ritualized non-damaging fights are so common in animals. Here is one result. Suppose we have a popu-

[6] See Giere (1988), Godfrey-Smith (2006), Weisberg (2007b, 2013).
[7] Maynard Smith and Price (1973), Maynard Smith (1982).

lation in which individuals meet at random, one on one, and fight over resources. The population contains two kinds of individuals, "hawks" who will fight until they win or are seriously injured, and "doves" who bluff initially but retreat if things get out of hand. Individuals who do well in these contests are assumed to reproduce more than those who do not, and to pass on their behavioral type to their offspring. What will happen in such a population? If the cost of injury from losing a hawk-on-hawk fight is high in relation to the value of the resource, and some other assumptions are met, the population will reach a stable state where it contains a mixture of both strategies. Each type does well when it is rare. When hawks are rare they exploit the doves; when doves are rare they are the only ones avoiding damaging fights.

I said, "What will happen in such a population?" But the first thing to note is that natural populations are never as simple as this—there are no populations with exactly two behavioral types, where all the "hawks" are behaviorally equivalent to each other, and so on. Even theories that have a less obvious role for imagined scenarios often have some of this character; many evolutionary models assume populations that are effectively infinite, that deal with a uniform environment, and have unrealistically simple genetics. I will describe all models that make use of deliberate simplifications as *idealized*. Idealization can be contrasted with *abstraction*, which does not involve imagining things to be simpler than they are, but merely leaving some factors out of a description. Abstraction, to some degree, is inevitable; you can't include everything. Idealization, in contrast, is a choice. The border between these two is not always obvious, though, and will be important to some issues discussed later in this book.

It is hard to work out, philosophically, how to think about modeling of a kind that seems to involve the investigation of imaginary systems. One approach is to treat it as analogous to work that uses scale models: sometimes a model system is built, and sometimes it is just imagined. This takes us into problems about fictions and possibilities. A good model system is similar to its target; how can a target be similar to something that does not exist? A different approach is to see a "model system" as an abstract mathematical object. One way or another, any analysis of modeling has to grapple with the importance of consideration of the merely hypothetical or possible. As R. A. Fisher, who developed some of the most influential models of evolutionary change, put it many years ago (1930, p. ix), "The ordinary mathematical procedure in dealing with any actual problem is, after abstracting what are believed to be the essential elements of the problem, to consider it as one of a system of possibilities infinitely wider than the actual. . . ."

The approach I will take is to set aside some questions about what models are, and focus on the *products* of this work. The usual product of a piece of scientific model building is a set of *conditional* statements, statements of "if . . . then . . ." form. Conditionals raise philosophical problems of their own (Bennett 2003), but I am going to take them for granted. In modeling, the "if . . ." can be freely invented—modelers can explore any scenario they like. But the choice is usually guided by two goals. First, the scenario should be one where it is possible to work out, in some rigorous way, what would happen if it obtained. The obvious way to do this is to make the scenario one whose consequences can be investigated by mathematical analysis, or by programming a computer. The second goal, which can pull against the first, is that the scenario specified should be one that is usefully close to the real world.

In making the transition from "if" to "then," computer simulation has become more and more important. People sometimes describe model systems, such as evolving populations or predator-prey interactions, as being "inside" computer simulations. Rather than trying to make sense of that kind of claim, the way to understand simulations of this kind is to see computers as aids to the rigorous use of the scientific imagination. A computer is a physical device whose operation can be exploited to trace out very complex networks of "if . . . then . . ." relationships. A modeler will specify a setup, some relevant configuration of organisms or cells or something else, and then look for a way to determine the consequences of the setup. Computers are useful because our ability to specify these setups outruns our ability to work out how they would behave. Regularities in the operation of the computer can be used to tell us the consequences of the scenario that has been imagined. That this is the role of computers is illustrated by the way modelers move freely back and forth between "analytic" methods (solving equations) and simulations.

Whatever method is used, the typical result is a claim of the form, "*If* there are one-on-one contests over resources and these further conditions are met . . . , *then* the population will come to contain a stable mixture of hawk and dove strategies." Given that a modeler has to start by making deliberate simplifications, there are two ways to try to give conditionals as much relevance to the actual world as possible. One is to minimize the departure from reality on the "if" side, thereby retaining some of the problem of the real world's complexity. The other is to start further away, but also look for ways to then make the "if" side as logically *weak* as possible—that is, as undemanding or easy to satisfy as possible. A good way to do this is to develop many variants of a model, each of which makes different assumptions—all the variants are idealized, but in different ways. If things go well, many variants will lead to the same outcome. In the best-case scenario, the modeler starts out with assumptions that involve significant departures from reality, but is then able to make the "if" side of the conditional so undemanding that the actual world is one of the ones that satisfies it, or is very close to satisfying it.

An example is "Volterra's Principle," which says that in a system with a predator and prey population, if some external factor is introduced that kills them both, such as a pesticide, this will increase the relative abundance of the prey population (Wilson and Bossert 1971, Weisberg and Reisman 2008). Volterra started out making a lot of deliberate simplifications (Kingsland 1995), but many would say that the resulting principle (more carefully formulated than above) is a true generalization about actual systems, one that explains why the application of pesticides in agriculture often makes problems worse, as the pesticide does more harm to the natural enemies of a pest than to the pest itself. Some think that many conditionals in biology, especially ones like this, include a tacit clause requiring *ceteris paribus*, which means "with other things equal." The idea is that intrusions into the system from outside, and freak events, are set aside as irrelevant.

One attitude in this area holds that modeling always aims to bring us eventually to a description that does not idealize.

Ideally, there would be no idealization. Another view is that even when all the details can be known, idealized models are useful because they can highlight similarities between different systems. Richard Levins, yet another influential modeler, argued that science will always make use of models that simplify, and will retain several models of any given system, as a result of facts both about nature and about ourselves:

> The multiplicity of models is imposed by the contradictory demands of a complex, heterogeneous nature and a mind that can only cope with few variables at a time; by the contradictory desiderata of generality, realism, and precision; by the need to understand and also to control; even by the aesthetic standards which emphasize the stark simplicity and power of a general theorem as against the richness and the diversity of living nature. These conflicts are irreconcilable. (1966, p. 431)

Models that apply to particular cases with great precision are good, and so are models that cover a wide range of cases. Pursuing one of these goals usually requires sacrificing the other. Simplicity is good, too, and simple models can sometimes be applied to a wide range of real systems—but only if the model is interpreted as fitting these real systems in a loose or approximate way.[8]

Suppose it is agreed that conditionals are the typical results of modeling work. Perhaps *these* are the "laws of biology"? They are generalizations, not restricted in space and time, and when the connection between antecedent ("if . . .") and consequent ("then . . .") is established mathematically, they surely have a high degree of resilience—perhaps even necessity.

One possible disanalogy between these conditionals and laws has been discussed by Elliott Sober (1993, 1997). He thinks that biology does have laws, uncovered by modeling, but these laws are not empirical. They are just pieces of mathematics, and hence are necessarily true. Laws of nature are usually seen as having

[8]For detailed discussion of these trade-offs, see Matthewson and Weisberg (2009).

empirical content, but Sober thinks we should get used to the idea that laws can be purely mathematical. However, I do not agree that the conditional statements we get from models are mathematically necessary. Modelers might use mathematics to work out what follows from a set of assumptions, but the conditionals they end up with do not have purely mathematical content. Compare: "7 + 5 = 12" is mathematically necessary, but "if you put seven marbles on a table and add five there will be twelve marbles on the table" is not mathematically necessary. Whether this is true depends on the physical characteristics of marbles and tables. The same applies to conditionals about what will happen in an ecological system where a certain kind of predator eats a certain kind of prey. The mathematics is often where the hard work is done, but the conditionals that result are not merely mathematical statements. They are statements about the behavior of organisms, populations, and other biological objects.

Something that does look like an important difference between many of these conditionals and laws in a traditional sense is that the antecedent often describes a situation that does not actually occur. It is only close to something that occurs. Laws have traditionally been seen by philosophers as applying more directly to real systems—having antecedents that are often literally true—rather than merely making claims about what *would* happen in a nonactual scenario. In some cases, like Volterra's Principle, this might not be much of an issue, but many conditionals derived from models do have an idealized character. At this point, though, it is worth casting a more critical eye on the assumptions being made about laws in sciences like physics. Some think that idealization is so pervasive that theoretical "laws" in physics rarely or never describe the behaviors of actual objects (Cartwright 1983, Giere 1999). To the extent that this is true, the apparent contrast with biology fades.

Putting things together, there are two kinds of generalizations in biology that both look a bit like laws. First, there are conditional statements derived from models—these are not beholden to historical contingency, are often very abstract, and tend to idealize to some extent. Second, there are general statements about what actual organisms are like—spiders are carnivorous, Mendel's First Law—that depend on historical contingencies and usually have exceptions.[9] Philosophers have often seen natural laws as independent of historical contingency *and* applying directly to real systems *and* highly general *and* having a kind of necessity. One possible position is that physics does have laws with this remarkable combination of properties, and biology does not. If so, perhaps the difference between these sciences is permanent, a consequence of the subject matter, or perhaps the gap will close. Another possibility is that laws in that sense are not found anywhere in science.

Setting laws aside, I will mention one other feature of model building before moving on. The aim of a modeler is often to come up with something whose departures from the real world do not matter too much: the "if . . ." is close enough to reality for the "then . . ." to be something we can expect to actually happen, at least approximately. Following the "internal logic" of a hypothetical scenario may become a goal in itself, however. This can lead to great theoretical creativity, but also to problems. After the financial crisis of 2008, the biggest crisis for banking and commerce since the Great Depression, some writers argued that economics had failed to predict and prevent the problem because it had become obsessed with the development of idealized models and had lost contact with reality. Paul Krugman, who had earlier won a Nobel Prize in economics, argued that the economics profession went astray "because economists, as a group, mistook beauty, clad in impressive-looking mathematics, for truth" (2009). Clarifying

[9] Within logic, generalizations about actual cases are also usually seen as conditionals: *if something is a spider, then it is carnivorous*. It is common then to distinguish several kinds of conditionals. A *material* conditional merely describes the layout of the actual world: *if something is a spider, then it is carnivorous* is true so long as there are no noncarnivorous spiders. This could be because there are no spiders at all. A *subjunctive* conditional asserts a connection between the two properties that goes beyond—somehow—the facts about which things actually exist. They can be expressed in a way that emphasizes this by saying, *if something were to be a spider, then it would be carnivorous*. There are also other kinds of conditionals, and the relations among them all are controversial: see Bennett (2003) and Edgington (2008). Here I will work within the idea that there are two (or more) kinds of scientific generalization, described in the text, without committing to an analysis of how they look from the point of view of logic.

this, the economists were probably making lots of true "if . . . then . . ." claims about markets and finance, using their high-powered mathematics, but the "ifs" were further from reality than they realized, and so were the "thens."

Further reading

On laws, Armstrong (1985), Carroll (2004); in biology, Turchin (2001), Waters (1998), Ginzburg and Colyvan (2004), McShea and Brandon (2010); on mechanistic explanation, Craver (2009), Glennan (2002) and see note 4; on emergence, McLaughlin (1992), Bedau (1997), Bedau and Humphreys (2008); on idealization, Weisberg (2007a); on models, Wimsatt (2007), Frigg (2010), Downes (2011), Toon (2012), and see note 6.